



Federal Office
for Information Security

Status of quantum computer development

Entwicklungsstand Quantencomputer



Document history

Version	Date	Editor	Description
1.0	May 2018		Document status after main phase of project
1.1	July 2019		First update containing both new material and improved readability, details summarized in chapters 1.6 and 2.11
1.2	June 2020		Second update containing new algorithmic developments, details summarized in chapters 1.6.2 and 2.12

Introduction

This study discusses the current (Fall 2017, update early 2019, second update early 2020) state of affairs in the physical implementation of quantum computing as well as algorithms to be run on them, focused on applications in cryptanalysis. It is supposed to be an orientation to scientists with a connection to one of the fields involved—mathematicians, computer scientists. These will find the treatment of their own field slightly superficial but benefit from the discussion in the other sections. The executive summary as well as the introduction and conclusions to each chapter provide actionable information to decision makers.

The text is separated into multiple parts that are related (but not identical) to previous work packages of this project.

Authors

Frank K. Wilhelm, Saarland University

Rainer Steinwandt, Florida Atlantic University, USA

Brandon Langenberg, Florida Atlantic University, USA

Per J. Liebermann, Saarland University

Anette Messinger, Saarland University

Peter K. Schuhmacher, Saarland University

Aditi Misra-Spieldenner, Saarland University

Copyright

The study including all its parts are copyrighted by the BSI–Federal Office for Information Security. Any use outside the limits defined by the copyright law without approval by the BSI is not permitted and punishable. This covers reproduction, translation, micro filming, and storing and processing in electronic systems.

BSI-Reference

BSI Title: Entwicklungsstand Quantencomputer

BSI Project Number: 283

Table of Contents

Document history.....	2
Introduction.....	3
Authors.....	3
Copyright.....	3
BSI-Reference.....	3
Index of Figures.....	9
Index of Tables.....	13
1 Deutsche Zusammenfassung.....	15
1.1 Was ist ein Quantencomputer?.....	15
1.2 Die Relevanz von Quantencomputern für die Kryptoanalyse.....	15
1.3 Hardware für Quantencomputer.....	16
1.4 Aktuelle Entwicklungen.....	19
1.5 Extrapolation.....	19
1.6 Aktualisierungen.....	20
2 Synopsis.....	23
2.1 Basic idea.....	23
2.2 Hardware platforms.....	23
2.3 Algorithmic goals.....	24
2.4 Computational Models.....	24
2.5 Evaluation along computational models.....	26
2.6 Evaluation of platforms.....	27
2.7 Extrapolation.....	31
2.8 Alternative platforms.....	33
2.9 Global activities and potential for development.....	33
2.10 Risks.....	33
2.11 Update Spring 2019.....	34
2.12 Update Spring 2020.....	34
Part I: Evaluating and benchmarking quantum computing hardware.....	36
3 Introduction to computational models and their requirements.....	37
3.1 Structure and requirements of an evaluation system.....	37
3.2 Basic principles and notions of evaluating fault tolerant gate-based quantum computing.....	38
4 Are recent quantum supremacy proposals relevant for cryptanalysis?.....	41
4.1 Fault tolerance vs. near-term quantum supremacy.....	41
4.2 Algorithmic perspective.....	41
5 Low-level analysis of qubit systems.....	44
5.1 Initial remarks.....	44
5.2 The DiVincenzo criteria reviewed.....	44
5.3 Coherence time scales.....	46
5.4 Qubit definition indicators.....	49

5.5	Qubit initialization indicators.....	50
5.6	Readout indicators.....	50
5.7	Final remarks.....	52
6	Benchmarking qubits.....	53
6.1	Introduction.....	53
6.2	Benchmarking and error mitigation techniques.....	53
6.3	Qualitative criteria beyond DiVincenzo.....	53
6.4	Benchmarking operations.....	54
6.5	Phenomenology of quantum errors.....	58
6.6	Application to the Google quantum supremacy experiment.....	60
7	Connecting physical and logical qubits with fault tolerant quantum computing.....	62
7.1	General observations on the role of fault tolerance.....	62
7.2	Introduction to surface code quantum error correction.....	67
7.3	Basic requirements.....	80
7.4	Performance discussion.....	81
7.5	Operational conclusions.....	87
7.6	Groups.....	89
7.7	Glossary for error correction.....	90
8	Benchmarking and fault-tolerance on non-standard architectures.....	93
8.1	Quantum annealing.....	93
8.2	One-way quantum computing.....	99
8.3	Oscillator encodings.....	102
9	Algorithmic innovations.....	105
9.1	Quantum circuit versus adiabatic quantum computation model.....	105
9.2	Quantum circuits without error correction?.....	106
9.3	Quantum circuits for cryptanalysis—Key players.....	108
9.4	Minimizing quantum circuits.....	109
9.5	Algorithmic innovations with relevance for symmetric cryptography.....	109
9.6	Algorithmic innovations with relevance for asymmetric cryptography.....	115
9.7	The cost of factoring and computing discrete logarithms in the Clifford+ T model.....	120
9.8	Translating algorithmic gate counts into fault-tolerant building blocks.....	123
10	Risks of our evaluation scheme.....	126
10.1	Risks that make quantum computers more reachable.....	126
10.2	Risks that make quantum computers less reachable.....	127
11	Global operational criteria for quantum computers.....	128
11.1	Extensive parameters.....	128
11.2	Critical parameters.....	130
11.3	Further descriptors.....	131
11.4	Addressing these operational challenges.....	131
11.5	Articulated architectural extrapolations.....	131
	Part II: Platforms.....	133

12	Quantum technology and computing platforms.....	134
13	Josephson qubits.....	136
13.1	Description.....	136
13.2	Evaluation: Planar transmon.....	146
13.3	Evaluation: 3D transmon.....	150
13.4	Evaluation: Flux qubit.....	152
13.5	Evaluation: Fluxonium.....	153
14	Ion traps.....	155
14.1	Description.....	155
14.2	Evaluation.....	160
15	Semiconductor platforms.....	164
15.1	Description.....	164
15.2	Evaluation: Spins in quantum dots.....	173
15.3	Evaluation: Other quantum dot platforms.....	175
15.4	Evaluation: Single dopants in Si/SiGe.....	177
15.5	Evaluation: NV-Centers.....	177
15.6	Evaluation: Topological qubits—Majorana fermions.....	178
15.7	Evaluation: Topological qubits—Other candidates.....	181
16	Neutral atoms.....	183
16.1	Description.....	183
16.2	Evaluation: Rydberg atoms.....	185
17	Photonic qubits.....	187
17.1	Description.....	187
17.2	Evaluation: Single photons.....	189
17.3	Evaluation: Continuous-variable and Gaussian encodings.....	191
18	Molecular approaches.....	194
18.1	Description.....	194
18.2	Evaluation: Molecular approaches.....	197
19	Exotic qubit candidates.....	199
19.1	Electrons on helium.....	199
19.2	Quantum nano- and optomechanics.....	199
19.3	Research groups.....	199
20	Big actors and programs.....	200
20.1	National and international programs.....	200
20.2	Academic institutions.....	202
20.3	Companies.....	207
21	Global conclusions.....	208
21.1	Update 2019.....	208
21.2	Update 2020.....	208
	Appendix.....	211

22	Details on error models.....	212
22.1	Treatment of random Pauli errors.....	212
22.2	Treatment of unitary errors.....	213
22.3	Mapping between small ubiquitous errors and error syndrome detection.....	215
22.4	Error correction experiments.....	217
22.5	Systematic unitary errors.....	217
22.6	Unrecoverable errors.....	218
22.7	Decoherence and entanglement with an environment.....	219
23	Fault-tolerance calculation.....	221
23.1	Errors from deviations in the measurement and initialization basis.....	221
24	Low-level error mitigation.....	223
24.1	Dynamical decoupling.....	223
24.2	Decoherence-free subspaces.....	223
25	Practicalities of fault tolerance.....	225
25.1	Fault-tolerant implementation of a small rotation: A practical example.....	225
25.2	Stabilizer measurement pattern for logical operations.....	225
26	List of records.....	229
	Table 26.4: Records for neutral atoms.....	230
26.1	Abbreviations.....	231
	Reference Documentation.....	232
	Keywords and Abbreviations.....	267

Index of Figures

- Figure 2.1: Levels of development towards a fault-tolerant quantum computer.....26
- Figure 2.2: Evaluation of the main platforms following the developed scheme.....27
- Figure 2.3: The currently leading quantum computing platforms - microscopic perspective. Left. Josephson processor (image: Julian Kelly, Google); Linear array of 9 qubits (crosses) with nearest-neighbor coupling; bottom: control lines, top: readout lines; Right. Linear ion traps (image Jürgen Eschner, Saarland University): trap electrodes (rods) and lenses for laser irradiation to implement quantum logic; same setup with its vacuum apparatus..... 28
- Figure 2.4: Comparison of algorithmic demands with currently achieved hardware performance. The plot shows required resources as number of qubits times rounds of error correction in the surface code for dlog (blue) and factoring (orange) for common key sizes (details in Chapter 9.5 and Figure 9.1) as a function of the physical error rate p . The squares show current realizations assuming one day run time (solid) or 100 days (empty), the yellow area shows expected near-term progress. Both scales are logarithmic..... 31
- Figure 2.5: Infrastructure units for quantum computers in leading platforms. Left: Dilution cryostat optimized for large cooling power and large wire-count for the operation of Josephson qubits (opened); qubits and other electronics units are mounted on the copper plates on different temperatures, the rack on the left contains control electronics. (Image: Edward Leonard jr. , University of Wisconsin-Madison); Right: Parts of a vibration-controlled optical table containing two vacuum chambers for separate ion traps (Image: Jürgen Eschner, Saarland University).....32
- Figure 3.1: sketch of interdependencies and connections of our evaluation scheme. Hardware needs to pass checkpoints from below, software is compiled from above.....37
- Figure 7.1: Schematic diagram of fault-tolerant quantum computing/storage. Boxes with S represent one cycle of syndrome measurement, in which all stabilizers relevant for the detection of errors are measured. The syndrome measurement consists of several physical operations on the bare physical qubits, which are necessary to get an effective multi-qubit measurement operator. Top: Protection of logical quantum information can be ensured by successive syndrome measurement rounds. The measurement outcomes of each round (double-line arrows) are collected in a classical memory and evaluated at the end of the process to identify and correct the effect of all errors at once. Middle: (Most) Clifford gates can be performed on the logical qubit states by modifying the syndrome measurements. This can be done by adding additional gates on the underlying physical qubits before the usual syndrome measurement circuit (as is the case for simple encodings like the Steane code) or by changing the set of operators that is measured, *i.e.*, including more stabilizer measurements or not measuring some of them. The surface code uses both modification types, depending on the logical gate that shall be performed, see Section 7.2.4. One logical gate can require several rounds of (modified) syndrome measurement rounds. Bottom: Non-Clifford gates like the T gate can not be performed directly with this method. The most common way is to use a (logical) magic state qubit which is produced in a different part of the circuit by state distillation (see Section 7.2.4.3) and combine it using Clifford gates. However, this is a non-deterministic process that needs classical feed-forward (blue box) of intermediate measurement outcomes to correct for the probabilistic effects.....65
- Figure 7.2: (a) Arrangement of physical qubits for the surface code. Data qubits are shown as open circles, measurement qubits as solid circles. The green and yellow crosses denote Z and X stabilizer measurements of the data qubits at the ends of the cross, respectively. At the boundaries, the stabilizer measurements include only three data qubits, represented by truncated crosses. (b) Circuit diagram for the Z stabilizer measurement. Identities are included to compensate for the Hadamards in the (c) X stabilizer measurement. Each step is performed simultaneously for all stabilizers. One round of such circuits for all Z and X stabilizers along the array corresponds to one syndrome measurement Box as

shown in Figure 7.1. Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society." 68

Figure 7.3: Three-dimensional space-time lattice of syndrome measurement outcomes. One horizontal layer corresponds to one round of syndrome measurement, where the signs indicate the outcomes. Red lines show where a change of measurement outcome occurs. A single error (X or Z) of a data qubit leads to a neighboring pair of sign changes in a spatial dimension—with the faulty data qubit lying in the middle, a single error on the measurement qubit leads to a pair in temporal dimension—with the error happening between the two changes (M). Error chains lead to pairs of sign changes lying further apart [FMMC12]. "Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society." 71

Figure 7.4: Figure 7.4: Implementation of logical qubits: (a) Double Z-cut qubit, (b) double X-cut qubit. The logical operators XL (ZL) consist of X (Z) operations on the physical qubits along the blue (red) lines [FMMC12]. Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society..... 72

Figure 7.5: Schematic protocol for creating and initializing a double X-cut qubit in a logical Z eigenstate. MZ denotes measurements in the basis of Z, $|g\rangle$ denotes initialization of the data qubits in the ground state [FMMC12]. "Reprinted figure with permission from. Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society..... 73

Figure 7.6: (a) Circuit diagram for a logical CNOT operation between two double Z-cut qubits, mediated by a double X-cut qubit. During the process, the target qubit is measured, and a new double Z-cut qubit is initialized in $|+\rangle$ to take the place of the target qubit. (b) Description of the braiding of holes that is done to perform the three CNOT steps: Every double Z(X)-cut qubit is represented by a pair of black (blue) lines, where the movement of the holes in time is shown along the x-axis. Two lines corresponding to two holes of the same qubit join when the qubit is initialized or measured. (c) Simplified representation of the braiding, showing the double X-cut qubit only as an intermediate tool for the gate. In fact, the double Z-cut qubits do not need to be moved at all and the new target qubit can be initialized at the position of the measured old one. (d)-(f) Equivalent representations for an indirect CNOT between two double X-cut qubits. [FMMC12] Reprinted figure with permission from. Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society..... 74

Figure 7.7: Implementation of S (top) and T (bottom) gate on the input state $|\psi\rangle$ with magic states $|Y\rangle$ and $|A\rangle$, respectively. In a more recent version, the S gate can also be performed without the final Hadamard gate, carrying a byproduct operator in the classical control [GF17]. The T gate additionally needs a conditional S gate to correct its non-deterministic nature. The classical process of deciding whether or not to perform the additional S gate after measuring MZ is represented by double lines. When the S gate is needed, the final state will be $XZT|\psi\rangle$, but the X and Z byproducts can be carried in the classical control. Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society..... 76

Figure 7.8: Performance below threshold for the surface code for distances 3,5,7,9,11,15,25,35,45 and 55. For distances 3,5 and 7, quadratic, cubic and quartic fit curves are shown as dashed lines. They only approximate the actual curves for low physical error rates p [FDJ13]. Reprinted by permission from Macmillan Publishers Ltd: Scientific Reports (A. G. Fowler, S. J. Devitt, and C. Jones, Sci. Rep., 3(1), 2013.), copyright (2013)..... 80

Figure 7.9: Another two threshold plots indicating the threshold at the crossing of the different lines.....83

Figure 8.1: Sketch of total time until the ground state is found with desired probability as a function of the problem size. The dotted lines show the performance for several fixed values of per-round run time t_f . The blue line shows the optimal result, reached if the run times t_f were optimized individually for each problem size. When measured with a fixed t_f (for example because of limitations of the annealing device), the slope of the measured curve (red) might indicate a wrong behavior: For small N, the slope

- is lower than optimal (possibly faking speedup where there is none), for large N , the slope is higher than optimal (which might mask possibly existing speedup) [Ami15]. "Reprinted figure with permission from [M. H. Amin. Phys. Rev. A, 92(5):052323, 2015.] Copyright (2015) by the American Physical Society."..... 95
- Figure 9.1: Physical volume required for performing dlog (blue) or factoring (orange) algorithms of different problem sizes (numbers on the right) as a function of the physical error rate. We use a threshold of $p_{th} = 10^{-2}$ and assume a symmetric error model (all kinds of errors equally likely) with equal rates p for initialization, gates, waiting and readout, which makes this plot hardware-agnostic and approximate. A surface code cycle consists of multiple gates, measurement, initialization, and classical processing. 124
- Figure 13.1: Double-logarithmic plot showing possible system sizes when implementing factoring (left) or dlog (right) algorithms for different runtimes, with experimental values from Table 13.2. (rather than the uniform error probabilities of the other chapters, see chapter 7.4.1 for background) The violet line gives the number of qubits for the time-optimized limit, *i.e.*, the minimal required runtime due to the T depth when having maximal parallelization. Implementations above the line would violate the temporal order of T gates and are thus not possible..... 149
- Figure 14.1: Double-logarithmic plot showing possible system sizes when implementing dlog algorithms for different runtimes, with experimental values from table 14.2. (rather than the uniform error probabilities used in the other chapters, see chapter 7.4.1 for clarification). The purple line gives the number of qubits for the time-optimized limit, *i.e.*, the minimal required runtime due to the T depth when having maximal parallelization. Implementations above the line would violate the temporal order of T gates and are thus not possible..... 162
- Figure 21.1: Comparison of algorithmic demands with currently achieved hardware performance. The plot shows required resources as number of qubits times rounds of error correction in the surface code for dlog (blue) and factoring (orange) for common key sizes (s . Chapter 9.5 and Figure 9.1) as a function of the physical error rate p . The squares show current realizations assuming one day run time (solid) or 100 days (empty), the yellow area shows expected near-term progress. Both scales are logarithmic. Note that unlike Chapters 13 and 14, we are assuming uniform error rates for gates, initialization, and readout, rather than those obtained from current experiments, see chapter 7.4.1.....209
- Figure 25.1: Translation of our schematic diagram (left of the "=") to stabilizer measurements on physical qubits as represented in [FMMC12]. One picture can correspond to one or more surface code rounds. Gray regions show where stabilizer measurements on the qubit array are turned off during these rounds. The edge color represents the type of active stabilizers next to the deactivated region. This is a distance 3 example, however, the schematic diagram is valid independent of distance, where gray regions can also include more than one deactivated stabilizer, *i.e.*, represent larger holes (proportions unaccounted for). Inside the holes, additional direct action on data qubits may be necessary, which is also not shown..... 226
- Figure 25.2: Sequence of different stabilizer measurement patterns for initialization of a Z -cut qubit in a Z eigenstate (a) or X eigenstate (b). Time runs from left to right..... 226
- Figure 25.3: Sequence of different stabilizer measurement patterns for measurement of a Z -cut qubit in the Z basis (a) or X basis (b)..... 227
- Figure 25.4: Sequence of different stabilizer measurement patterns for a logical Hadamard. Uncolored pictures represent additional physical gates between the stabilizer measurements (number of qubits chosen arbitrary). Here, the physical Hadamards in step 4 can be done during the measurement of the stabilizer cycle before, the physical SWAP gates (step 5–6) can be performed while stabilizer measurements are running in the rest of the array. This sequence basically shows a transformation to a different qubit encoding (step 3), where the logical X and Z operators are just vertical and horizontal lines, on which a Hadamard can be performed easily (steps 4–6) and a transformation back to the original shape of the qubit..... 227

Figure 25.5: Sequence of different stabilizer measurement patterns for a logical CNOT, in analogy to the braiding diagram in Figure 7.6 (d). Logical qubit operators are defined as horizontal operator chains between the holes, i. e., the two top holes build a logical qubit, and the two bottom ones build another qubit. Steps 2–3 show the creation of an X-cut qubit (green) in a Z eigenstate, step 4 is a projective measurement (with initialization following in step 5) of the upper Z-cut qubit in the X-basis, step 6 a measurement (destructive) of the X-cut qubit in the Z basis..... 228

Index of Tables

Table 7.1: Space-time volume in fundamental surface code building blocks for Clifford and T gates. One building block represents $4(d+d/4)^2$ physical qubits undergoing $d+d/4$ error correction cycles. The count for the S gate is an approximation, including the costs for a two-level distillation with constant distance d . The distillation cost for the T gate v_{dist} must be optimized separately for every problem. 87	
Table 9.1: Clifford+T gate counts for Grover-based key searches according to [LPS20] and [AMG+16a].....	111
Table 9.2: Toffoli gate counts for a dlog computation over an elliptic curve over a prime field $\text{GF}(p)$ with $n = \lceil \log_2(p) \rceil$, according to [RNSL17c, Table 2], [Roe17]. With recent work by Ekerå [Eke18] one may hope to halve the gate and depth counts, the number of qubits is not affected.....	122
Table 9.3: Toffoli gate counts for factoring an n -bit number according to [HRS17], [RNSL17c, Table 2]. With work by Ekerå and Håstad [EH17] one may hope to reduce the gate count by a small factor (4), the number of qubits is not affected.....	122
Table 9.4: Logical gate counts for elliptic curve Shor attack using dlog. With recent work by Ekerå [Eke18] one may hope to halve N_b , N_T , and $N_{T,s}$	123
Table 9.5: Logical gate counts for factoring an n -bit number. The Toffoli depth has been approximated for low parallelization similar to the dlog algorithm. "With work by Ekerå and Håstad [EH17] one may hope to reduce N_b , N_T , and $N_{T,s}$ by a small factor (4).....	124
Table 9.6: Minimal number of surface-code cycles for performing dlog or factoring algorithms with problem sizes n with one surface code cycle typically corresponding to a time of $\sim 10t_M$. Thus, during one surface code cycle one can run 10 subsequent T gates.....	125
Table 13.1: Summary of DiVincenzo criteria for planar Transmon qubits. ✓: Met routinely, ?: Met sometimes or meeting them is controversial, ×: not met.....	148
Table 13.2: Optimistic assumptions for different error rates and operation times combining the best reached values for each operation. This does not describe a current available setup, but shows what is in principle possible right now or in near future. Times are initialization, 1- and 2-qubit gate and measurement time. Probabilities are error probabilities for the respective processes. A surface code cycle contains 4 two-qubit gates, 2 one-qubit gates, measurement and initialization as well as classical processing.....	148
Table 13.3: Summary of DiVincenzo criteria for 3D Transmon qubits.....	151
Table 13.4: Optimistic assumptions for different error rates and operation times combining the best reached values for each operation. This does not describe a current available setup, but shows what is in principle possible right now or in near future.....	151
Table 13.5: Summary of DiVincenzo criteria for flux qubits.....	153
Table 13.6: Summary of DiVincenzo criteria for fluxonium qubits.....	154
Table 14.1: Summary DiVincenzo criteria for trapped ions.....	161
Table 14.2: Optimistic assumptions for different error rates and operation times combining the best reached values for each operation. This does not describe a current available setup, but shows what is in principle possible right now or in near future. Times are initialization, 1- and 2-qubit gate and measurement time. Probabilities are error probabilities for the respective processes. A surface code cycle contains 4 two-qubit gates, 2 one-qubit gates, measurement and initialization as well as classical processing.....	161
Table 15.1: Summary DiVincenzo criteria for spins in quantum dots.....	175
Table 15.2: Short summary double-dots.....	176

Table 15.3: Short summary single-dots.....	176
Table 15.4: Short summary triple-dots.....	176
Table 15.5: Summary DiVincenzo criteria for other Quantum dot platforms.....	176
Table 15.6: Summary DiVincenzo criteria for Single Dots in Si/SiGe.....	177
Table 15.7: Summary DiVincenzo criteria for NV centers.....	178
Table 15.8: Summary of the experimental status of DiVincenzo criteria for Majorana Fermions.....	180
Table 15.9: Summary of the experimental status of DiVincenzo criteria for fractional quantum Hall states.	181
Table 15.10: Summary of the experimental status of DiVincenzo criteria for topologically protected qubits in strontium ruthenate.....	181
Table 15.11: Summary of the experimental status of DiVincenzo criteria for topologically protected qubits in Josephson junction arrays.....	182
Table 16.1: Summary DiVincenzo criteria for Rydberg atoms.....	186
Table 17.1: Summary of DiVincenzo criteria for single photons.....	191
Table 17.2: Summary DiVincenzo criteria for CV and Gaussian encodings.....	192
Table 18.1: Summary DiVincenzo criteria for liquid- and solid-state molecular approaches.....	198
Table 26.1: Records for Josephson qubits.....	229
Table 26.2: Records for trapped ions.....	230
Table 26.3: Records for semiconductor platforms.....	230
Table 26.4: Records for neutral atoms.....	230
Table 26.5: Records for photonic qubits.....	231
Table 26.6: Records for molecular approaches.....	231

1 Deutsche Zusammenfassung

1.1 Was ist ein Quantencomputer?

Heutige Computer behandeln Informationen gemäß den Gesetzen der klassischen Physik: Register und Speicherinhalte haben zu jedem Zeitpunkt einen einzigen Wert. Dies gilt ungeachtet der Tatsache, dass die Bauelemente eines Computers wie Transistoren auf den Gesetzen der Quantenphysik basieren.

In einem Quantencomputer wird die Information selbst quantenmechanisch behandelt: Register und Speicherinhalte können mehrere Werte gleichzeitig in Überlagerung enthalten, und Maschinenbefehle wirken sich simultan auf all diese Werte aus. Damit ist bereits ein einziger Quantenprozessor intrinsisch massiv parallel, ohne parallelisierte Hardware wie mehrere Prozessorkerne zu benötigen.

Nutzung dieser Parallelität erfordert allerdings Umgang mit dem probabilistischen Charakter der Quantenphysik und das Kompilieren von Algorithmen in quantenmechanisch erlaubte Gatter (Quantenschaltkreise). Aus diesem Grund erfordert die Nutzung der Quantenbeschleunigung zunächst die Entdeckung geeigneter Algorithmen. Zu diesen gehören bisher die schnelle Datenbanksuche, das Durchsuchen von Graphen, die Lösung linearer Gleichungssysteme, Anwendungen der schnellen Fouriertransformation einschließlich Faktorisierung und Berechnung diskreter Logarithmen, und die Simulation von Quantensystemen einschließlich Chemikalien und neuer Materialien, sowie Maschinenlernen und Optimierung. Für einige dieser Anwendungen, insbesondere die letztgenannten, ist die Quantifizierung der erreichbaren Quantenbeschleunigung noch Gegenstand aktueller Forschung. Quantencomputer sind darum – aufgrund der möglichen Anwendungen aber auch aufgrund der aufwändigen Hardware, auf der Ebene von Rechenzentrumstechnologie und Höchstleistungsrechnen anzusiedeln und keine Büro- oder gar mobile Technologie.

Quantencomputer wurden zunächst als hypothetische, theoretische Konstruktion eingeführt. Inzwischen, nach mehr als 20 Jahren Entwicklung seit den ersten Laborexperimenten, beginnt sich das Feld der Hardwareplattformen zu konsolidieren und Zugriff auf Quantenprozessoren wird als Dienstleistung von mehreren Firmen angeboten, eine sehr spezielle Quantencomputerplattform wird auch kommerziell angeboten. Diese Quantenprozessoren erlauben die Entwicklung und Evaluation von Quantenalgorithmen, sind aber noch in keiner Anwendung klassischen Rechnern überlegen. Lediglich für ein (rein) akademisches Problem ohne offensichtliche Anwendung konnte eine Hardwareplattform von Google diesen als Quantum Supremacy bezeichneten Schnittpunkt bislang erreichen.

1.2 Die Relevanz von Quantencomputern für die Kryptoanalyse

Ein Großteil der heute auf breiter Basis eingesetzten asymmetrischen kryptographischen Verfahren kann nicht mehr als sicher betrachtet werden, sobald die Faktorisierung großer Ganzzahlen und die Berechnung sogenannter diskreter Logarithmen effizient möglich ist. Diese einfache Beobachtung erklärt direkt das signifikante Interesse an Quantencomputern in der kryptoanalytischen Forschung – Peter Shor zeigte erstmals, dass beide Probleme asymptotisch effizient gelöst werden können, wenn ein hinreichend großer und verlässlicher Quantencomputer verfügbar ist. Die Effizienz der Shorschen Algorithmen beruht unter anderem auf der geschickten Nutzung von Überlagerung, einer Technik die mit klassischen Bits nicht realisierbar ist. Quantencomputer verwenden als elementare Einheit Qubits, bei denen den klassischen Werten 0 und 1 lediglich die Rolle von Basiswerten zukommt, und der Wert eines Qubits komplex gewichtete Anteile beider Basiswerte simultan innehaben kann. In ähnlicher Weise werden klassische Bitregister durch komplexe Quantenregister ersetzt, die effiziente hochdimensionale Berechnungen ermöglichen. Aus praktischer Sicht stellt sich die Frage, wie groß ein Quantencomputer sein muss, um real eingesetzte kryptographische Verfahren, etwas aus der RSA-Familie oder aus der Elliptischen-Kurven-Kryptographie, zu gefährden. Hierzu ist eine genauere Analyse bekannter Quantenalgorithmen erforderlich.

Die abstrakten Schritte eines Quantenalgorithmus müssen für das konkret angegriffene Verfahren (effizient) in Elementarschritte umgesetzt werden, die wiederum auf reale Hardware abbildbar sind.

Detaillierte Kostenanalysen sind erst in geringem Umfang in der Literatur verfügbar, und es ist realistisch anzunehmen, dass die bislang veröffentlichten Quantenschaltkreise noch weiter optimiert werden können. Aber die verfügbaren Arbeiten lassen es bereits machbar erscheinen, die Shorschen Algorithmen für kryptographisch interessante Parameterwahlen in Quantenschaltkreise moderater Komplexität zu übersetzen.

Konkret werden für einen Angriff auf 2048 Bit RSA insgesamt $5.5 \cdot 10^{15}$ Elementarschritte auf 4098 logischen Qubits benötigt; andere Abwägungen zwischen der Anzahl der Rechenschritte und der Anzahl der Qubits sind möglich. Für den diskreten Logarithmus auf einer elliptischen Kurve über 256 Bit werden 10^{14} Rechenschritte auf 2330 logischen Qubits benötigt. Dies sollte nicht mit den physikalischen Qubits verwechselt werden, deren Zahl wir in Abschnitt 1.5 besprechen, und die um Größenordnungen höher liegt. Nach einer aktuellen Abschätzung werden 20 physikalische Megaqubits als hinreichend für einen Angriff auf 2048 Bit RSA mit einer Laufzeit von acht Stunden betrachtet.

Für die symmetrische Kryptographie bieten Quantencomputer ebenfalls neue kryptoanalytische Möglichkeiten, aber mit den momentan bekannten Algorithmen sind die Auswirkungen deutlich weniger spektakulär als im asymmetrischen Fall. Auch hier kann davon ausgegangen werden, dass die besten vorhandenen quantitativen Aussagen, etwa zur Schlüsselsuche bei AES-128, noch verbessert werden (es wurden bereits mehrere Optimierungen vorgeschlagen), aber eine Vergrößerung der Schlüssellänge erscheint momentan eine wirksame Gegenmaßnahme. Spektakulärere Quantenangriffe auf symmetrische Primitive sind bekannt, aber gerne wird hier ein Angriffsmodell verwendet, welches bei heute genutzten Implementierungen nicht realistisch ist.

1.3 Hardware für Quantencomputer

Diese gesicherten Erkenntnisse über Quantenalgorithmus wären nicht relevant, würde nicht parallel Hardware entwickelt. Es wird eine ganze Reihe von Hardwareplattformen weltweit verfolgt, die durchaus dramatisch unterschiedlich sind—etwa vergleichbar mit dem Übergang von mechanischen zu elektronischen Computern. Die augenblicklich führenden Plattformen sind

1. Gefangene Ionen—eine Plattform die u.a. mit der Technologie von Atomuhren verwandt ist.
2. Integrierte Schaltkreise aus Supraleitern—eine Plattform, die Ähnlichkeit mit aktuellen Computerchips hat, jedoch aus anderen Materialien besteht und bei sehr tiefen Temperaturen betrieben wird.




Abbildung 1: Die augenblicklich führenden Quantencomputing-Plattformen—mikroskopische Perspektive. Links: Josephson-Prozessor (Foto: Julian Kelly, Google); Mitte: Lineare Ionenfalle (Foto: Jürgen Eschner, Universität des Saarlandes); Elektroden zum Fangen (Stäbe) und Linsen zum Einstrahlen von Lasern für Quantenlogik (links und rechts); Rechts: Die gleiche Ionenfalle eingebettet in ihre Vakuummapparatur.

Es wird eine Vielzahl weiterer Plattformen erforscht, die im Augenblick weniger weit fortgeschritten sind, darunter Nanostrukturen in Halbleitern, gezielt dotierte künstliche Diamanten und Atome in Laserfeldern. Es ist durchaus möglich, dass diese in den kommenden Jahren aufholen. Insbesondere der Fortschritt von Halbleitern war im Jahr 2018 sehr eindrucksvoll.

Die führenden Plattformen erleben gerade den Übergang von naturwissenschaftlicher Grundlagenforschung in angewandte Forschung unter Industriebeteiligung und inhaltlich den Übergang von grundlegenden wissenschaftlichen zu technologischen Herausforderungen bei der Entwicklung komplexer, gesteuerter Quantensysteme.

Die strukturelle Herausforderung des Gebietes ist dabei die Fehleranfälligkeit von Quantencomputern. Diese geht über das Technologische hinaus und ist grundsätzlicher Natur—der besondere Glücksfall der Fehlertoleranz von klassischen Digitalrechnern tritt hier nicht ein. Quantencomputern können kryptoanalytische Aufgaben realistisch nur dann bewältigen, wenn sie aktiv fehlerkorrigiert werden. Ein konsistentes theoretisches Gerüst dieser Fehlerkorrektur wurde entwickelt. Seine praktische Umsetzung ist Gegenstand intensiver Forschung und erste Erfolge wurden erzielt. Diese Fehlerkorrektur beeinträchtigt die grundsätzliche Effizienz von Quantencomputing nicht, ist aber trotzdem durch einen enormen Overhead gekennzeichnet—die *logischen* Quantenbits (Qubits) die einen Algorithmus beschreiben bestehen aus einer großen Zahl von Bauelementen, die *physikalische* Qubits darstellen. Auch bei großem Fortschritt ist davon auszugehen, dass der Bau eines leistungsfähigen fehlertoleranten Quantencomputers nicht nur eine wissenschaftlich-technische Herausforderung darstellt, sondern im Ergebnis eine Großanlage vom Umfang eines Rechenzentrums wäre.

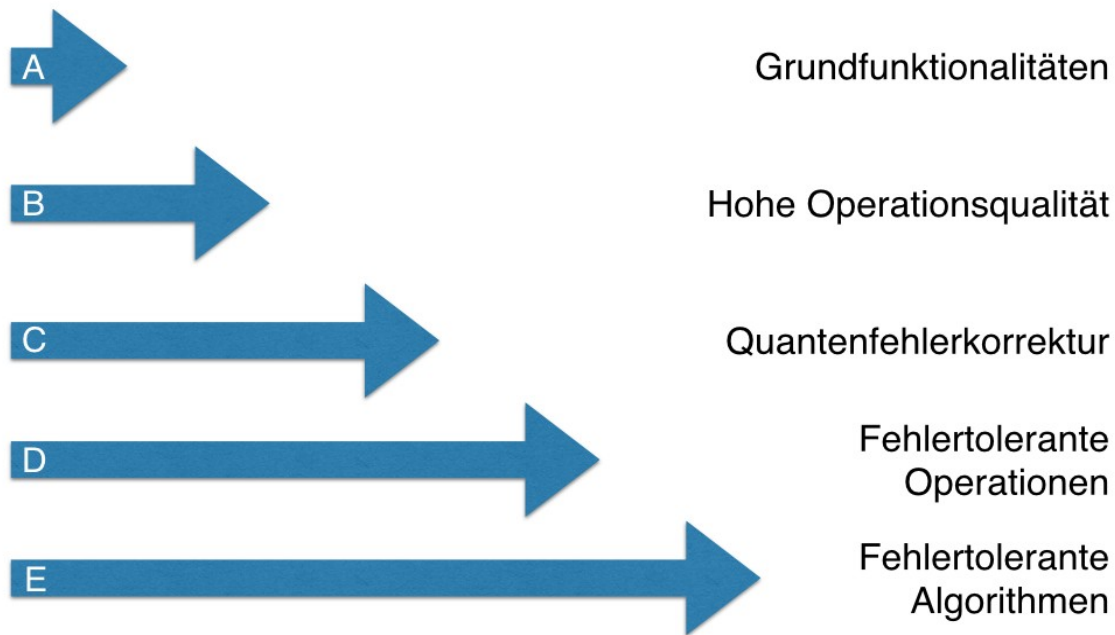


Abbildung 2: Schichtenmodell zur Bewertung von Quantencomputerplattformen anhand demonstrierter Schritte zur Fehlertoleranz

Forschungsergebnisse und die sie begleitenden Schlagzeilen können im Kontext der benötigten aktiven Fehlerkorrektur evaluiert werden. Die Studie enthält darum ein Schichtenmodell zur Bewertung von Quantencomputer-Kandidaten, beginnend mit der Demonstration von Grundfunktionen (Schicht A) bis hin zur fehlertoleranten Implementierung von Algorithmen (Schicht E), Abbildung 2. Augenblicklich wird von den führenden Plattformen Schicht C erreicht. Das Feld an Plattformen ist dicht und schnelle Veränderung der Bewertung wird erwartet, Abbildung 3.

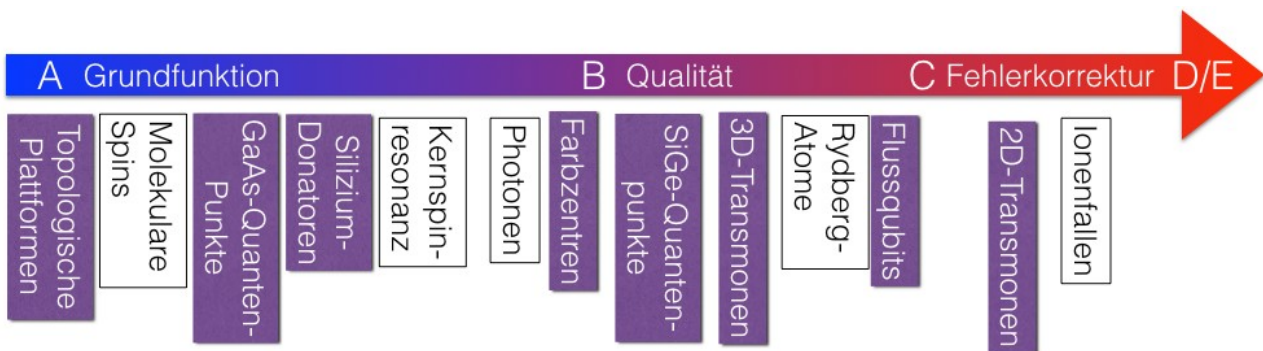


Abbildung 3: Einordnung verschiedener Plattformen im Schichtenmodell. Atomphysikalisch/optische Systeme sind weiß und Festkörpersysteme lila hinterlegt.

1.4 Aktuelle Entwicklungen

Quantencomputer und ihre Anwendungen sind augenblicklich fast wöchentlich in den Schlagzeilen der Technikpresse. Insbesondere wird dort ein „Rennen“ zur Realisierung von Quantencomputern mit den Hauptakteuren Google, IBM und—in etwas anderen Rollen—D-Wave Systems aus Kanada sowie Microsoft beschrieben. Dies spiegelt die reale rasante Entwicklung des Gebietes—auch während dieser Studie—wider. Die dort beschriebenen Systeme beschreiben 20–50 *physikalische* Qubits, deren Qualität bei weitem nicht ausreicht, Kryptoanalyse ohne Fehlerkorrektur zu erreichen. Ihr Interesse liegt in einem eingeschränkten Bereich von nicht-kryptographischen Anwendungen von Quantencomputern z.B. in der Quantensimulation.

Unsere Studie schätzt, dass bei robustem technologischem Fortschritt (entsprechend einer Fehlerrate von 1:10000) zum Brechen von 2048-bit RSA in 100 Tagen etwa eine Million physikalische Qubits, in 1 Stunde etwa eine Milliarde physikalische Qubits benötigt werden, siehe Abbildung 4c. Die unterschiedlichen Zahlen erklären sich aus einem Platz-Zeit-Kompromiss in der Quantenfehlerkorrektur. Dies liegt noch Größenordnungen über den aktuell realisierten Zahlen. Bei heutigen Fehlerraten wäre die Zahl der benötigten physikalischen Qubits hundert mal so groß, siehe Abbildung 4a und 4b. Auch mit dem Overhead der Fehlerkorrektur stellt dies eine *Schwelle* dar: Wenn Quantencomputer einmal asymmetrische Kryptographie erfolgreich angreifen können, dann ist eine weitere Vergrößerung der Schlüssel in der gleichen kryptographischen Technik nicht erfolgversprechend—der zum erfolgreichen Angriff benötigte Quantencomputer ist nur wenig größer. Dieser Zusammenhang wird verdeutlicht in Abbildung 4c. Dort ist der Zusammenhang zwischen der Größe kryptoanalytischer Aufgaben, der Fehlerrate, und der Zahl der physikalischen Qubits dargestellt. Es zeigt einerseits die Größe des zu überwindenden Weges, andererseits aber auch durch die enge Lage der Kurven dass nach dem Erreichen der untersten Kurven sehr schnell die ganze Kurvenschar erreichbar sein kann.

Es ist allerdings zu beachten, dass die aktuelle Entwicklung sich vermutlich noch beschleunigt: Nachdem im ersten Jahrzehnt der ernsthaften Hardwareforschung an Quantencomputern vor allen Dingen akademische Akteure in traditionellen Programmen aktiv waren, treten jetzt einerseits starke Industrieakteure auf den Plan, andererseits werden große Forschungsprogramme wie das EU-Quantentechnologieflaggschiff aufgelegt. Eine Sonderrolle nehmen hier die USA ein, wo Forschungsförderung von Militär und Heimatschutzministerium schon sehr lange die Erforschung und Entwicklung von Quantencomputern fördern. Da die technologischen Schritte, die für einen fehlertoleranten, kryptoanalytisch relevanten Quantencomputer relevant sind (große Zahl Qubits, geringe Fehlerrate) auch die aktuellen Herausforderungen prägen, ist jeder Fortschritt hier auch ein Schritt näher an einen kryptographisch relevanten Quantenrechner. Bereits mit der Aktualisierung 2018 hat sich angedeutet, dass Fortschritt in der Fehlerkorrektur und der Kompilation etwa eine Größenordnung an Quantenbits für eine gegebene Aufgabe einsparen könnten, wenn sie bestätigt werden.

1.5 Extrapolation

1.5.1 Stand 2018

Der große Aufwand für Fehlerkorrektur macht es auf absehbare Zeit unwahrscheinlich und vermutlich auch wirtschaftlich uninteressant für akademische und industrielle Labors, einen kryptographisch relevanten Quantencomputer zu realisieren. Wenn jedoch eine große Industrienation ihre Forschungsanstrengungen auf dieses Ziel konzentrieren würde, ähnlich den Manhattan- und Apollo-Projekten des 20. Jahrhunderts, so erscheint ein Quantencomputer mit wenigen Millionen physikalischer Qubits, der zumindest in 100 Tagen 2048-Bit RSA brechen kann, erreichbar, wenn auch die physikalische Fehlerrate angemessen sinkt und in einen Bereich von 1:10000 gebracht werden kann. Dies wäre eine Großanlage, die in mehrerlei Hinsicht technologische Rekorde benötigen würde und ggf. Zugriff auf seltene Materialien erfordert.

Die Forschung an Quantencomputern entwickelt sich sehr schnell. Für supraleitende Qubits lässt sich diese Entwicklung unterteilen in eine Entwicklungsperiode von 1999–2015 und die Ära der Cloud-Quantenprozessoren seit Frühjahr 2016. Sollte die Ankündigung von 50 bzw. 72 Qubits als nächsten Meilenstein von IBM und Google im Jahr 2019 ebenfalls eintreten, was durchaus wahrscheinlich ist, entspricht dies mehr als einer Verdoppelung jedes Jahr.

Auch im aktuellen Zeitalter erster kommerzieller Anwendungen ist die Forschung an Quantencomputern noch jung genug, dass überraschende qualitative Basisinnovationen, die zu einer Beschleunigung ihrer Entwicklung führen könnten, weiterhin möglich sind.

1.5.2 Aktualisierung 2020

Die Meilensteine, die 2018 angekündigt wurden, sind zwar mit Verspätung eingetreten, wurden aber erreicht. Hier ist die Qubitzahl bei nahezu gleicher Fehlerrate gestiegen, und die damit verbundenen operationellen Herausforderungen wurden gemeistert. Die Ergebnisse dieser Experimente zeigen aber auch deutlich, dass die Fehlerrate das deutlich limitierende Element ist, was den Fokus zumindest in den nächsten Jahren weg von immer größeren Qubitzahlen verschieben dürfte.

1.6 Aktualisierungen

1.6.1 Frühjahr 2019

Die Überarbeitung im Frühjahr 2019 betrifft sowohl den Inhalt als auch die Darstellung.

Es werden zwei algorithmische Innovationen diskutiert, die ohne Fehlerkorrektur auskommen und die auf Fortschritten im Quantencomputing für andere Anwendungen basieren: Variationelle Faktorisierung (9.2) und Faktorisierung auf adiabatischen Quantencomputern (9.1). Beides sind gültige Algorithmen, deren Beschleunigung gegenüber klassischer Faktorisierung jedoch weder erwiesen noch wahrscheinlich ist, so dass sich die Studie weiterhin auf fehlertolerantes Quantencomputing konzentriert. Innerhalb dieses Paradigmas sind seit der ersten Version dieses Reports keine strukturellen (asymptotischen) Durchbrüche zu vermelden, welche Kostenabschätzungen für Angriffe mit Quantenrechnern komplett in Frage stellen. Es sind jedoch durchaus algorithmische Verbesserungen zu vermelden, Kapitel 9.6, die konkrete Kostenschätzungen, etwa für die Anzahl der Quantengatter zum Faktorisieren eines RSA-Modulus, senken. Weitere derartige Verbesserungen "um kleine konstante Faktoren" in den kommenden Monaten wären nicht unerwartet, und in Kombinationen mit Verbesserungen bei der Fehlerbehandlung könnte hier eine Verringerung der Angriffskosten um einen moderaten konstanten Faktor realistisch werden. Auch bei Kostenabschätzungen für die Schlüsselsuche bei symmetrischen Verfahren sind weitere moderate Verbesserungen durch effizientere Quantenschaltkreise zu erwarten.

Im Bereich der Quantenfehlerkorrektur wurden vielversprechende Ansätze zur Reduzierung des Overheads durch Gitterchirurgie besprochen (7.2.5), deren tatsächliches Potenzial aber noch nicht geklärt ist und darum nicht die Extrapolation beeinflusst. Ein spektakuläres Ergebnis mit sehr hohen Korrekturschwellen bei asymmetrischen Fehlern wird kritisch diskutiert, da die verwendete Methodik nicht skalierbar ist, Kapitel 7.4.3.

Die Diskussion der Rolle unitärer Fehler wurde stark ausgebaut, Kapitel 7.4.1.3: Einerseits ist jetzt deutlich präzisere Literatur erschienen als in der ersten Version, andererseits haben wir in verschiedenen Anhängen die Fehlermodelle präziser abgegrenzt. Es zeigt sich, dass überwiegend unitäre Fehler zwar die Grundaussagen der Fehlerkorrektur nicht verändern, aber einen Einfluss auf die quantitativen Ergebnisse haben können – welchen genau ist noch unklar.

Insbesondere sind die Besonderheiten der Fehlermodelle jetzt im zentralen Kapitel 7.4.1 zusammengestellt mit technischen Details in einem ausgebauten Anhang, Kapitel 22, die es auch erlauben die Annahmen hardware-spezifischen Ergebnisse in Kapiteln 13.2.3 und 14.2.3 und der Extrapolation auf zukünftige Hardware (Kapitel 21) zentral zu verifizieren. Zur Verbesserung der Lesbarkeit der vielen technischen Begriffe haben wir ein Glossar zur Fehlerkorrektur eingefügt, Kapitel 7.7.

Die Tabellen zur Performance der Hardware wurden aktualisiert. Große Unterschiede ergaben sich bei den Halbleiterplattformen, Kapitel 15, die zwar weiterhin der Bewertungsstufe „B“ anzusiedeln sind, deren Fehlerraten sich aber dem für Fehlerkorrektur interessanten Bereich genähert haben.

1.6.2 Aktualisierung Frühjahr 2020

Auf der algorithmischen Seite sind Fortschritte in der Anwendung von Grovers Algorithmus zur Schlüsselsuche bei AES in Kapitel 9.5.1 dokumentiert; Clifford+T Ressourcenabschätzungen wurden aktualisiert, um effizientere Implementierungen der AES S-Box als Quantenschaltkreis abzubilden. Auf symmetrischer Seite gab es weiterhin Fortschritte bei der Nutzung von Simons Algorithmus, auf die wir in Kapitel 9.5.4 kurz eingehen.

Bei der Faktorisierung von RSA-Moduli und der Berechnung diskreter Logarithmen in Primkörpern finden jetzt Ergebnisse von Gidney und Ekerå Berücksichtigung, und Abschnitte 9.6 und 9.7 wurden entsprechend aktualisiert. Bei der Diskussion der Berechnung diskreter Logarithmen auf elliptischen Kurven über Primkörpern sind neue Ergebnisse von Häner et al. eingearbeitet. Diese verwenden zahlreiche Messungen als Teil der Berechnung und verringern hierdurch die Anzahl der T-Gatter. Ferner wurde in Abschnitt 9.6.3 eine kurze Diskussion des Verfahrens von Harrow et al. (und Nachfolgearbeiten) zur Lösung linearer Gleichungen eingefügt.

In der Zeit der Aktualisierung wurde von Google das vielbeachtete Ergebnis zur Quantenüberlegenheit veröffentlicht. Auch darüber hinaus haben sich Aktivitäten in Forschung und kommerzieller Anwendung im Bereich Quantencomputing stark entwickelt. Diese konzentrieren sich auf nicht-fehlerkorrigierte (NISQ) Ansätze. Da, siehe Kapitel 4, die tiefen kryptoanalytischen Algorithmen mit NISQ unerreichbar sind, haben diese Ergebnisse zwar keine unmittelbare Auswirkung auf Kryptoanalyse und finden vor allen Dingen auf der zweiten Schicht unseres Bewertungsschemas statt, jedoch ist zu erwarten dass die hier entwickelten Technologien auch die Fehlerkorrektur und -toleranz weiter beflügeln.

Wir beschreiben das Google-Experiment in Kapitel 6.6 und die darunter liegende Methode des Benchmarkings in 6.4.5 sowie eine Verifikation von Fehlerkorrektur am Ende von Kapitel 22.3. Infolge von NISQ sind neue Heuristiken zur Optimierung entstanden, u.a. diabatisches Quantenannealing. Eine Methode zur Entschlüsselung von gitterbasierter Kryptographie wurde für diese Plattform vorgestellt, jedoch (folgend dem Charakter einer neuen Heuristik) besteht kein überzeugender Hinweis auf echte Quantenbeschleunigung, Kapitel 8.1.3.

Neben diesen neuen Abschnitten wurden zahlreiche Diskussionen im Bereich Fehlermodelle überarbeitet und Teile vorheriger Anhänge in Kapitel 6.5 konsolidiert und die Akteursbeschreibungen aktualisiert.

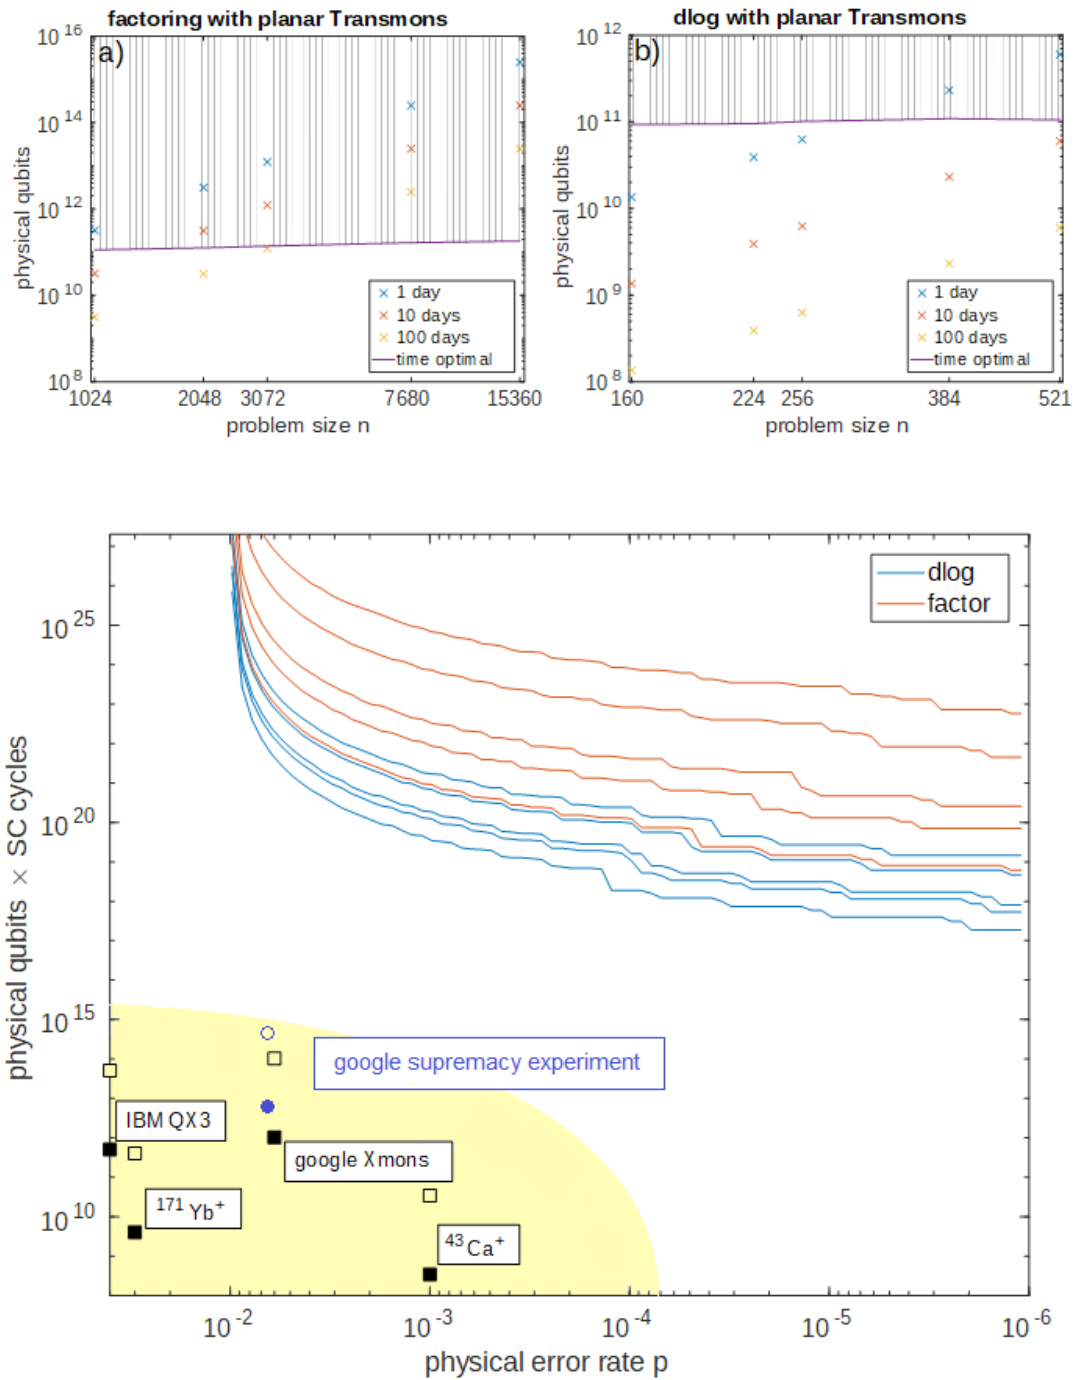


Abbildung 4: Standortbestimmung: Anforderungen an die Größe von Quantencomputern a) zur Faktorisierung und b) für diskreten Logarithmus bei heutigen Fehlerraten für verschiedene Laufzeiten als Funktion der Problemgröße c) Wechselspiel von Fehlerrate und Zahl der Qubits für algorithmische Anwendungen. Gelber Bereich: Stand der Forschung anhand einzelner Experimente bei Laufzeit von einem Tag (gefüllte Quadrate) oder 100 Tagen (offene Quadrate); Linien: Anforderungen von dlog für 160,224,256,384 und 521 Bit bzw. Faktorisierung für RSA von 1024, 2048, 3072, 7680 und 15360 bit.

2 Synopsis

2.1 Basic idea

The often counter intuitive concepts of quantum physics are well understood and precisely confirmed in science. First applications of simple quantum physics have been known for a long time—transistor, laser, magnetic resonance, nuclear technology and others. These applications use a few quantum properties of an otherwise macroscopic system (Quantum Technology 1.0). Currently, a new generation of Quantum Technologies 2.0 is emerging, which uses many more unique properties of quantum physics and addresses single quantum systems—one of them is the concept of a quantum computer. Quantum computers use the feature of quantum physics that the system state can simultaneously occupy many if not all classically permitted states. This way, it can internally process data in a massively parallel way without the need for parallel hardware instances. Algorithms for quantum computers must conform to various constraints, most notably does one need to remove (uncompute) parallel state occupation in the end of an algorithm in order to avoid random output.

There are a number of known quantum algorithms for cryptographic tasks. Most prominent are Shor's algorithms for factoring integers and for the computation of discrete logarithms. Shor's algorithms represent significant progress for standard asymmetric cryptographic protocols (including RSA-based and common elliptic curve-based methods). In principle they permit the efficient reconstruction of a secret key from public data; there are no known analogous classical attacks. Quantum algorithms also permit improvements compared to classical techniques when analyzing symmetric cryptographic protocols. Grover's method for the acceleration of a complete key search is probably the most well-known of such algorithms. Nevertheless, cryptanalytic progress through quantum algorithms is significantly less spectacular in the field of symmetric methods if one remains restricted on established threats and they do not endanger existing symmetric protocols from what is currently known.

Quantum computing was first proposed by Nobel laureate Richard Feynman in 1982 as a tool to simulate quantum systems. This research field has expanded since the discovery of Shor's factoring algorithm 1995 which can be viewed as the starting point of the global activities towards constructing a quantum computer. Since then, a number of physical platforms to realize such a computer are being pursued. Quantum computing is an interdisciplinary research area between physics, computer science, and engineering which is being pursued in universities, research centers, and companies. Milestones such as the establishment of a division for quantum information in the American Physical Society and the announcement of a European Quantum Technology Flagship program, both in 2016, have made quantum computing an established research discipline.

2.2 Hardware platforms

Similar to the early days of classical computing there exists a wide variety of hardware platform candidates for quantum computing today. These, on the one hand, need to display detectable quantum effects—which means they need to be small and isolated—on the other hand they need to be operated as computers, i.e., their technology needs to be scalable and permit access to write, read, and control. This brings together challenges within science and engineering—isolation and access need to be provided simultaneously.

The structuring element for the selection of platforms by researchers and their evaluation is their sensitivity to operational errors—quantum error corrections is driving architectures and overhead.

2.2.1 Global categories

Atomic platforms use elementary quantum systems such as single atoms in which the laws of quantum mechanics can be naturally resolved but where scaling and control are a challenge.

Solid-state platforms use various types of integrated circuits which are naturally scalable and controllable, in which the main challenge is the realization of quantum effects and their stabilization over a long time.

Momentarily, the most advanced platform in atomic physics are trapped ions. This technology is related to the development of atomic clocks and already shows precise quantum operations on large registers up to 14 quantum bits (qubits). In the area of solid-state platforms, the currently leading approach are Josephson-qubits—integrated circuits made from superconducting metals such as Aluminum and Niobium: Currently 17 qubits have been demonstrated and 22 are being discussed on conferences.

Beyond those momentary leaders, there is a range of candidate that have the potential to catch up and overtake. These include silicon-based nanotechnology, trapped neutral atoms, and with some degree of speculation topological states in nanowires.

2.3 Algorithmic goals

Attacks using quantum computers frequently aim at the direct reconstruction of a secret key under rather moderate assumption—only access to a public key or a few text-cipher pairs is assumed. Beyond that also complex attacks using quantum technologies have been proposed, which on the one hand have impressive potential but which on the other hand are based on assumptions that are not satisfied by real implementations. If one allows the attacker to run the targeted implementation with inputs in superpositions, theoretically interesting models of attack can be formulated, but this type of access is not given in classical implementations.

A current focus in the literature on quantum cryptanalysis is a detailed cost analysis of (abstractly) known attacks applied to relevant cryptographic instances (such as 2048-bit RSA, AES, or SHA-2). Grover's and Shor's algorithms are fundamentally based on performing computations within the symmetric primitive under attack or within the algebraic structure behind an asymmetric method on a quantum computer. The relevant computations are expressed as quantum gates. The quantum gate model can then be the interface to the underlying computational models. Even though the fundamental efficiency of Shor's algorithms is not based on the details of the cryptographic protocol under attack, the details of the underlying quantum circuit are essential for a quantitative cost estimate. In the case of computing a discrete logarithm on an elliptic curve, for example, the curve arithmetic is mapped on quantum gates, which can be done in different ways. Analogously, in factoring with Shor's method it is necessary to implement modular arithmetic with quantum gates and in a Grover-based key search for AES, AES-encryption is implemented with quantum gates. In the design of these rather complex circuits, one usually designs a classical reversible implementation first, which then is translated into elementary quantum gates (Clifford+T).

Typical goals of optimizations are the reduction of the number of qubits, circuit depth, and/or number of gates. There, one typically differentiates between gate types in order to take into account different complexities in physical implementations. We list relevant cost estimates from the literature. If robust quantum processors are provided, it is realistic to realize cryptographically relevant computations of discrete logarithms. Also factoring of larger integers appears realistic. Key search in AES on the other hand appears to remain a large challenge even with reliable quantum processors—the asymptotically exponential scaling of Grover's algorithms represents a serious obstacle.

2.4 Computational Models

The concrete realization of quantum algorithms is discussed in different computational models. The by far most relevant model for cryptanalysis is the fault tolerant implementation of the quantum gate model.

The quantum gate model resembles the operation of a classical computer: A sequence of logical operations, gates, in a simple machine code is applied on a data register which is read out in the end. The quantum superpositions which contribute to the performance of quantum computers only appear in intermediate steps of the computation. As these superpositions are related to output only probabilistically, the superpositions have to be mapped on a result that can be read out with certainty in the end of the computation. This constitutes a key challenge in the discovery of quantum algorithms which, as explained above, has been solved for a number of cryptanalytic scenarios. For an ideal implementation of this models, quantum speedup is mathematically proven. As a perfect, error-free implementation of such an algorithm is impossible, it is the goal of a physical realization to approximate it as close as possible.

The fault-tolerant implementation of the gate model relates to the observation that quantum operations and hardware are much more susceptible to error than their classical counterparts. It is thus necessary to correct errors during the operation. This can in principle reduce the probability of an error in the final result to an acceptable, predetermined size such that the computer potentially needs to run only a few times. Quantum error correction has a number of peculiarities based on the analog character of quantum operations as well as the invasive nature of quantum measurements. Still, a mathematical framework of quantum error correction has been formulated, culminating in the use of the surface code. The overhead imposed by quantum error correction is significant and determines size and speed of potential quantum computers without challenging basic speedup. In particular, they set a threshold for the physical error rate, below which error correction is possible and effective. Hardware below this threshold can thus be used to simulate an ideal quantum algorithm using error correction.

Currently, results on quantum advantage or quantum supremacy make regular headlines. This describes imminently reaching state at which quantum computers cannot be simulated by current classical supercomputers any more. This point has been reached in 2019 and it is expected that (given the exponential overhead of simulating a quantum computer with classical hardware) will not be reversed. On these expected platforms which are assumed to be operated without error correction, one can only execute algorithms with low gate counts. Applications of such processors on non self-referential problems are currently developed. They are found in the area of quantum simulation where on classical architectures the memory needs are a limiting resource. Cryptanalytic applications of this approach are not known and rather unlikely.

The technique of quantum annealing and adiabatic quantum computing makes headlines based on the products of the Canadian company d-Wave systems. As this hardware is less demanding to build than that for the gate model, large processors up to 2000 units have been realized. The products of d-Wave systems are designs for optimization problems and can be programmed in a versatile way. Systematic analysis have so far negated that they show quantum speedup. In principle, quantum annealing can be applied to cryptanalytic problems and can lead to acceleration but a key hardware element necessary for that has so far not been realized. A number of platform-specific models such as one-way quantum computing are evaluated separately.

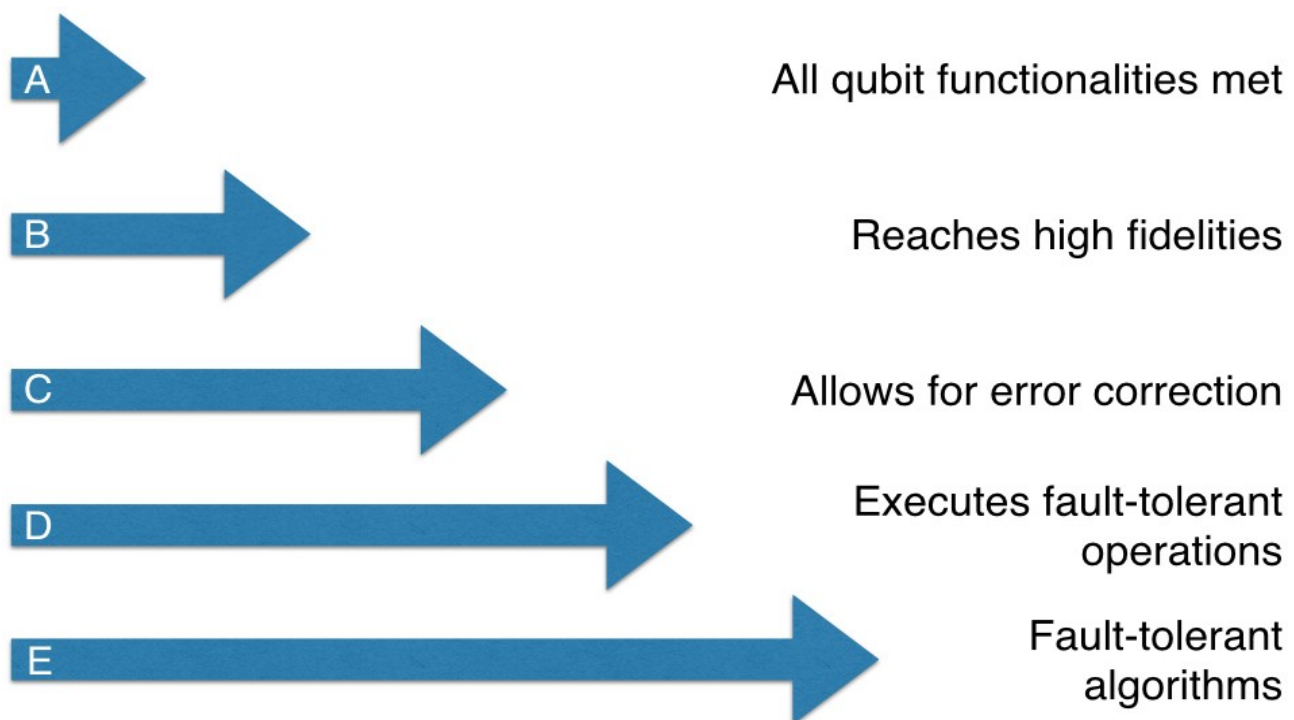


Figure 2.1: Levels of development towards a fault-tolerant quantum computer

2.5 Evaluation along computational models

Similar to the software stack of modern computer architectures (from machine code to a user interface) we can organize quantum computer evaluation from the bottom up: We propose to use five levels. If the quality of operations identified on level B allows to implement cryptanalysis without error correction, the following levels could be skipped—but this appears impossible following current thinking:

A: Basic functionality the has candidate for a quantum computing already demonstrated all basic functionalities of quantum processor (qubits, gates, initialization, coherence, readouts)? Were all these functionalities demonstrated in the same experiment containing more than two qubits?

B: Quality of operations Has the error rate of all relevant operations been measured? Are they compatible with error correction thresholds? Have all ingredients of a fault-tolerant architecture been demonstrated?

C: Error correction Has quantum error correction been demonstrated and is it effective? Are logical error rates smaller than physical error rates?

D: Fault tolerant operation Have operations on logical qubits been implemented in a fault-tolerant way? Has these been achieved for a universal set of gates (Clifford+T)?

E: Algorithms Have complex fault-tolerant algorithms and operations been implemented? Quantum error correction requires spatial and temporal redundancy without reducing the efficiency of quantum computers. Information gleaned on levels B and C allow to project the size and temporal overhead of future quantum computers—this overhead is directly determined through the error rate of the underlying operations.

2.6 Evaluation of platforms

In order to evaluate the potential of different platforms, this study describes all known platforms for quantum computing and categorizes them into the above scheme.

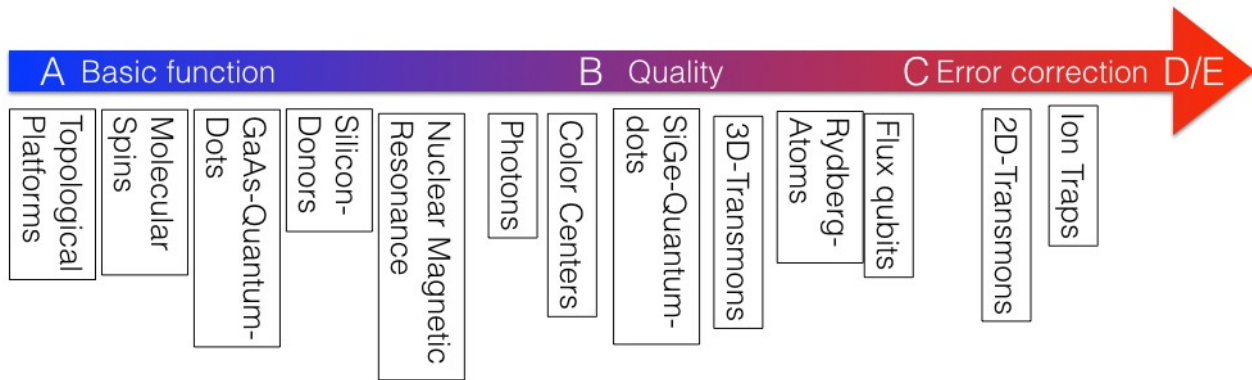


Figure 2.2: Evaluation of the main platforms following the developed scheme.

2.6.1 Trapped ions

This is an atomic platform, in which single ions are floating in ultra high vacuum held by slowly varying electric fields. It is a very well controlled quantum and well isolated quantum system. Research on trapped ions has already been applied, e.g., in metrology in the area of atomic clocks—an ideal starting point for low-error operation. The quantum information is stored in loosely bound outer electrons whose states can be manipulated through laser or microwave fields. Ions can be trapped in chains of mutually repelling objects and they can interact through vibrations in order to implement qubit-qubit logic operations. This is possible in very high quality. Further scaling requires to change from chains to complex two-dimensional arrays for which the electrostatic trap is implemented as a chip surface. Research in this area is impressive but the high operational quality of one-dimensional traps has not been reached yet. Alternatively, one-dimensional traps can communicate optically. Reaching this type of scaling is the biggest current challenge in trapped ions. All ingredients of a quantum processor and high operational quality along with simple error corrections have been demonstrated: Level C.

2.6.2 Superconducting circuits

This is a solid state platforms. It consists of integrated circuits made from superconducting metals and has hence to be operated at extremely low temperatures. Its key element are superconducting Josephson junctions. Their typical size is in the range of a micrometer or below—orders of magnitude larger than current semiconductor transistors. This basic technology is based, among others, on a superconducting measurement technology and its application in current and voltage standards—this again provides a good starting point for reaching high operational quality. Superconducting elements can be assembled into different quantum processor architectures, whose evolution has largely been driven by the requirement to maintain quantum coherence as a necessary ingredient for error avoidance. Next to the necessary cooling infrastructure (which is not an obstacle per se, but a complication) they have the control by microwaves in common. This platform is currently attracting the most industrial interest.

Flux qubits are superconducting loops in which logical states are represented by circulating currents. They resemble classical superconducting electronics more than other architectures. In some cases, flux qubits can reach very long coherence times and they can be easily coupled. It is challenging to fabricate these qubits consistently and with predictable properties which makes realizing the gate model a challenge. Their

superior connectivity makes them the leading platform for adiabatic quantum computing. For gate based computing their are on level B—operational quality for two-qubit gates needs to be improved, everything else has been shown.

Two-dimensional (or planar) transmons are single Josephson resonators whose electromagnetic oscillations states carry the quantum information. This design is an evolutionary development from charge qubits. It allows coupling through microwave resonators. Planar transmons reach very long coherence and can be flexibly coupled. They are planar on a chip surface and so far, chains and simple networks have been demonstrated. Further integration requires to build control and read-out lines into the third dimensions. Planar transmons have demonstrated simple instances of error correction, so they are on level C.

Three dimensional transmons are resonators similar to their planar version, but they are surrounded by a superconducting cavity at all sides. This further increases coherence times, but also makes control more complicated and gates more slow. Together with relatively large size, this makes scaling a challenge. So far, level B has been reached. A variation of quantum error correction tailored to this platforms has shown elements of level C, but cannot be firmly extrapolated as existence and numerical value of a threshold for this type of error correction have not been shown.




Figure 2.3: The currently leading quantum computing platforms - microscopic perspective. Left. Josephson processor (image: Julian Kelly, Google); Linear array of 9 qubits (crosses) with nearest-neighbor coupling; bottom: control lines, top: readout lines; Right. Linear ion traps (image Jürgen Eschner, Saarland University): trap electrodes (rods) and lenses for laser irradiation to implement quantum logic; same setup with its vacuum apparatus.

2.6.3 Neutral Atoms

In contrast to ions charge neutral atoms cannot be trapped by electrical fields alone. However, trapping is possible with much weaker, light-induced forces. Especially Rydberg states—atomic states with huge outer shell radii—allow for long distance interactions. Relatively large arrays of qubits have been demonstrated, but operations still show a lack of fidelity.

A bigger challenge is to stabilize such processor long enough—qubits are getting lost due to weak coupling. Error correction can absorb qubit losses, but is penalized with a massive overhead. Trapped atoms and atomic clouds remain important platforms for quantum metrology and quantum simulations of lattice models.

Cavity-QED is a platform using neutral atoms without trapping them. Beams of atoms cross resonators and interact with trapped photons. This is a powerful platform to demonstrate basic quantum effects and points

to an interface with quantum communication. The complexity to produce such cavities and their size prohibit research to pursue them as quantum computing platform.

2.6.4 Semiconductors

Semiconductor technology as an industrially relevant, spectacularly miniaturized and highly integrated has a strong potential for quantum computer development. There is currently a variety of semiconductor platforms—we describe the currently most interesting ones in this synopsis.

Semiconductor quantum dots are small isolated areas, “artificial atoms”, in which single electrons can be trapped so their spin degree of freedom can be used as quantum bit. Multi-qubit logics can be realized with interactions similar to those in magnetic materials. This platform has operational similarity with superconducting circuits. For a long time, they were held back by material-induced decoherence, which recently has been overcome by changing material to Silicon. During this study, all basic requirements for quantum processors were reached, but performance is still limited due to charge noise—Level B although progress in 2018 has been impressive. Current interest of semiconductor research institutes and companies promises that this platform can advance quickly.

Color centers are isolated defects in artificial diamond. They can be used similar to trapped ions, where the diamond crystal acts as a trap. These defects carry a nuclear and an electronic degree of freedom, i.e., a center potentially contains two qubits and in some cases up to four. Color centers are leading in quantum sensing and an important platform in quantum photonics. Targeted, scalable implantation of single centers is currently still very unreliable, which forces optical networking as the most likely strategy for scaling—which practically has large overhead (level A). Should this be overcome without compromising the high quality of single color centers it is expected that this platform quickly catches up to level C.

Single donors in Silicon have shown excellent single-qubit properties, reaching good two-qubit operations is the current frontier—level A.

Topological qubits are systems whose discrete quantum character originates from topological properties. They can be realized on a variety of platforms. Currently, semiconductor nanowires are pursued as a promising candidate. The motivation for the use of this platform is that the core concept of topological error protection—which is sensitive only to extremely unlikely correlated errors—is implemented already on the level of the material, not during quantum error correction. Alas, not even the existence of these qubits is established beyond doubt (although evidence is mounting), this platform clearly is on level A. If level A is mastered and if theoretical predictions are applicable, this platform could reach level D rather quickly.

2.6.5 Photonic platforms

Light can not only be used as a control and communication channel for quantum computers but also host quantum information. A number of important ingredients for quantum photonics has been developed in neighboring areas. Its key challenge is the implementation of two-qubit-gates, given that quanta of light (photons) do not interact. A number of indirect strategies can simulate this interaction, such as the use of special media or measurement and post-processing. So far, the gate qualities of two-qubit gates are not high enough for error correction (level B). Beyond this, there is a challenge to build scalable systems with a compact footprint given the large speed of light—but integrated optics appears to provide a solution.

This peculiar balance of resources has led to a range of alternative quantum computing protocols such as one-way or continuous variable quantum computing, which are better adapted to the physical situation of this platform.

2.6.6 Molecular platforms

The molecular platform of liquid state nuclear magnetic resonance (NMR) has historically been a first candidate for quantum computing. It is based on decades of experience in NMR in chemistry and medical diagnostics. Despite otherwise great parameters, there is currently no known method to scalably and robustly initialize such a quantum computer, which is why it has been given up as a candidate for scaling.

Other molecular platforms, molecular cages and magnetic molecules, are still in their early stages of development.

2.6.7 Overview

1. Topological platforms and molecular platforms meet few of the level A criteria, GaAs quantum dots, Si-Donors and nuclear magnetic resonance satisfy most but not all of them.
2. SiGe quantum dots, Rydberg atoms and two of the superconducting platforms meet all basic criteria, but the error rate is still far above the threshold for quantum error correction
3. One superconducting platform (2D transmons) and trapped ions are at or slightly below the error threshold—quantum error correction in simple instances has been demonstrated.

Based on the fast development of a few platforms, it is expected that both the absolute evaluation and the relative ranking can still change significantly, as there are in many cases only a few obstacles to be overcome. This has just happened in SiGe quantum dots. On the other hand, a few platforms such as liquid state NMR are stagnating for many years already.

The leading platforms, planar transmons and ion traps, fall in category C, so extrapolation on the development potential of quantum processors are difficult, but a scalable quantum computer is not imminent. On level C, a lot of scientific challenges have been identified or mastered and many next steps are of technological nature. As many technology-focused actors such as companies are entering the arena at this level, one can expect that the development in this area will accelerate. This means, that with a concerted effort, scalable quantum computers can be reached.




Figure 2.4: Comparison of algorithmic demands with currently achieved hardware performance. The plot shows required resources as number of qubits times rounds of error correction in the surface code for dlog (blue) and factoring (orange) for common key sizes (details in Chapter 9.5 and Figure 9.1) as a function of the physical error rate p . The squares show current realizations assuming one day run time (solid) or 100 days (empty), the yellow area shows expected near-term progress. Both scales are logarithmic.

2.7 Extrapolation

2.7.1 Status in 2018

Currently realized quantum processors are many orders of magnitude away from cryptography attacks, see image. Given the fast development in the field, there is accelerated progress towards this goal. This makes extrapolation challenging, but here is an attempt.

Scaling as described above is mostly controlled by the need for error correction. Also, it turns out that the most time-consuming step within error correction allows a time-space trade-off. Thus, some of the extrapolation depends on the time allocated to an attack on cryptography: While attacks in one hour are enormously far on the horizon, attacks in 100–1000 hours require much more spatial overhead. The size of the cryptosystem under attack is less relevant—the advantageous scaling of Shor attacks is preserved within its fault-tolerant implementation.

2.7.2 Update 2020

These 2018 milestones have been met, albeit with some delay. Notably, the number of qubits has been increased with a nearly constant error rate and the associated operational challenges have been met. The results show that error probability is currently the performance-limiting factor, which is expected to be the focus of research for the next couple of years, rather than qubit number.




Figure 2.5: Infrastructure units for quantum computers in leading platforms. Left: Dilution cryostat optimized for large cooling power and large wire-count for the operation of Josephson qubits (opened); qubits and other electronics units are mounted on the copper plates on different temperatures, the rack on the left contains control electronics. (Image: Edward Leonard jr., University of Wisconsin-Madison); Right: Parts of a vibration-controlled optical table containing two vacuum chambers for separate ion traps (Image: Jürgen Eschner, Saarland University).

Currently, such a quantum computer would be, even with an optimistic view of the near-term progress, a major piece of research infrastructure—such as a soccer-field size hall with vibration-controlled optical tables or a large arrays of cryostats containing the scarce isotope ^3He .

It is an interesting exercise to extrapolate, what a concerted research program for building a quantum computer could reach within the foreseeable future. With “concerted program” we mean that an industrialized nation pools a lot of its research and development effort into such a project, comparable with the Apollo and Manhattan programs in the US. Assuming that the current technical challenges are met—somewhat better operations, sparse use of voluminous periphery, larger chip areas, inter-chip connects and upgrades to cryogenic technology—it seems to be possible to have a computer with a Million planar transmons and a physical error rate of 1:10000. This would allow to attack 2048 Bit RSA in a few hundred days. A faster attack (in one day) would require to connect up to 1000 such units. This would require new technological solutions to connect these units—which have been demonstrated but currently would be too slow. Also, the initial filling of these machines with Helium 3 would require roughly the full annual industrial demand of Helium 3, likely requiring new nuclear facilities to produce this isotope. The financial and human investment in such an effort would be by far larger than current efforts in quantum computing.

An analogous activity in ion traps would require to bring the currently developed scalable trap technology to the same quality as linear traps. If successful, building the required quantum processor occupying roughly a soccer field would again require a concerted program.

2.8 Alternative platforms

A number of quantum computing platforms and protocols cannot be evaluated by direct application of our scheme. Most prominently, adiabatic quantum computing/quantum annealing is less technologically demanding than the fault-tolerant model. In particular, Canadian company D-Wave Systems offers such a machine and serves the international hype machine. At the moment, there is no proof of quantum acceleration of this machine and it is unlikely. Thus, the direct impact of this machine on cryptography is low, but these machines are interesting test-beds for quantum technologies.

A further direction are efforts, to demonstrate quantum supremacy with error correction. Application of the first generation of these processors is expected to be the simulation of quantum physics. For cryptanalysis this is merely an early milestones.

The third direction are optimized photonic protocols that are making optimal use of the capabilities of these platforms—specifically continuous variables and cluster states. These do reach impressive coherence and qubit numbers, but integration of all required functionalities is currently between levels A and B.

2.9 Global activities and potential for development

Quantum computing is progressing fast. Traditionally, this area has been sponsored by the funding agency of the US military and intelligence community (IARPA, ARO, DARPA). There, one can perceive increased focus on very few leading platforms and larger research teams, as well as an increasing role of government laboratories.

The engineering challenges starting on level C at the latest go, in most places, beyond the capabilities of university research. It is thus all the more important that laboratories outside universities and companies enter the field, which are currently driving progress in particular for Josephson Qubits. These are established large technology corporations (IBM, Google) as well as well-capitalized startups and SMEs (D-Wave, Rigetti) and a range of small companies. There is some, but much less, business interest in other platforms, with a notable startup in Level C. This should however not lead to the conclusion that the technological challenges for ion traps cannot be mastered - but industry is less experienced in integrating such systems.

Significant investment of Microsoft and Intel goes into semiconductor platforms. This is still basic research on levels A and B right now.

There are notable government investments in quantum computers in a few countries. Australia continuously invests in semiconductor platforms. The EU has launched a flagship initiative for quantum technologies, one of which is quantum computing. This is accompanied by large national programs.

These have in common that they typically do not directly aim at cryptanalysis, but in many cases a universal, fault tolerant quantum computer is the long-term goal, which can be used for this application.

2.10 Risks

This study reflects the current state of knowledge and expects continuous progress. There can be disruptive discoveries that would dramatically change its evaluation, the main ones being cryptographic algorithms that can be run on near-term non error-corrected machines or dramatic breakthroughs in the error rate of some platforms.

2.11 Update Spring 2019

The spring 2019 Edition updates both content and presentation.

We are discussing two algorithmic innovations that do not require active error correction, building on progress in other quantum computing applications: Variational factoring (9.2) and factoring on adiabatic quantum computers (9.1). Both are valid quantum algorithms whose speedup relative to classical factoring, however, is not proven and improbable.

The rest of the study thus focuses on fault tolerant quantum computing. On that end, no structural (asymptotic) breakthroughs have become available since the first version of this report. There have been algorithmic improvements, mostly described in (9.6), which impact quantitative estimates, however, e.g., to factor an RSA modulus, but these improvements do not put in question the overall magnitude of cost estimates for quantum attacks. Additional such improvements by "small constant factors" over the course of the coming months would not come as a surprise. In combination with progress in fault tolerance and error correction, an improvement of quantum attack costs by a moderate constant factor might become realistic. Similarly, for a key search against symmetric primitives, further moderate improvements through the identification of more efficient quantum circuits are to be expected.

In the area of quantum error correction, we are discussing promising early results for the reduction of overhead by lattice surgery, section (7.2.5). The improvement through this approach is not quantified yet in a scalable manner and does thus not influence the final extrapolation. A spectacular result with a very high error threshold for the case of asymmetric errors is discussed in chapter 7.4.3 but considered preliminary as the methodology is not scalable.

The discussion of the role of unitary errors has been greatly extended, chapter 7.4.1.3. On the one hand, more precise literature has appeared over the last year, on the other hand we have added a clearer discussion delineating the different error models. It turns out that the general results of fault tolerance are not different from unitary as compared to random errors, but the quantitative performance can be influenced if the errors are predominantly unitary. It is not yet clear what that quantitative change is.

The peculiarities and assumptions underlying the error models are now bundled in the central chapter 7.4.1 which also allows it to more clearly identify the assumptions of extrapolations for current hardware in chapters 13.2.3 and 14.2.3 vs. the extrapolation on future hardware in chapter 21. To improve readability involving the many technical terms in error correction, we have inserted an error correction glossary, chapter 7.7. Technical details in chapter 22 are greatly expanded.

The hardware-specific performance tables were updated. Major differences can be identified in semiconductor platforms, chapter 15. These are still on level „B“, but the error rates get closer to the regime where error correction can be attempted.

2.12 Update Spring 2020

On the algorithmic side, progress in Grover's algorithm for exhaustive key search is documented in section 9.5.1. Clifford+T gate resource estimates were updated accordingly to reflect improved implementations of the AES S-box as a quantum circuit. On the symmetric side, there has also been progress in leveraging Simon's algorithm, and we address this briefly in Section 9.5.4.

For factoring RSA moduli and computing discrete logarithms in prime fields, work by Gidney and Ekerå is taken into account now. Sections 9.6 and 9.7 have been updated accordingly. In the discussion on computing discrete logarithms on elliptic curves over prime fields, we incorporated recent work by Häner et al., which uses a large number of measurements as part of the calculation and thereby reduces the number

of T-gates. Moreover, in Section 9.6.3 we added a brief discussion of Harrow et al.'s (and follow-up) work on solving linear systems of equations.

During the update period, Google has published their celebrated result on quantum supremacy. In a wider context, activities in quantum computing research and commercialization have greatly expanded. These are focused on non fault-tolerant (NISQ) activities. Consequently, as outlined in chapter 4, these results do not have an immediate impact on cryptanalysis and are thus mostly influencing the second level of our evaluation scheme, yet the technological developments will certainly also be useful in order to pursue error correction and fault tolerance.

We are describing the Google experiment in section 6.6 along with the underlying method of benchmarking in section 6.4.5. We are mentioning an experimental verification of error correction in the end of section 22.3. Based on NISQ, new quantum heuristics for optimization have emerged, including diabatic quantum annealing. A method to decrypt lattice-based protocols on this platform has been presented. Yet, following its character as an early heuristic, there is no convincing indication for real quantum speedup, see chapter 8.1.3.

Next to these new sections, we have updated discussions around error models and phenomenology, consolidating text that also includes materials from the appendices into section 6.5 as well as some of the group descriptions.

Part I: Evaluating and benchmarking quantum computing hardware

3 Introduction to computational models and their requirements

3.1 Structure and requirements of an evaluation system

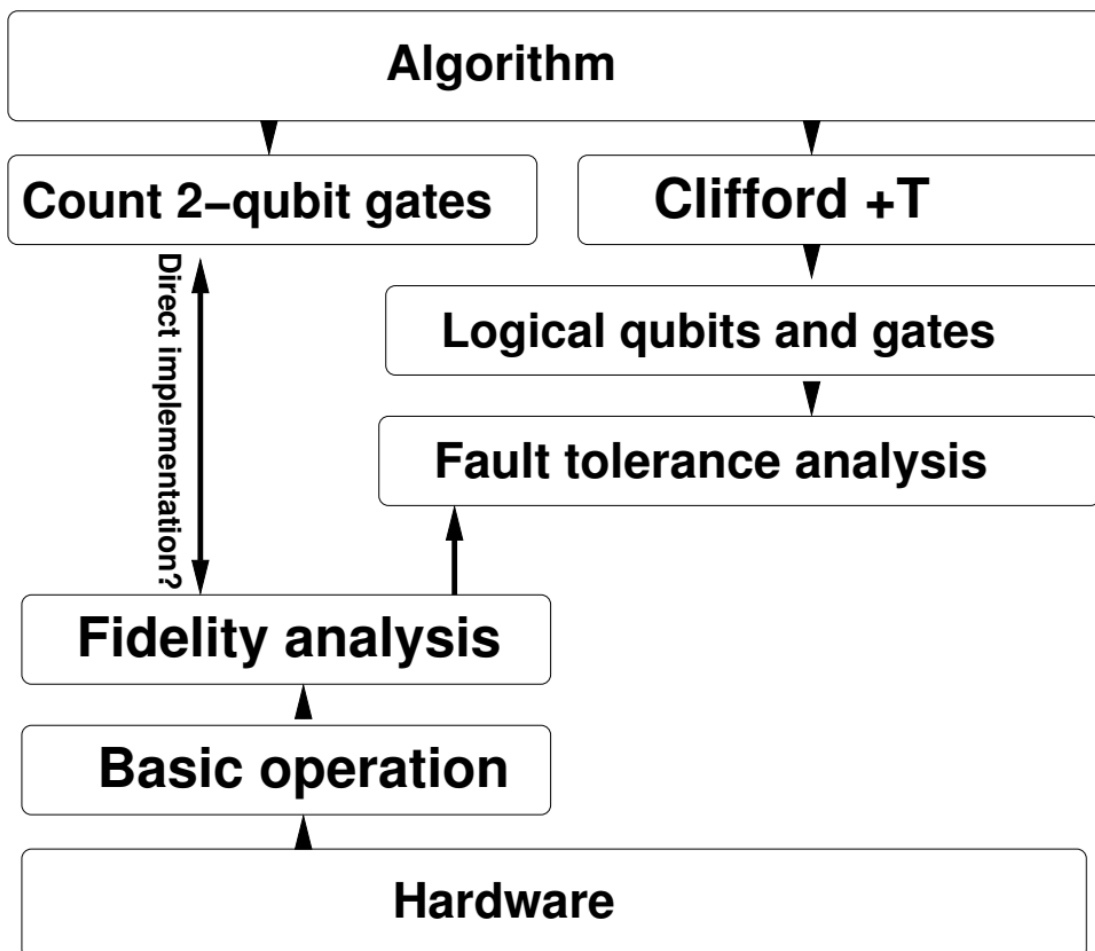


Figure 3.1: sketch of interdependencies and connections of our evaluation scheme. Hardware needs to pass checkpoints from below, software is compiled from above.

A quantum computer is a complex piece of technology that needs to function on many levels. Its basic components—qubits—are intricate physical objects based on pushing some experimental modality to its extremes. On the other hand, while obviously not as huge as some modern classical applications due to the power of quantum computing, algorithms running on it are already complex pieces of code. These elements are connected by a stack with intermediate level of detail. The challenge of evaluating where construction of a quantum computer stands is essentially an exercise in evaluating the machine on all of these levels and connecting them. This situation is not unlike the construction of a classical computer, which no person can do on their own—somebody with expertise in semiconductor lithography is hardly able to develop an operating system and vice versa. Also there, different levels communicate by standards they have to mutually meet and that build onto each other.

This principle drives the scheme we propose here. It is constructed *bottom up* in the sense that low-level requirements need to be met before high-level discussions even make sense. We feel that this is important,

in particular given some types of engagement in quantum computing research trigger hyperbole and press releases that often highlight advantages on one level only while omitting failure on other levels. We first divide the topic up into three categories:

1. Fault tolerant gate-based quantum computing—the main stream of quantum computing research and the centerpiece of this report.
2. Non fault-tolerant quantum computing—quantum computing that needs a limited number of gates but still outperforms classical computers—we will indicate why this has no reference for cryptanalysis in the foreseeable future
3. Nonstandard protocols like quantum annealing and cluster states—we evaluate these separately in Chapter 8.

3.2 Basic principles and notions of evaluating fault tolerant gate-based quantum computing

Before introducing the core level of this evaluation scheme, we would like to provide a big picture overview and background for fault tolerance.

Reliably executing a quantum algorithm requires to run it at a fixed and usable error rate of the binary input and output of the algorithm. By *fixed* we mean that the error rate does not grow with a longer algorithm, by *usable* we mean that a small number of runs of the algorithms should lead to an acceptable result. The complexity of quantum algorithms that outperform classical supercomputers requires a large number of gates, hence the number of logical gates needs to be small as we will quantify. Now in classical computers, where data encoding is strictly binary even during the computation and energy barriers lower the error rates to negligible values, this can be reached in hardware. Quantum computers, however, while certainly operating in a binary data space and accepting simple binary data as in- and output, use superpositions and entangled states during the computation that are fragile to continuous errors. Similar to classical computers, there is a measurement of binary registers in the end where any accumulated errors will result in a probability of the wrong outcome. As there is no self-correcting energy barrier for quantum computing (we will discuss a topological barrier below), intrinsic error rates of physical qubits cannot be expected to ever be as low as required by algorithms, so one stands before the challenge of executing an algorithm with faulty hardware.

This can be addressed with *fault-tolerant quantum computing*. Fault tolerant quantum computing draws a distinction between the faulty *physical qubits* that are used in a laboratory and the low-error *logical qubits* in which an algorithm is written down. Logical qubits are redundantly encoded in physical qubits, which are steadily error-corrected, giving logical qubits lower error rates than its physical components. We will make sure that we clearly label—often by chapter—whether physical or logical qubits are addressed. We will analogously talk about physical and logical operations depending on whether these are operations on physical or logical qubits.

These two layers are connected by fault-tolerant quantum computing protocols. These have been developed for more than two decades and their basic ideas are written in textbooks [NC00]. The efficiency of these techniques has been dramatically improved by the introduction of the *surface code*, an error-correction scheme that uses topological ideas to protect data—only errors that change topological properties of a state are not noticed. Note that topological qubits (see section 15.6) use these ideas on an elementary physical level, whereas the surface code is assembled from devices. We will describe basic notions of fault-tolerant computation in the introduction of chapter 7 but already highlight that its main ingredients are: i) Error syndrome extraction (and in some cases correction) in a stabilized code state which includes reducing analog error probabilities to digital errors, ii) storage of logical qubits which in parts requires temporarily turning off error correction in the stabilized state iii) implementation of “easy” logical operations, those that can be performed within the error correction code (typically the full set of Clifford gates) and iv) implementation of the remaining gates for forming a physical gate set, typically the T-gate.

Now these operations generally introduce a large overhead—a logical gate requires repeated error correction and generally consists of many physical operations on many physical qubits, all of which are in general faulty. For a well-designed code, there is a *threshold theorem* stating that under generic assumptions of the error model, the logical error rate can be made arbitrary small with finite overhead as long as physical error rates are below a certain threshold.

Given that the T-gate and other external gates are much more difficult to implement than the rest, one would like to stick with a minimal set—the T-gate alone. Hence, any quantum algorithm needs to be broken down to Clifford+T gates. Now note that the Solovay-Kitaev theorem states that any gate can be efficiently (logarithmic gate count in the required precision) approximated. This compiled count of Clifford + T gates (with some refinements outlined below) together with the target logical error per gate and the physical error per gate can be put together in a resource estimate for the overhead needed to implement a quantum algorithm in given hardware.

We now outline the structure of the scheme in a preview that highlights how its different components work together.

3.2.1 Lowest level (A): Basic operation—do we have working qubits?

At the lowest hardware level, physical modalities encoding qubits can and will be vastly different. In order to make them upwardly compatible, they need to function as qubits in the broadest sense. Here, we would like to know if all basic functionalities that allow to even think about running a low-level error correction algorithms are present. hence, all operations here are deemed physical, not logical. The criteria we propose are an extension and quantification of the famous DiVincenzo criteria. Platforms passing this test quantitatively will typically be able to demonstrate some basic quantum algorithms with 2–5 qubits. There are high-potential platforms like spins in semiconductors that are just passing this lowest level.

3.2.2 Intermediate level (B): Benchmarking—does our hardware meet fault tolerance criteria?

Once basic qubits functionality is established, it is important to quantitatively evaluate the performance of given hardware in a matter that is compatible with fault tolerance—but largely agnostic to hardware. Still, all operations discussed here are physical operations. Hardware may drive the choice of computational model (circuit based, adiabatic, cluster states) and fault tolerance scheme (surface or color code) but performance needs to be quantified in a way that is compatible with the analysis of fault tolerance. These numbers are essentially some qualitative statements about the architecture (how many operations can be parallelized? Can measurement be used as qubit reset?) but boils down to *fidelity measures* of the basic operations in fault tolerant computation—initialization, gate operations, and readout. Extracting these numbers is currently at the cutting edge of research as this is the level where the leading implementations of quantum computers currently stand. As reliable estimation of these parameters requires a quantum processor with some basic functionality, in particular faithful measurement and the ability to run at least in principle a long pulse sequence, it is important that processors have passed the level A in order to make meaningful statements. Passing level B is not only required in order to proceed to level C in general, benchmarking also determines the design parameters of the fault tolerance algorithms

Now there is a possible (but so far inconceivable—see below) outcome that no error correction is required because gate fidelities are large enough to execute a full algorithm. This is based on the count of two-qubit gates in the algorithm which consistently are limiting the total fidelity.

3.2.3 Central element (C): Fault tolerance analysis—how much quantum volume can we execute?

Once fault tolerance criteria are met from the intermediate level, it is known that adding more error correction (i.e., larger codewords, larger code distances) will reduce the logical error rate. Thus, with information from the next higher level (number of logical qubits and logical gate count) as well as below (architectural constraints and operation fidelities of the physical qubits) we can estimate the number of physical qubits and the time to execute an algorithm on given hardware thus estimate the *effective physical size* of the quantum computer that can execute the *effective logical volume* of the algorithm of interest. We describe its principles along the main technique, the surface code. We provide concrete numbers allowing physical resource estimates.

3.2.4 Compiled level (D): Elementary, fault-tolerant gates

Transitioning to the software layer, algorithms need to be broken down into elementary gates on logical qubits. The gate set of interest depends on whether active error correction is required (as is assumed throughout). In that case, transversal gates—gates that can be executed without decoding or any other construction outside the code space—are easy to implement—in the best current codes these are all the Clifford gates. Executing a general quantum algorithm that *cannot* be classically simulated requires at least one non-Clifford gate that needs to be produced outside the code. As this is for all known examples the by far most resource-intensive step (and needs to be designed for all individual non-Clifford gates) one restricts the gate sets to Clifford plus a single non-Clifford gate, typically the T gate (a phase shift of $\pi/4$ on one of the two basis states). Given these requirements, algorithms that need to be executed fault-tolerantly are broken down into Clifford+ T , i.e., gate counts for both Clifford and T gates are given.

3.2.5 Algorithmic level (E): Fault-tolerant algorithms

In a first step, cryptanalytic algorithms are commonly formulated at a high abstraction level. Details of implementing the necessary arithmetic, e.g., on an elliptic curve, or how to perform the round function of a block cipher with a superposition of inputs are not considered. To bring a quantum computer to use, the portions of the algorithm that cannot be run on classical hardware need to be identified, and design decisions on how to map abstract operations onto the available hardware need to be made. Just as with classical implementations, different algorithmic choices are possible, e.g., for computing an inversion modulo a prime number or for implementing an S-box. Different optimizations can be pursued—like minimizing the number of logical qubits or reducing the circuit depth of a computation. Cryptanalytic algorithms tend to involve complex operations, and as long as reliable libraries for elementary tasks are lacking, it seems prudent to organize the algorithm at hand in such a way that debugging remains feasible when passing to the gate level.

3.2.6 Conclusions and application

Once this is determined and we have a firm understanding of size, we can estimate extensive operational parameters (volume, heat dissipation, power consumption, amount of rare substances etc.) Moreover, we can evaluate if scaling up requires hitherto non-existing technologies, for example if multiple experimental infrastructural units need to be connected (multiple cryostats, multiple optical tables, multiple UHV systems). This will give an assessment of the feasibility and scope of building such a machine. Criteria will be laid out in chapter 11.

4 Are recent quantum supremacy proposals relevant for cryptanalysis?

4.1 Fault tolerance vs. near-term quantum supremacy

The enormous computational power of quantum computers comes with the drawback of also having a wealth of error mechanisms. On the one hand, these are based on the analog character of data being stored, on the other hand on the exponential capacity of quantum computers that allows for much more places for errors to occur. The standard remedy to this is some kind of quantum error correction that is used for fault-tolerant computation. Quantum error correction makes the key prediction that if the error rate is lower than some threshold p_c the logical error rate can be reduced arbitrarily by introducing overhead. The use of this standard paradigm will be the core of the evaluation system in Part I.

Recently, an alternative route for quantum computer applications emerged. It was asked, under what minimal requirements it is possible to outperform a current classical supercomputer using a quantum computer, i.e., at what point one can expect to reach *quantum supremacy* [Pre12]. As a benchmark, current supercomputers can maximally simulate the time evolution of about 50 qubits ([RMR⁺07,SSAG17], more recent [HS17]). One would first expect these to be *logical* qubits in the sense of error correction. On the other hand, given that even the best (surface) code has a threshold of $p_c \simeq 1\%$ and that this means that for manageable overhead the community aims for operating at no more than $p \simeq 0.1p_c \simeq 10^{-3}$ one can ask whether quantum supremacy can be reached without error correction, by merely executing about $p^{-1} \simeq 10^3$ gates. Surprisingly, the answer is positive ([BIS⁺16]). It has been shown that such a system can simulate quantum chaos in an exponentially large dimension, and that simulating quantum chaos, specifically sampling from the Porter-Thomas distribution is likely an NP-hard problem. This has now been experimentally demonstrated by Google [AAM⁺19]. Similar results have been obtained for variational quantum simulation [OBK⁺16] (Aspuru-Guzik, Dallaire-Demers [DDW16], Troyer [BWM⁺16]), where the quantum advantage comes from the need to store a complex quantum state (i.e., problems that on classical computers are memory-limited). This is a main driver of near-term quantum computer development.

4.2 Algorithmic perspective

We thus asked ourselves the question whether the same approach is relevant for cryptanalysis, whether a quantum computer whose gate set is *not* restricted to Clifford+T and which can execute around p^{-1} two-qubit gates with unlikely error can in fact attack any meaningful cryptographic code. We assumed that $p \simeq 10^{-5}$, i.e., far below threshold and two orders of magnitude lower than the current world record. We do not expect this to be beaten in the foreseeable future: It is very hard to even characterize operations to that precision and it is not critical in fault tolerance: Once the threshold has been passed, the qubit overhead scales only logarithmically in the physical error.

A promising direction to watch in this area is the direct implementation of Toffoli gates in hardware. These have been demonstrated in ion traps [FML⁺17]. The observed error rates are not disruptive and not affecting our conclusion, but deserve further attention.

Quantitative Analysis

For low-depth circuits with few gates, one could in principle consider a scenario without error correction, based on the idea that imperfections in the experimental realization might not significantly reduce the success probability. For cryptanalytic applications, the literature does at this point not offer quantum circuits that have been shown to meet these criteria. Low-depth solutions have been considered for solving the discrete logarithm problem on particular elliptic curves [RS14], but this comes at the cost of a large number of gates and qubits. In a similar vain, Cleve and Watrous [CW00] show that the Quantum Fourier

Transform (QFT)—which is at the heart of Shor’s algorithms—can be realized in logarithmic depth, but for the number of gates needed, only a polynomial bound is available.

Still, in view of the overhead incurred by error correction (cf. Section 7.5), one may ask if we might be able to tolerate errors on the gate-level without impeding the logical correctness of a cryptanalytic algorithm. Indeed, Nam and Blümel (see [Nam17,NB15b, NB15a]) make a case that a QFT implementation can perform very well even in the presence of noise and gate defects – thus suggesting that if the QFT is performed at the end of Shor’s algorithm, one could try to be lenient with error correction. One may also hope to simplify the QFT by passing to an approximate QFT (see [Cop94]), but for state-of-the-art implementations of Shor’s algorithms the logical gate cost is dominated by the arithmetic portion (see Section 9.2). State-of-the-art implementations of Shor such as [RNSL17c] save qubits by using a semi-classical QFT variant, with repeated (single qubit) measurements, where the rotations needed are chosen adaptively (in dependence on measurement outcomes so far), and savings/avoidance of error correction in the arithmetic would be particularly valuable. One can expect that any “accidental error tolerance” of arithmetic operations will depend on specific algorithmic choices (e.g., how exactly is a modular multiplication implemented, or how exactly is a point addition on an elliptic curve realized?). A common approach for the arithmetic tasks is to start with a reversible circuit which is then further decomposed into Clifford and T -gates—resulting in various options, e.g., to decompose a Toffoli gate (see [AMMR13,Jon13]). Having said this, there is very limited literature on error tolerance of arithmetic in Shor’s algorithm. Notably, in [Nam17, Chapter 9], Nam considers an implementation of Shor’s algorithm for factoring in the presence of errors in the angles occurring in elementary gates used. Due to resource constraints, the simulations he reports are restricted to very small examples (Chapter 9 discusses a factorization of 21), and making extrapolations for the arithmetic in cryptographically relevant factorization problems from this limited data set seems problematic. In recent work [NB17] on working with imperfect gates, the question to what extent errors can be tolerated in a large-scale (cryptanalytic) computation still remains open. In [NB15a], one particular adder design is considered and identified as quite robust against gate errors, but it remains open to what extent this can simplify a full-scale implementation of Shor’s algorithm. Taking into account debugging considerations, implementing a Toffoli-based arithmetic (cf. [HRS17, RNSL17c]) may in fact be considered as preferable over a (QFT-based) adder design as considered in [NB15a].

Work predating Nam and Blümel’s on the robustness of Shor’s algorithm in the presence of errors is due to Devitt et al. [SJD06]. They look specifically at the quantum period finding (QPF) subroutine of Shor’s algorithm and explore if a more lax error bound than imposing a precision of about $1/(\text{depth} \times \#\text{qubits})$ can be achieved. To test this, they apply three different discrete errors (bit flip, phase flip, both) randomly to the QPF portion of Shor’s algorithm. Each number of errors was simulated 50 times for specific factorable numbers with a binary length L ranging from 5 to 10 (invoking $2L + 4$ qubits) to determine how many errors were allowable until the result was no longer useful. Their result suggest that for larger L , more errors were acceptable. For example, when $L = 5$, at most 15 errors were acceptable before the result was unrecognizable from random, but with $L = 8$, up to 40 errors could be allowed. However, even a single error for $L = 5$ reduced the probability of success to 0.34. These results suggest that the precision of $1/(\text{depth} \times \#\text{qubits})$ can be reduced to $p(L)/(\text{depth} \times \#\text{qubits})$ where $p(L)$ is monotonically increasing and at least linear in L . Devitt et al. note that the greatest benefit of these results is for small simulations of QPF where observing the quantum process is the goal and extensive quantum error correction may not be feasible. However for large factoring problems (such as attacking cryptographically relevant RSA parameters) extensive error correction will still be required since the overall size of the quantum algorithm grows much faster ($O(L^4)$) than this error rate.

Overall, the question of intrinsically fault-tolerant cryptanalytic algorithms is still open. The existing literature does not really provide a road map to avoid quantum error correction in a serious cryptanalytic application. With the current state of the literature, explicitly budgeting overhead for error correction in quantum circuits for cryptanalysis appears adequate.

This analysis notwithstanding, research into these non-error corrected devices describes the progress of the field in a way that is very informative for this study. Its performance characteristics inform the input to the fault tolerance analysis. We are thus going to describe the accomplishment of [AAM⁺19] in section 6.4.5.

5 Low-level analysis of qubit systems

5.1 Initial remarks

5.1.1 Scope and motivation

This is the lowest (in the sense of being closest to hardware) level of a cascaded evaluation system for quantum computing candidates. It talks about physical qubits and operations only. It contains parameters that are easily characterized experimentally and serve as a stepping-stone for mid-level evaluation schemes (see Chapter 6) that are in turn the basis for analyzing fault tolerance requirements (see Chapter 7).

Such a low-level scheme has been published a long time ago in the form of the DiVincenzo criteria [DiV00]. These criteria were giving a succinct summary of what it takes for a qubit candidate to be serious, mostly in order to help new and then-emerging (condensed-matter) platforms to evaluate themselves and ask the right questions. Notably, these criteria are not quantitative (which they do not have to be, only the next level should) but they do not even give suitable numbers to use. As the field has matured since then, this part of our survey explores these numbers as they are typically given in experimental papers. It also compares different quantifiers used in different experimental traditions and develops relations between them. We review the DiVincenzo criteria and then modern ways to clarify and quantify them.

For a large-scale analysis of quantum computing candidates, this serves as an entry ticket. If these criteria and parameters cannot be verified and measured satisfactorily, development of architectures and measurement of performance parameters that are relevant for fault tolerance are usually futile—they require a functional qubit to at least have some understanding what design operates under what condition. This is thus the lowest-level performance check for quantum computing platforms.

Notable special cases are adiabatic quantum computing/quantum annealing and cluster state quantum computing, which, although not fundamentally different, put different priorities on hardware and are thus not easily connected to these criteria and therefore need to be treated differently. We will describe how to evaluate them in a separate Section 8.1.

5.1.2 Limitations

The next level beyond these low criteria that will be the core of medium-level analysis in a further deliverable and will be referenced here is Randomized Benchmarking (RB) [KLR⁺08, RLL09, MGE12, ECMG14, MLS⁺15, XLM⁺15, ATB16]. It plays a connecting role as it is relatively easy to use experimentally given basic qubit functionality. It consists of preparing a convenient initial state, running a sequence of random Clifford gates, invert it by a single further Clifford gate (relying on the fact that these gates form a group that can be efficiently simulated classically) and measure the survival probability of this initial state. It can be shown that this maps out the average fidelity of the sequence and can hence be a reliable estimator for the error per gate. Usually, the survival probability does not extrapolate to unity thus capturing state preparation and measurement (SPAM) errors. Low-level performance indicators covered in this section are discussed up to the point where performing RB would be the more adequate choice. A detailed description of RB as well as its limitations is given in a subsequent deliverable.

5.2 The DiVincenzo criteria reviewed

The 5+2 criteria [DiV00] for quantum computation are:

1. a scalable physical system with well characterized qubits,

2. the ability to initialize the state of the qubits to a simple fiducial state
3. long relevant decoherence times, much longer than the gate operation time
4. a universal set of quantum gates
5. a qubit-specific measurement capability
6. the ability to interconvert stationary and flying qubits
7. the ability faithfully to transmit flying qubits between specified locations

A few initial remarks are in order

Well-characterized qubit array The requirements that the qubits are *well characterized* means that the physical parameters should be accurately known, including the internal Hamiltonian, couplings to other qubit states, interactions with other qubits and coupling with external fields. Higher qubits states should be avoided (leakage) so the physical qubits represent mathematical qubits—abstract two level systems. The proper identification of the qubit needs to be done carefully. Remedies to imprecise characterization can be found in robust control, which are however generally less efficient than controls for precisely characterized system.

Initialization The need for initialization arises from the straightforward computing requirement of known initialized registers. Now the evolution of a closed quantum system is unitary, hence invertible, whereas initialization is not invertible. Thus, initialization requires to open the quantum system operation to achieve, e.g., cooling or measurement. Initialization is also important for quantum error correction, where a continuous supply of fresh qubits for re-encoding is a real headache for many implementations. The speed of initialization is an important issue in experiments. The main approaches for initialization are cooling to the ground state of the Hamiltonian or projective measurements. Cooling works if the energy gap between the ground and first excited states of the quantum computer is much smaller than the temperature in appropriate units. In practice, it is hard to define that temperature in some cases—e.g., the effective temperature of a Josephson circuit is usually higher than the temperature of the surrounding Helium bath—which can be mitigated by making temperature margins wide enough. Unfortunately, natural cooling is on the same timescale as energy relaxation, which is just the bit flip error rate described below, posing a conundrum when using this method within error correction. This is mitigated if the relaxation rate can be switched or otherwise manipulated. In some optical approaches (ions and neutral atoms), where qubits are encoded in hyperfine states, relaxation is so slow that it needs to be manipulated by optical pumping: Selective excitation of one of the qubit state to a metastable excited state. An alternative approach to fast initialization is projective measurement and feed-forward correction, i.e., we measure the state and apply an additional gate depending on the measurement outcome [RvLK⁺12].

Coherence Error correction can be applied in quantum computation putting more reasonable requirements for decoherence times, for decoherence times of the order of 10^4 to 10^5 times the clock time of the quantum computer [DiV00] (see Section 7.3). Note that also errors other than decoherence enter that threshold, e.g., systematic errors such as gate axis misalignment or over- and under-rotation. The latter two errors are unitary errors where in the first case the rotation axis n' is tilted compared to the ideal axis n , and in the second case the rotation angle $\theta' = \theta \pm \epsilon$ is too large or too small. The former occur, for example, in gates driven by resonant radiation if the resonance condition is not met perfectly, the latter occurs based on errors of amplitude of the drive field or timing. A more detailed description of those is given in the Appendix 22.

Universal set of gates The universal set of gates is the heart of quantum computing. In principle, the desired Hamiltonians to perform quantum gates are turned on and off via external controls, with somewhat smooth pulse shapes. These have to address all the interactions that cannot be turned off, e.g. in NMR, i.e., in the presence of spurious coupling there is some control required to simply keep qubits or registers idle, typically in the form of refocusing operations. Refocusing consists of designing control sequences such that the impact of undesired term averages out in the end. Its simplest example is Hahn spin echo [Lev01,VC05]. It has been invented in original NMR and can be interpreted in quantum computing as a way to protect quantum memory from inhomogeneity. Turning off all couplings between the spins is known as

decoupling, and turning on specific couplings is called refocusing, and the latter can be done efficiently [LCYY00,VC05]. The drawback of these techniques is that they make gate sequences longer thus making operations more susceptible to unitary error. Refocusing is compatible with error correction, see Chapter 24.1. In the context of our evaluation scheme, experimentalists will decide whether or not to use refocusing for their operations and benchmark them accordingly on level B.

Also auxiliary systems are used for gate implementations, e.g., in ion traps, where direct interactions between qubits cannot be turned on. Also, fully parallel operations are needed for quantum error correction, which can be a problem when a single bus is used to mediate the interaction between arbitrary qubits, while nearest neighbor interactions allow for sufficient parallelism. Systematic errors due to imperfect gates should be below the error correction threshold [Pre97b, DiV00]—see measures for error rates described on level B in Section 6.4.1 and thresholds described on level C in Section 7.2.7. Note that we are talking about physical gates here that are meant to execute the operations underlying fault tolerance—logical error-corrected gates are treated later, in Chapter 7. The use of coding the qubit can reduce the gates required. A standard choice for physical qubits are single qubit rotations and a perfect two-qubit entangler [Mak02,ZVSW03] (often a CNOT). For *logical* qubits one typically relies on the minimal set of Clifford gates and the T gate, a $\pi/8$ phase shift with opposite signs for its basis states.

Measurement The qubit specific measurement capability is minimally needed to read out the result of the computation. If the measurement is an ideal quantum measurement (restrictions to this are described below), it can be used for fast state preparation, e.g., for recycling qubits in quantum error correction—but this is not necessary as quantum error correction can be done only with final measurements but other overhead. In threshold calculations, a single quantum-efficiency parameter is an often used to summarize the fidelity of a quantum measurement, whereas the reality is more complex. Improving the efficiency can be done by a “copy” of the single qubit to three, which is done by initializing two qubits to $|0\rangle$, applying CNOTs and measuring all of them [DiV00]. Also, perfect initialization of maximally entangled states in the form of cluster states leads to a protocol that only requires single qubit gates for computation, making measurement a resource for actual computation.

Communication-related criteria The last two criteria play a role for communication, i.e., transmission of qubits. Requirement (vii) is important for cryptography. Proposal for flying qubits usually assume photons as flying qubits, but also electrons traveling through solids. Potential candidates are described in the nonstandard architectures Section 8. They are important if quantum processors are used as or in quantum repeaters (not part of this study) or in distributed quantum computing (see Section 7.4.3).

5.3 Coherence time scales

Decoherence describes the process of the loss of quantum information through interaction with an environment. Its nomenclature is not unique throughout literature. In this review, we are going to describe both decoherence and energy relaxation in a unified language and do not discriminate that the former governs the quantum-to-classical transition whereas the latter can also occur in a purely classical system—after all, both are contributing to errors of the quantum algorithm so they both are of relevance for reaching the threshold. As a first characterization of qubits that implement the circuit model in time, any of these time scales should be much longer than a typical gate time.

We assume a qubit with some capability for single qubit gates. If that does not exist yet, coherence time scales are also related to spectroscopic line widths. The latter is proportional to the decay rate of the state through spontaneous emission.

5.3.1 Single-qubit level

The nomenclature of coherence times in quantum computing has largely been adapted from nuclear magnetic resonance (NMR), where they were introduced with the Bloch equation [Lev01,VC05]. This nomenclature assumes a preferred basis set by some static energy splitting that is larger than any time-

dependent controls. In most quantum computing architecture, this is also the basis in which the qubit states are encoded. A further assumption behind the Bloch equation is, that errors are Markovian, i.e., the noise process does not have any temporal memory. The Markov assumption implies that terms decay exponentially, hence an error occurring over a time T_e and within a gate time T_g leads to an error $1 - \exp(-T_g/T_e) = T_g/T_e + O((T_g/T_e)^2)$. Exceptions to this assumption are discussed later. Most of these rates can be estimated from the noise of the qubit environment, if known, using Fermi's golden rule [MKT⁺00,SW03,SHKW05].

In this framework, we identify the following time scales as being relevant:

Energy relaxation time T_1 The time T_1 describes the time of energy relaxation, i.e., bit flip errors. It is dominated by noise at the transition frequency of the qubit. Note that long T_1 can always be reached by inhibiting transitions of the qubit between its logical states (including coherent gates), hence on its own it alone is not a clear performance indicator. The standard experiments to get T_1 are inversion or saturation recovery [VC05], e.g., one prepares a non-stationary mixture of energy eigenstates and measures their decay time.

Phase coherence time T_2 and pure dephasing time T_ϕ The time T_2 describes the time of phase randomization, i.e., the time it takes to transfer a superposition of the qubit states into a statistical mixture. This is not independent of T_1 errors and in fact it can be shown that $T_2 \leq 2T_1$ based on the constraint that the qubit density matrix has to remain positive. The difference as a rate (inverse time) can be identified as the pure dephasing time T_ϕ as $T_2^{-1} = (2T_1)^{-1} + T_\phi^{-1}$. The rate T_ϕ^{-1} is proportional to the low-frequency energy fluctuations of the qubit. Formally, the relevant frequency is zero, however, practically this is set by one over the duration of the experiment. Applying this argument to $1/f$ -noise produces a short T_ϕ that formally diverges in a long experiment. While this formal divergence shows the limitations of this simple argument [MS04], see also our discussion in Subsection 5.3.3.2, this motivates that the impact of $1/f$ noise needs to be avoided. This can be achieved if the fluctuations do not impact qubit energy - which can be arranged, e.g., in Josephson qubits, by choosing an optimum working point [VAC+02,CW08]. T_2 can be measured by Ramsey interferometry that is corrected for homogeneous effects (see below) by some type of echo. Ramsey interferometry consists of preparing the qubit in an energy eigenstate, then performing a $\pi/2$ rotation into an equal superposition of eigenstates, wait for a time t , and repeat the $\pi/2$ rotation. The decay of the resulting signal shows the decay of a superposition hence directly gives T_2 .

Note that while it is intuitive that T_1 limits T_2 (energy relaxation through an environment breaks the phase), the factor 2 arising above has been subject to much argument. Its existence is well established and many experiments reach $T_2 \simeq 2T_1$ one should keep in mind that in the Bloch equations, T_2 appears twice and T_1 appears once—and that T_2 describes the decay of a probability *amplitude* whereas T_1 describes the decay of a probability.

Ensemble phase coherence time T_2^* Measurements of T_2 for example by Ramsey interferometry require collecting data from an ensemble as they are based on expectation values. This type of ensemble averages is collected on single quantum systems by repeating the experiment in time. Now in principle, the parameters of the experiment can slowly fluctuate between these ensemble members, i.e., between different runs of the experiment, which after performing the ensemble average looks like a short T_2 . In some realizations (in spin ensembles in NMR) the ensemble is built in one temporal run but inhomogeneity between ensemble members arises because of variations of the magnetic field across the test tube. This phenomenon is also called inhomogeneous broadening (from the broadening of the Fourier transform, i.e., the spectral line below saturation).

Inhomogeneous effects can be suppressed by the spin echo technique (the NMR Hartmann-Hahn echo can be viewed as stabilization of quantum memory). Logical operations need to incorporate echo in the form of composite pulses or robust controls, which are typically longer than uncompensated pulses. Experimental designs thus decide whether the savings of going from T_2^* to T_2 are overcompensated or not by the echo technique. Basic notions are described in Ref. [VC05] and its application to quantum computing is outlined in this study in Appendix 24.1.

Rotating frame decay time $T_{1\rho}$ In many cases it is useful to visualize qubits in three-dimensional affine space by plotting the expectation values of the three Pauli matrices $\sigma_{x/y/z}$ on the respective coordinate axis—the Bloch sphere. Pure qubit states are represented by points on the Bloch sphere, mixed states by points in its interior, the Bloch ball. In this representation, a basis change to a time dependent with continuously evolving phase factors can be visualized as a changing into a co-rotating frame, which is often very useful to understand and describe qubit dynamics. Specifically, most quantum computing platforms realize off-diagonal single-qubit gates by resonant external fields that are easily described by quasi-static terms in a frame rotating with that resonant field, and that drive Rabi oscillations. Moreover, based on a phenomenon called spin-locking, the relevant decoherence time for these gates is not T_2 but $T_{1\rho}$ which probes the environment at the Rabi frequency scale rather than at very low frequencies as T_ϕ would do. In particular, in systems with strong $1/f$ or other low-frequency noise this time can be much longer than T_2 , hence leading to more optimistic performance estimates for these gates [VC05].

5.3.2 Properties unique to multi-qubit noise

Qubit noise metrology becomes much more complex on the multi-qubit level. The commonly used mathematical structure to describe this is the Lindblad equation (of which the Bloch equation is a special case)—a master equation that describes general strictly memoryless (Markovian) and completely positive quantum dynamics. We will describe limitations to this method below. A novel component that needs to be considered is the question whether noises are correlated across quantum bits or whether they are separate and uncorrelated between qubits or whether they are correlated. On an operational level, correlated noise is less harmful and in some cases allows for decoherence-free subspaces [LW03] see also Appendix 24.2. These are in fact a guiding principle behind the design of single-triplet and triple-dot qubits in semiconductors. On the other hand, it is known that the increase of sensitivity to uncorrelated noise is a measure of entanglement, so the effective dephasing rate of a maximally entangled N -qubit state is N -times faster than the individual dephasing rates, making uncorrelated noise a worst-case scenario. This is taken into account in fault tolerance.

Discussion of noise correlations is mostly driven by noise modeling—when the primary source of noise is known one can assess their spatial correlations. This includes knowledge that the long-range of nuclear magnetic fields makes noise in GaAs mostly correlated, the same is true for anomalous heating in ion traps—it also includes knowledge that materials-induced noise in superconducting qubits is mostly uncorrelated.

Measurements of noise correlations are rare (example given for ultra cold atoms in [Föl14]) as they would require partial process tomography, see process tomography in Section 6.4.2. Rather, given that multi-qubit operations rest on the shoulders of coherent single qubits, they are obtained by an intermediate-level characterization method, specifically can be inferred from error budgets gleaned from RB. Randomized benchmarking methods have been realized in ion traps [GMT⁺12, HAB⁺14, MKC⁺15], at IBM [MGJ⁺12], in NMR [RLL09] and in semiconductors [MLS⁺15]. However, it is often assumed that the correlation in noise between qubits either is small or can be ignored in fault-tolerant estimates [MGE12]. The information contained in RB will be discussed in Section 6.4.3 and the precise nature of multi-qubit errors is described in Appendix 22.

5.3.3 Non-Markovian effects and other caveats

5.3.3.1 General observations

As discussed above, characterization of coherence decay in terms of exponential decay and single time scales relies on a number of assumptions. The most crucial of those is the Markov assumption—the assumption that temporal correlations of the environment are short-lived. This is a central assumption behind the description of decoherence in terms of the Lindblad equation [Lin76,BP02]. At first, this appears very unreasonable, given the low temperature most qubits operate at. Low-temperature operation is not an

experimental accident, it is often needed to avoid thermal noise, it is also needed to allow initialization into the ground state by thermalization. However, if done properly [WSHG06] it turns out that the environmental correlation time needs to be shorter only than the typical coherence decay time. This implies that for serious qubit candidates, where the latter is long, naturally can be described with Markovian decay pictures.

A few exception to this general equation observations need to be noted.

5.3.3.2 1/f noise and nuclear spin noise

Pink noise with a frequency spectrum that diverges roughly as $1/f$ at low frequency f are ubiquitous in condensed phases [VC76, DH81, Wei88, SMS02] leading to very slow correlation decay in the time domain. It turns out [ICJ'05, SMSS06] that coherence decay here is Gaussian $\propto \exp(-T^2/T_2^2)$ and still a time scale T_2 can be defined. Now note that this bounds a short-term error rate by $1 - \exp(-T^2/T_2^2) = T^2/T_2^2 + O(T^4/T_2^4)$ seeming lower. While this is established in single experimental runs, it is theoretically understood that this assumes starting from non-entangled qubit and environment and thus only applies to the first operation applied to a freshly initialized qubit.

Similarly, but with a much richer set of details, decoherence due to a nuclear spin bath in electron spin qubits (typically in GaAs) can be described. Nuclear spin baths are also intrinsically slow owing to the large nuclear mass compared to the electron mass. They cannot be described by their correlation function alone due to their localized nature and restricted spectrum. Still, their impact can be put into a time scale that can be gauged similar to T_2 [FTCL09].

5.3.3.3 Slippage and other non-Markovian affects

Another assumption of Markovian decoherence is that the initial conditions between qubit and heat bath are uncorrelated, which is rather artificial. It is known [SSO92, Wil08] that this mostly leads to short-time effects or even loss of initial visibility. These phenomena contribute to state preparation and measurement (SPAM) errors on the next higher level (see Chapter 6).

In general, the notion of non-Markovianity is lacking a standard model comparable to the Lindblad equations. Criteria to quantify it have been introduced [BLPV16] and are currently actively researched. Generally, these effects are subtle and occur only if they are not masked by Markovian effects, hence, systems that are affected by this type of non-Markovian decoherence can be analyzed with RB, the key tool on the next level.

5.4 Qubit definition indicators

5.4.1 Qubit longevity

For some platforms, the qubits themselves maybe short-lived, primarily in neutral atoms where trapping forces are weak. An efficient way of detection of atom loss is the 'knock-knock' protocol [Pre97a, Saf16], measuring the fluorescence signal for trap lifetimes [WL95] One can trap individual atoms with optical tweezers [BBL16]. These are tightly focused dipole traps, $\sim 1\mu\text{m}^3$ using aspherical lenses, such that at most one atom is trapped at a time. The atoms are precooled in a MOT and loaded into the tweezers through spatial overlap with the MOT. The inelastic two-body collisions omit loading a second atom into the tweezers. Real-time imaging is possible measuring the fluorescence photons of the trapped atom illuminated with near-resonant laser light. Thus, one can determine if an atom is present in the trap via the toggling fluorescence signal between high and low photon counts.

Other qubit longevity issues, e.g., from ultraslow (sub-millihertz) drifts that can occur in solid-state qubits can usually be mitigated by calibration.

5.4.2 Leakage

Mathematical qubits—systems that can be completely described as two-state quantum systems, do not exist in nature, not even as elementary particles [WL02]. Typically, the computational states are either one degree of freedom of an elementary particle (e.g. the spin-1/2 of a proton in NMR, which also possesses motional degrees of freedom) or they are taken as low-energy states of a more complex energy spectrum (e.g. in ions or Josephson circuits). In order to still operate these devices as qubits, one needs to guarantee that the state returns to the computational subspace (CSS) after operations (whereas non-computational states can in fact be useful in gate operations or for readout). Deviations from this are referred to as leakage. Leakage errors are particularly difficult to error-correct. In some platforms, leakage is not a problem—specifically when non-computationally states are far separated in energy from computational states. Nuclear motion in molecules, for example, has frequencies in the infrared range whereas spin dynamics is between radio frequency and the low end of the microwave spectrum. In some platforms, most notably Josephson qubits, leakage is an ever-present challenge as the energy splitting between qubit states differs from that to the leakage levels by typically only 10–20% in the transmon.

A first indicator for leakage resistance are these energy splittings. If they are critical, one can rely on a sophisticated array of detection techniques: Leakage measurement by IBM [WG17] introduce two new criteria, the leakage rate L_1 describing leakage from the computational states to other states, and the seepage rate L_2 , population transfer from other states to computational subspace. The latter rate introduces memory effects to the system. These two leakage rates also appear in [WBE16]. However, the protocols to measure leakage differ. The sum of the leakage errors in [WBE16] is measured via a protocol using a random sequence of length m of gates applied to a completely mixed state in the CSS, without a final inversion. At IBM measurement is done with leakage RB [WG17]. Also the incoherent leakage subspace is investigated for states that are an incoherent mixture of CSS and non-CSS states. Earlier work on measuring leakage done in the Martinis group with introduction of the Ramsey interference error filter [LHA*08, LKB*10], first applied on phase qubits: applying π -pulses to excite the leakage level and varying the time between a second π -pulse. Leakage is referred as amplitude and phase errors: amplitude error for population out of the CSS and phase errors due to interaction with leakage levels.

5.5 Qubit initialization indicators

If qubit initialization is done by cooling in the ground state, one can upper-bound its population by $1 - \exp(-\Delta E/k_B T)$ where ΔE is the energy gap to the first excited state and k_B is the Boltzmann constant. In initialization by measurement, the maximum initialization fidelity is limited by the projection fidelity of the measurement and the fidelity of the gate that needs to be applied to correct the measurement if necessary [RvLK*12, RBLD12]. In initialization via optical pumping [Saf16] a high contrast of rates is required for good initialization.

A posteriori, initialization can be measured by measuring right after initialization. As measurement is typically more restricted than initialization, this is not often done.

5.6 Readout indicators

Readout is a crucial part of quantum computers (and of quantum physics). It does not only serve the final analysis of the outcomes of the algorithms, but is pivotal in syndrome extraction for quantum error correction and thus an important ingredient of threshold calculations. Also, readout can influence architecture decisions.

Quantum measurement is probably the most intriguing part of quantum physics leading to a lot of foundational arguments. Also, quantum measurement science is related to precision, quantum-limited measurement—a lot of modern quantum measurement science and engineering has for example been originally driven by the application of gravitational wave detection[CDG⁺10, BSV01, DK12, BBV⁺16]. We will only touch to these two related tangents in a minimal way, to the extent that they are relevant to qubit measurement—for example because some critical element in the overall measurement chain.

5.6.1 Binary-outcome detectors

In a POVM (positive operator-valued measurement) [NC00, BP02], the probability to measure the value i is $P_i = \text{Tr}(M^{(i)}\rho M^{(i)*})$, and the state after the measurement reads $\rho_i = M^{(i)}\rho M^{(i)*}/P_i$, with the linear operators $M^{(i)}$. An ideal binary-outcome detector, $i \in \{0,1\}$, is described by ten real parameters, and a general binary-outcome detector by 28 real parameters [Kor08]. Therefore, the quantum efficiency of a general binary-outcome detector is describe by 18 parameters. This is impractical—most threshold calculations described the detector efficiency as a *single* parameter that can be measured with RB. Here, we outline basic notions of quantum measurement that are necessary for basic detector validation.

A *Quantum Non-Demolition* detector (QND) does not induce any random transition between the eigenstates of the measured operator, hence the post-measurement state is uniquely determined by the measurement outcome. This is a necessary requirement for the re-use of the measured state in error correction and the use for measurement in state initialization. QND-ness can be verified by repeating measurements. Practically, measurements are QND when the measured operator commutes with the qubit Hamiltonian, i.e., if one measures in the energy eigenbasis.

Restriction to QND detectors leads to a quantum efficiency characterization through six parameters only.

For an ideal QND detector the pure initial state $\alpha|0\rangle + \beta|1\rangle$ is transformed to $(\alpha c_0^{(i)}|0\rangle + \beta c_1^{(i)}|1\rangle)/\sqrt{P_i}$, with fidelities $F_i = |c_i^{(i)}|^2$, probabilities $P_i = |\alpha c_0^{(i)}|^2 + |\beta c_1^{(i)}|^2$ and phases $\phi_i = \arg(c_0^{(i)}) - \arg(c_1^{(i)})$. In general, a non-ideal QND detector transforms $\alpha|0\rangle + \beta|1\rangle$ to

$$\frac{1}{P_i} \begin{bmatrix} |c_0^{(i)} \alpha|^2 & c_0^{(i)} c_1^{(i)*} e^{-D_i} \alpha \beta^* \\ \text{c.c.} & |c_1^{(i)} \beta|^2 \end{bmatrix}$$

with decoherences D_i . Given the fidelities F_i , phases ϕ_i and decoherences D_i , the quantum efficiency is completely described.

The quantum efficiency is defined as $\eta = D_{\min}/D_{\text{avg}}$, with the lower bound D_{\min} of the ensemble decoherence D_{avg} . The ensemble decoherence is defined as $e^{-D_{\text{avg}}} e^{-\phi_{\text{avg}}} = \sum_i c_0^{(i)} c_1^{(i)*} e^{-D_i}$ and from $D_i \geq 0$ one gets $D_{\text{avg}} \geq D_{\min} = -\ln(\sum_i |c_0^{(i)} c_1^{(i)*}|)$. If $D_{\min} = 0$ then the measurement doesn't give any information about the qubit state. Outcome dependent efficiencies are defined as $\eta_i = D_{\min}/(D_i + D_{\min})$. A perfect measurements has $D \rightarrow \infty$. The coefficients $c_{0,1}^{(i)}$ describe a potential state-dependent unitary rotation executed by the classical back-action of the detector.

We now describe the application of these quantifiers to two paradigmatic measurement protocols

Indirect projective measurements The qubit interacts with an ancillary qubit (in $|0_a\rangle$), and a projective measurement is performed on the latter. The decoherences are $D_i = 0$ and the detector is ideal for the individual efficiencies $\eta_i = 1$, but $\eta = D_{\min}/[-\ln(\sum_i |c_0^{(i)} c_1^{(i)*}|)] < 1$ if the phases mismatch $\phi_0 \neq \phi_1$.

Linear detector in binary-outcome mode This is a common example for solid-state qubits. A linear detector is characterized through the average output signals I_0 and I_1 for the two measurement outcomes, with $\Delta I = I_1 - I_0$, and the spectral density S of the white noise. $\tau_m = 2S/\Delta I^2$ is the time needed for the signal-to-noise (SNR) ratio to reach 1. The quantum efficiency reads $\eta = D_{\min}/(\gamma t + s^2[1 + (\tau_m K \Delta I/2)^2])$, with the measurement strength $s = \sqrt{t/2\tau_m}$, the decoherence rate γ and the correlation between output and back-action noise K . Even for an ideal linear detector ($\gamma = 0, K = 0$) the efficiency $\eta_0 < 1$ is not ideal, and $\eta < 2/\pi$.

5.7 Final remarks

This low-key application of DiVincenzo's criteria is the first qualifier for quantum computing platforms. Passing them with structures containing at least a few qubits will enable a more quantitative performance discussion as that done in the next section, which essentially proposes small quantum algorithms that allow to extract quantitative performance indicators.

6 Benchmarking qubits

6.1 Introduction

When we have basically functioning qubits, we would like to know if they have the potential to be scalable, i.e., if they meet the threshold for quantum error correction. The question is if a device is given to us how to know if the threshold has been surpassed. This is nontrivial because one needs to know which measure returns the threshold. While the low-level criteria can give bounds on achievable errors, they crucially depend on very finicky models to be accurate and complete - models that address human-made systems and hence would need to be re-evaluated over and over. Thus, in order to validate qubits and refine these models, it is important to have a way to measure the error of quantum operations on a real qubit. Next to evaluating the distance to the error correction threshold, this type of characterization also helps to improve quantum processor elements and assists in calibration of operations.

6.2 Benchmarking and error mitigation techniques

Benchmarking can be used to evaluate low-level gate design and error mitigation techniques. Specifically, dynamical decoupling and spin echo as described in Chapter 24.1 and Chapter 5 can be used to remove systematic errors and inhomogeneities on the expense of longer gate times. On the other hand, decoherence-free subspaces (DFS, see Section 24.2) use symmetries of the noise mechanism to protect qubit states and their effectiveness influences gate fidelities that can be benchmarked.

6.3 Qualitative criteria beyond DiVincenzo

A lot of this has been discussed in the surface code chapter already, see surface code chapter.

6.3.1 Connectivity

Error correction codes need the right connectivity of physical qubits to carry out operations, e.g., a nearest-neighbor lattice for the 2D surface code. A 1D-architecture with nearest neighbor-connectivity needs to face extremely low thresholds—full connectivity such as, e.g., in Monroe's ion trap [LMR⁺17] Figure 1 allows to implement surface codes of high dimension with high thresholds.

6.3.2 Parallel operations

Error correction is envisaged to be done in parallel or with at most constant overhead on all qubits. Sequential error correction cycles would render error correction ineffective. An example of a non-parallelizable architecture is coupling all qubits to a single bus, which can typically only mediate a single two-qubit operation.

6.3.3 Supply of fresh qubits

Fresh initialized ancillae in error correction are needed in all cycles, requiring either a large supply or fast reset. Time for this needs to be factored into the determination of time constants for error correction.

6.4 Benchmarking operations

6.4.1 Gate fidelities

There exist a wealth of fidelity functions allowing to estimate the proximity of two quantum operations. In its simplest form, the fidelity can be written as a state overlap $\langle \psi_F | \psi_T \rangle$ between a desired state $|\psi_F\rangle$ and the final state $|\psi_T\rangle$. The final state is obtained by applying some operation on the prepared input state $|\psi_0\rangle$, and the desired state is the one we would get if the operation on the input state would be ideal. This state overlap basically defines the fidelity of the quantum process, but it depends on the given input state, leading to a large range of obtainable fidelities.

There are two natural routes to lift this input state-dependence: one is to average over all input states. Such an average can be reduced to the Hilbert-Schmidt scalar product of the two corresponding operators, like the trace fidelity $|\text{tr}(U_F^\dagger U_T)|^2 / \mathcal{N}^2$ for a unitary process with the desired evolution U_F and the implemented evolution U_T . This can be extended for general maps, to a trace fidelity $|\text{tr}(\Phi_F^{-1} \circ \Phi_T)| / \mathcal{N}^2$, where $\Phi(\rho_0) = \rho$ maps density matrices onto density matrices, and the indices stand for desired (F) and actual (T) gates. The average gate fidelity is yet another way to measure the fidelity of a process and is defined through

$$F_g = \int \langle \psi | U_F^\dagger \Phi_T(|\psi\rangle\langle\psi|) U_F |\psi\rangle d\psi .$$

The integration is done over all possible input states in the computational subspace. We will see in the end of this section that the average gate fidelity can be estimated efficiently through RB. However, for high dimensional gates, like an n -qubit controlled-phase gate $C\dots CZ$, i.e., a lot of control qubits to perform a Z gate, leads to a high fidelity even for the identity operation. Without loss of generality this gate is a diagonal matrix where all but the last entries are 1, and the last entry is -1 . Now the average gate fidelity for a final gate being unity (quantum memory) is $1 - 4(2^n - 1)2^n(2^n + 1) / (1 - 4(2^n - 1)2^{2n})$ for the trace norm, increasing with the number of qubits.

The other way is to look at the worst input state—the one producing the largest error. That combined with the possibility to augment the operation with a unit operation (hence finding the worst input state over a large set) defines the diamond norm. This measure depends on an input state that is defined through the norm itself. The diamond distance of two quantum channels Φ_1 and Φ_2 is defined as the trace distance of the channels for the worst case input state ρ

$$\|\Phi_1 - \Phi_2\|_D = \sup_{\rho} \|(\Phi_1 \otimes 1)(\rho) - (\Phi_2 \otimes 1)(\rho)\|_1 .$$

Then the diamond norm is 1 minus the diamond distance. The diamond norm measures the fidelity of the worst case possible, which maximizes the diamond distance. The diamond norm returns a significant error for this wrong implementation of the n -qubit controlled-phase gate $C\dots CZ$.

It is the diamond norm that enters the threshold theorem of fault tolerance. We will see that it is cumbersome and inefficient to measure, so one needs to rely on bounding it by feasible measurements. This statement will be made more formal later.

6.4.2 Process tomography—idea and pitfalls

The historic first benchmarking procedure proposed has been quantum process tomography (QPT), which is based on quantum state tomography (QST) [NC00]. QPT aims at reconstructing the full quantum process, from which operation errors and fidelities can be computed (for a caveat see next subsection). The goal of QST is to reconstruct the full density matrix of a state through measurement. For a system of qubits, it consists of measuring the expectation values of all combinations of Pauli matrices, including the identity, in a given state. Quantum state tomography is a procedure to measure the complete density matrix. For a system of qubits, it consists of measuring the expectation values of all combinations of Pauli matrices (including the identity) in a given state. It thus requires a number of measurement operators that is

exponential in the number of qubits. Practically, in most cases with the possible exception of photon polarization, physical detectors are set up to measure only one specific observable. Measuring any other Pauli operator requires additional operations between the operation of interest and measurement, introducing an additional error source. Practically, measurement imperfections can easily lead to non-physical density matrices (e.g., with negative eigenvalues), which can be mitigated by advanced data analysis, need to be determined through many repeated experiment for each generalized Pauli operator, then the Pauli decomposition of the state ρ can approximately be reconstructed. Although it is conceptually easy, it needs accurate state preparation and measurement (SPAM), and is prone to errors in these.

Quantum process tomography (QPT) consists of preparing a complete set of pure initial states spanning the space of input density matrices, which is of size d^2 . Applying the process $\epsilon(\rho)$ for each initial states and performing full QST on the output [NC00] determines the full quantum process. This increases the needed resources by a factor of d^2 , and is still prone to SPAM errors. Given that in quantum processors there is usually a single fiducial initial state, preparing the input state adds another operation leading to an artifact analogous to that in QST. These errors are called State Preparation And Measurement (SPAM) errors.

Compressed sensing allows to reduce the number of experiments for QST [GLF⁺10] for low rank density matrices with dimension d and rank r . Since the density matrix is sparse in this case, it is uniquely determined by a small number of random Pauli measurements and can be reconstructed efficiently via convex optimization.

6.4.3 Randomized Benchmarking and Interleaved Randomized Benchmarking

QST and QPT need a lot of resources to characterize even small quantum systems. Additionally, they require accurate state preparation and measurement (SPAM), and are vulnerable to errors in these. An efficient way to estimate quantum gates is RB [KLR⁺08]. It does not need that many measurements and is stable under SPAM errors. Therefore, it is good candidate for characterization of large systems and the de facto standard tool for such a task.

The basic RB protocol works as follows: first, one chooses a fixed sequence length m , where a sequence contains $m+1$ Clifford gates. The Clifford gates form the Clifford group and are the normalizers of the Pauli group: they map Pauli operators onto Pauli operators. The generators of the Clifford group are the phase gate, the Hadamard gate and the CNOT gate. The last gate in a sequence is set as the inverse of the concatenated preceding m gates, which is feasible given the group structure. For the chosen m , one builds K_m random sequences, each with an error Λ , and calculates the average of the K_m fidelities which are the measured survival probabilities of the initial state. This is repeated for each m , and fitted to a decay curve model of the form $A_1 p^m + B_1 + C_1(m-1)p^{m-2}$ [EAA⁺05, MGE12]. The coefficient A_1, B_1 and C_1 are fitting parameters and p is the depolarizing parameter. The factor A_1 is interpreted as the total SPAM error and $1-p$ as the error per gate. The average error rate is then given by $r = (d-1)(1-p)/d$ with the dimension d . RB approximates the average fidelity function. This remarkable result is based on the possibility to interchange the averaging over all input states to a channel by averaging over random sequences. The error channel to be characterized is twirled by the random Clifford gates and it has been shown that it reduces to a one-parameter depolarizing channel $\rho \rightarrow (1-p)\mathbb{1} + p\rho$ which can be extracted from the survival probability of any (convenient to prepare) initial pure state. The original proof makes the vastly unrealistic assumption of a gate-independent error that has later been lifted to weakly gate dependent errors [MGE11, MGE12]. Note that this technique allows to measure even small errors by making the sequence very long in order to bring the sequence error into a range that can be conveniently detected. Its convergence is rather fast, which has later been quantified [CW15].

As described, RB measures the average error of the whole Clifford group. Interleaved randomized benchmarking (IRB) [MGJ⁺12] allows characterization of a specific Clifford gate by detecting the lowering of the sequence fidelity because of twice as many applied gates in the random sequence with same nominal length m . In IRB after each random Clifford gate the target Clifford C_T is applied, again with errors Λ and Λ_F , respectively, and the last gate is again the inverse of the ideal concatenated sequence. IRB is preceded by a

parameter estimation through RB, and can be obtained from IRB by using the identity operator as specific gate. Comparison of the sequences allows to extract the error per gate for that specific Clifford gates.

RB has been implemented in many systems, and typically requires a modest number of measurements mostly controlled by the sheer size of the two-qubit Clifford group. It is also possible to characterize leakage errors with RB separately and several protocols have been proposed [ECMG14, CW15, WG17] to do it. But it should be pointed out that T gates cannot be benchmarked efficiently, which is a consequence of the Gottesmann-Knill theorem [Got98], that states that non-Clifford gates are computationally hard to simulate classically. There are some attempts to include non-Clifford gates, or at least trying to reduce complexity by forgoing the last inverting Clifford gate and performing optimized state tomography instead [CMB⁺16, CRKW17]. Another proposed idea is Randomized Benchmarking Tomography (RBT) [JdSR⁺15]. The protocol is compared to IRB, where non-Clifford gates for RB are written as linear decomposition of Cliffords [KdSR⁺14], and the latter are benchmarked with IRB. For characterization at the logical level the idea of logical RB [CGFF17] has been proposed recently.

It needs to be noted that the difficulty of characterizing non-Clifford gates with RB is not considered to be a major problem. The physical T -gate is not more difficult than the Z gate, which is a Clifford gate and one should not expect these errors to be vastly different. This is in sharp contrast to their difference in complexity as *logical* gates.

6.4.4 Gate set tomography

A complementary tomography tools to characterize qubits is gate set tomography (GST) [BKGN⁺13, MGS⁺13]. It is designed as a black-box characterization tool, such that the quantum device is accessible only through classical controls and measurement outcomes. In contrast to QST and QPT, it does not rely on accurate state preparation and measurement. Compared to RB, it needs much more resources: about 10^3 sequences for a single-qubit and 10^5 sequences for two-qubits. But it returns full tomography of gates, state preparation and measurement simultaneously, and an estimate of the diamond norm. GST has been successfully tested, for example in ion traps [BKGN⁺13, BKGN⁺17] and semiconductors [DMBK⁺16]. A Python implementation of GST (pyGSTi) can be found on GitHub [NER16].

In the black-box description the device contains some buttons to apply quantum gates, including an initialization button to prepare the (probably unknown) state ρ , a measurement button that returns a binary outcome, and K gates G_i . GST then works as follows: the state ρ is initialized, followed by a sequence of gates $s = \{G_{s_1}, \dots, G_{s_L}\}$ with length L , and a final positive-operator valued measurement (POVM) E . Each such experiment is repeated N times to gather sufficient statistics of the recorded outcome, and this is done for M different sequences. The number M scales with Kd^4 , where d is the Hilbert space dimension and K the number of gates one can apply directly (i.e., the number of gate buttons). Then linear inversion (LGST) provides rough estimates of the gates, state preparation and measurement (simultaneous state and process tomography)[Gre15], and is used as a starting point for maximum likelihood estimation (MLE). Each sequence consists of three parts: an initial fiducial sequence, a short germ sequence which is repeated several times, and a final fiducial sequence [BKGN⁺17]. The fiducial sequences effectively change the initial state and the measurement basis. Repeating the germs allows to enhance specific errors, such as over-rotation, tilt or dephasing. GST is therefore more sensitive to coherent errors compared to RB, which randomizes over gates. GST assumes that the gates are Markovian and non-Markovianity is obtained from deviations in the fitting model, where short sequences are less prone to non-Markovianity. Up to the choice of basis (gauge) the gate set $\{\rho, E, G_i\}$ is self-consistently determined. A consequence of the gauge invariance is that the gates do not have to be completely positive and trace preserving (CPTP) maps in an arbitrary basis. Therefore, GST does not enforce the CPTP condition and the gauge is usually chosen such that the estimated gates are as close as possible to the target gates.

6.4.5 Cross-entropy benchmarking

Cross-entropy benchmarking has been introduced by the Google group [AAM⁺19] as a means to benchmark large quantum processors as a whole rather than individual components or gates, and to be able to naturally include non-Clifford gates. Such an undertaking needs to make sure that the benchmarking operations are representative hard for the quantum hardware as they scale, and that classical simulation of a large quantum computer is not required or at least kept to a minimum.

These constraints are implemented by using a sampling problem as a synthetic benchmark, i.e., by running a quantum algorithm that does not have a unique result but rather a distribution from which measurement is sampling. If those distributions have both a quantum and a classical limit, then comparing the output distribution of the circuit to both these distributions through their cross-entropy, a well-known model testing method, allows to determine if the device is still a quantum processor.

In the case of Google, the algorithm implemented is a random set of gates. It is shown that the output distribution of that set of gates is described by the Porter-Thomas distribution, a notion from the field of quantum chaos (i.e., the quantum physics of classically chaotic systems). Its classical counterpart would be a uniform distribution of output values and it is argued, that simulating the system on a classical computer with a polynomial-time algorithm needs to take shortcuts equivalent to reaching only that distribution. It is in fact shown that simulating quantum chaos of this type is computationally equivalent to problems in NP.

In this sense, cross-entropy benchmarking is suitable to its mission in that it allows to benchmark large quantum processors by certifying that they are quantum. It certainly also gives insights into error rates even though these are currently under debate. It is not a (socially motivated) application of a quantum computer – in this sense XEB is a purely synthetic benchmark – nor is it a replacement for the benchmarking methods (RB, IRB, GST).

6.4.6 Risks at mid-level

The average gate fidelity measures how well a channel Φ performs a desired unitary gate U , averaged over all pure input states $|\psi\rangle$ in the computational subspace. But as we have already seen, there exist gates with bad implementation but high average fidelity. The average error per Clifford gate, defined as $r = 1 - F_g$ [MGE12], is not a measure for scalability [SWS16]. For example, the average gate fidelity underestimates unitary errors compared to the diamond norm. The proper threshold is defined through the diamond norm, but the latter lacks of an efficient estimation like RB gives for the average error per gate. The question arises how useful the average error is. In [SWS16], an upper bound η^{up} is given derived from r and the Hilbert space dimension d , such that the following condition holds if errors can be efficiently corrected

$$\eta \leq \eta^{\text{up}} \leq \eta^{\text{lb}} \leq \eta_0 \quad ,$$

with the threshold η_0 for FTQC, the error rate of the device η , and η_0^{lb} the lower bound error rate.

Furthermore, it's not completely clear if RB really estimates the average gate fidelity, since there are some problems with the gauge invariance [PRY⁺17]. It is currently debated whether cases where diamond norm and average fidelity vastly differ are practically relevant or pathological. While no example to the former has been found, it turns out that fidelities measured by GST are typically slightly lower than those obtained by RB.

Another point of concern is the Pauli twirling approximation (PTA) [GZ13] of arbitrary channels Λ . Given a map which reads $\Lambda(\rho) = \sum_i E_i \rho E_i^\dagger$, Pauli twirling (also full Clifford) takes the input state ρ , rotates it by a Pauli operator σ_i , applies the map Λ , and finally counter rotates the final state. The approximation is then performed by taking the average over all Pauli (Clifford) matrices, which can be cast into the form $\Lambda'(\rho) = \sum_{\sigma_i} p_{\sigma_i} \sigma_i \rho \sigma_i$. The Pauli-twirl of any channel is mapped onto a Pauli channel, Pauli channels are mapped onto themselves, and the channel Λ and its Pauli twirl Λ' have the same average gate fidelity. Since Clifford twirling of a quantum operation leads to a depolarizing channel [MGE12], RB can effectively use a fit model for the depolarization parameter p to estimate the average error per gate. However, for coherent errors the

Pauli twirling is in general not sufficient [GSVB13] and one has to take into account the difference between a coherent error and its Pauli twirl for threshold calculations (Pauli distance) [SWS16].

A currently under investigated error in RB is the role of errors that error correction does not catch [WF14a], such as extreme non-Markovian or highly correlated errors. These will all affect the RB result but are difficult to treat with error correction, hence the impact of error correction may be lower than the estimates of the next chapter predict if the physical error rate is estimated with RB.

6.4.7 Recommendation

This intermediate state is where most of the work of building scalable quantum processors is currently performed. It allows to predict what a fault tolerance implementation would do. It should be applied as follows:

1. Verify if one- and two-qubit RB experiments been done. Do they find error rates that are below the fault tolerance threshold for the envisaged error correction code?
2. If the first step has been met, verify whether the error rates have been at least in a sample verified by another method—modern process tomography or gate set tomography and whether the results from these methods still allow for fault tolerance. Monitor the impact of a subsequent error correction experiment (for first implementations, see [WL17, SBM⁺11, KBF⁺15, ARL⁺17] also described in section 22.4). If it is not as effective as expected, this points to temporally or spatially correlated errors.

6.5 Phenomenology of quantum errors

As a foundation for our error correction in chapter 7 we describe a classification of error mechanisms that relate to the simplified (but, as it turns out, sufficient) error models studied there.

It is helpful to develop a phenomenology of errors in a quantum algorithm. This is needed specifically to better relate to error propagation through a quantum algorithm and the state of knowledge of quantum error correction performance. This phenomenology needs to build on quantities accessible by benchmarking methods such as the average gate fidelity. It also needs to be compatible with the intrinsically stochastic nature of quantum physics. Furthermore, it needs to focus on predicting the probability of output error of a complete algorithm from the fidelity of elementary operations gates. For this, it is crucial to remember that cryptanalytically relevant quantum algorithms are aiming at producing a single output or a sparse distribution, i.e., the complete success quantifier is the probability to measure the correct binary output. It is thus, for example, not necessary for the real quantum computer to be in the precisely correct final state in order to produce the correct result – as long as the state is not in its orthogonal complement, there is still a nonzero probability for the correct output.

Specifically the stochastic nature of quantum physics must not be forgotten in this phenomenology. Stochastic processes such as random errors can be described either in an event-by event trajectory (akin to the Langevin equation) or by computing distributions at every step (akin to a Fokker-Planck equation). These make equivalent mathematical predictions. For stochastic errors, phenomenology and verbalization are often based on a Langevin description. The standard quantum formalism of quantum states in a Hilbert space is closer to a Fokker-Planck description – it is similar in the sense that it makes probabilistic predictions that can only be completely tested in an ensemble yet it is different in the sense that it describes a probability amplitude – a complex quantity whose absolute square amounts to probability – rather than the probability itself. The trajectory-style description of quantum evolution as a sequence of transition would be done with Feynman path integrals, which are uncommon in quantum information [PKS17].

In order to evaluate algorithms, it is sufficient to analyze errors after every quantum operation, i.e., at discrete times. Some error mechanism, including drift from spurious fields, T_1 and T_2 -errors, are usually described by continuous-time differential equations (such as the Schrödinger equation, the Liouville-von-

Neumann equation, the Bloch equation or more general master equations), in which case a time-discrete description amounts to the formal solution of said underlying differential equation between the beginning and the end of the discrete time unit. Examples are discussed in section 22.3.

Furthermore, it is natural to first decompose the real implemented gate operations into a sequence of the ideal gate operation and an error, see below. This is always possible as gates are unitary, hence they are invertible. This needs to be a permissible quantum map of density matrices onto density matrices, i.e., it needs to preserve hermiticity and trace and map positive matrices onto positive matrices. The distance of that map to the identity defines the gate error. It can be measured with the benchmarking operations described in section 6.4.

By assuming that this map applies to the density matrix in each step of the algorithm and it is in this sense assuming that this error map on a density matrix is systematic, i.e., occur on all steps, but it can have different size. However, the density matrix makes statements about statistical ensembles – on the one hand, a statistical mixture of quantum states. Furthermore, even pure quantum states only make probabilistic predictions about measurement outcomes and require an ensemble measurement to extract the full quantum state (quantum state tomography). It is still helpful to understand how to unravel these ensembles into an event trajectory in order to create a tangible error phenomenology [WM11]. It needs to be pointed out that this is a heuristic for interpretation, quantitative mathematical statements are only based on the size of the error, reflecting that fact that quantum physical predictions are stochastic in nature. In other words: An error that is systematic in a density matrix description need not happen in every realization of the experiment, and quantitative analysis should focus on the output of the algorithm only.

To this end, we resort to the Kraus operator decomposition for the error map which, describes the most general map of density matrices onto density matrices [NC00], i.e., a trace-, hermiticity, and positivity-preserving linear map $\hat{\rho} \rightarrow \sum_{i=1} \hat{K}_i \hat{\rho} \hat{K}_i^\dagger$ with $\sum \hat{K}_i^\dagger \hat{K}_i = \hat{1}$.

The different terms of this operator sum decomposition can be understood as different stochastic error channels whose probability needs to sum to unity. The fidelity measurements described in section 6.4 measure the average distance of this channel to the unit channel, which is much smaller than unity for a good quantum computing candidate. We first describe two extreme cases, the limits of purely stochastic and purely coherent errors.

If it is possible to have one of the Kraus operators to be proportional to unit matrix, WLOG the first one, then we can interpret its square as $1-p_0$, the probability that no error occurs, and all other Kraus operators describe errors occurring with some classical probability that can be extracted from normalizing the individual contributions. In the nontrivial case (i.e., with at least two distinct Kraus operators), it changes the entropy of the state and is thus connected to system-environment entanglement as highlighted in section 22.7. In that case the Pauli operators form a basis of all possible errors. An example of error correction of such errors is spelled out in section 22.1.

In the opposite case, we only have a single Kraus operator, which then due to the normalization of the Kraus decomposition has to be unitary (but not necessarily the unit matrix). We call this the limit of purely coherent errors. Here, we can in principle define the matrix elements of the error map as transition amplitudes into erroneous basis states as well as the accumulation of erroneous phases, accounting to a path-integral-type description of a quantum algorithm by formally inserting a projection on a complete orthonormal basis in state space in every step. A typical coherent error is a systematic unitary rotations (e.g. due to miscalibrated pulses at higher/lower frequency than the transition leading to over-/underrotation). Note that there is sometimes a subtle distinction in the open quantum systems literature between coherent and unitary errors: Whereas a unitary error is described by a unitary matrix on the qubit systems, the more general concept of a coherent error can include a coherent interaction with an unobserved degree of freedom outside the qubits. When one is restricted to observing only the qubit degree of freedom (as is the

case for error correction), this would manifest itself as an incoherent error with potentially time-dependent Kraus operators and thus we are not making this distinction in this study.

The correction of single unitary errors is remarkably similar to that of single stochastic errors with probabilities computed by Born's rule (highlighted in section 22.5), see section 22.2. An important distinction of these two scenarios is in the error propagation through a quantum circuit, i.e., how errors occurring in an earlier step of the algorithm affect later stages in the algorithm and how they get combined with later errors. In the case of purely stochastic errors, a conservative lower bound on the probability of error-free computation is given by multiplying the probabilities of error-free execution of all individual gates. In practice, errors of two-qubit gates are larger than those of single-qubit gates, but they are constant, so we can focus on these and take the probability of their error-free execution to the power of the number of these gates in the algorithms. For coherent errors, error amplitudes propagate through the circuit that get squared to yield the probability so, in the worst case, these will all interfere constructively. Whether this constructive interference is realistic (given that interference can also be destructive) or whether a compiled quantum circuit will rather randomize the phases of the errors, i.e., remove the coherence of errors, is subject of debate in literature which we will summarize when we discuss the impact of error correction on coherent errors in section 7.4.1.3. Ultimately, this research needs to be connected to experimental verification.

In a realistic quantum computer, of course, any viable Kraus decomposition of the error channel will contain more than one term and no Kraus operator will be precisely proportional to the unit matrix. We can thus interpret realistic errors as a combination of both scenarios – a statistical mixture of errors with a coherent component – and realistic error propagation to be within the margins of both extremes.

Note that these scenarios allow for errors to be time-dependent (the Kraus operators need not be constant) as well as non-Markovian (the Kraus-operators can be correlated over different time-steps). Furthermore, errors can be classified as single- or multi-qubit errors along the structure of Kraus operators, specifically, if they can be factorized or not.

Specifically, the Kraus operators provide information about error correlations that is hidden in their weight, i.e., the number of qubits in which they are different from unity. Errors are not recoverable if those correlations are larger than the code distance and correlated errors lower the code threshold. We outline in section 22.6 however that due to the stochastic nature of quantum physics, many classically correlated error models do not necessarily lead to correlated errors and that high-weight errors usually require high-weight interactions.

As a final remark it needs to be re-emphasized that the main accessible quantity in experiments is the fidelity, i.e., the overlap of the error with identity, see section 6.4. Rather than attempting to directly measure Kraus operators, the response of quantum processors to protocols with known sensitivity to different types of errors is a promising tool to clarify this phenomenology.

In general, it is crucial input for future extrapolation of fault tolerance to analyze the error budgets of experiments, specifically the contributions of stochastic vs. coherent errors.

6.6 Application to the Google quantum supremacy experiment

The Google group [AAM⁺19] has published a highly celebrated experiment claiming quantum supremacy, i.e., the execution of a computational problem that would be impractically long on a classical computer given its restrictions. The problem studied as laid out in [BIS⁺16] is the creation of a Porter-Thomas output distribution from a random quantum circuit, which is characteristic of quantum chaos, and then sampling from it. The classical simulation of this problem is claimed to take around 10000 years on the currently

largest supercomputers in the world largely because of a memory wall. It has been argued [PGMG19] that by more compact memory use that time could be brought down significantly. Yet, it has not been disputed that enlarging the processor would double the classical memory per added qubit. The algorithm was executed without error correction and the quantumness of the output was verified with cross-Entropy benchmarking (XEB), i.e., it was verified that the error rate was low enough so this is a bona fide quantum computation.

The hardware platform are superconducting transmon qubits arranged in a 9 x 6 array. Of the 54 qubits, 53 are functional and are used in experiment. This grid layout requires a three-dimensionally integrated chip with control and readout in the third dimension. Gate fidelities are not at but close to the previous records in superconducting qubit, held by the same group.

As laid out above in section 2.12 based on chapter 4, the gap in error rate between these experiments and those required by cryptanalysis is significant, thus, cryptanalysis should not be seen as a direct application of this technology, rather, the experiment informs the extrapolation of fault tolerance overhead. Specifically, it defines a best-in-class benchmark on processor size, gate fidelity and speed. Its architecture is a square two-dimensional array as required by fault tolerance using the surface code. Specifically, the experiment demonstrates that the system integration required at that scale can be mastered and are not an operational limitation. The technological requirements include a large number of control wires in the cryostat, a three-dimensional multi-chip module with control and readout moved away from the quantum plane, and adequate low-level software

One may still ponder whether executing a purely synthetic benchmark such as XEB is a valid milestone for quantum supremacy. On the one hand, unlike some more purpose-built platforms in the world of neutral atoms (as described in section 16.1.2.1) this is a freely programmable, universal quantum processor, so it could execute other tasks that follow a purpose other than pure benchmarking. On the other hand, it is a valid and intensely researched question what it takes to reach quantum advantage, i.e., the solution of a non-benchmarking but application-driven problem on a quantum processor. It is currently expected that hybrid quantum-classical algorithms with applications in theoretical chemistry will be the first to reach that level. There have been validations of that concept which are however not even close to quantum advantage yet. Precise roadmaps are intensely researched and are beyond the scope of the present study. In this sense, the Google experiment outperforms classical supercomputers, but does not represent quantum supremacy in an “applied”, i.e., externally motivated, situation.

7 Connecting physical and logical qubits with fault tolerant quantum computing

This chapter gives an overview of the criteria that need to be fulfilled in order to perform quantum error correction to a set of qubits, and on the improvement of logical error rates that can be reached this way. We will focus on the surface code, which has turned out to be the most promising candidate. In the original and hardly implementable idea by Kitaev [Kit97a] the qubits are arranged on the surface of a torus, however, it turned out [Kit97a] that the two dimensional surface code works efficient also with nearest-neighbor interactions on a square lattice with (non-periodic) boundaries, so there is no need for further consideration of the toric code¹. There are several other promising quantum error correcting schemes, like the color code or topological cluster states which will be discussed in short at the end of this chapter. These codes provide additional correcting possibilities, but often come with higher requirements to the physical implementation.

In order to assist with processing the rather specialized language of quantum error correction, we have added a glossary of terms in the end of this chapter, section 7.7.

7.1 General observations on the role of fault tolerance

Albeit accepting and producing regular binary output, quantum computers store intermediate information in the probability amplitudes of complex quantum states, which are analogue quantities. While these analogue quantities are not read out in a running quantum computer which, as in Shor's algorithm, only has binary input and output, they do matter as they provide probabilities for actual algorithmic errors. It is thus imperative to correct potential output errors which, in practice, means reducing the probability of a wrong output to a fixed, acceptable value. Also, the number of distinct probability amplitudes is exponential in the number of quantum bits. This alone opens up a multitude of possibilities for errors in computation. One way to battle these errors would be to completely isolate the quantum bits—however, in practice this is not possible as one needs access to them to control, program, and read-out the quantum computer, hence necessarily open up the quantum processor to error channels. As described earlier, achievable error rates are vastly overcome by the demands of cryptographically relevant algorithms. There are applications for quantum computers in this regime discussed under the heading of quantum supremacy, cryptanalysis is not one of them (see Chapter 4). For all other algorithms, fault tolerant quantum computing centered around error correction has been proposed. Error correction allows to use error-prone hardware to efficiently simulate a perfect quantum computer with a predefined precision, i.e., its goal is to reduce the error of the final result to a predefined value. In this sense, error correction is not complete and fault tolerantly implemented algorithms are to this degree stochastic.

Before describing the current gold standard of quantum error correction and its applications as an evaluation tool for quantum computers, we start with a few basic clarifying observations and terminology.

7.1.1 Redundancy and measurement

Similar to classical error correction, quantum error correction is based on redundant information encoding: A single *logical* qubit (qubits in which the fault-tolerant algorithms on higher levels are expressed) is encoded in several *physical* qubits. Information about errors is extracted by *syndrome measurements*. Given the intrinsic invasiveness of quantum measurements, this cannot be done by reading out all qubits first and then classically comparing the outcomes, rather, the syndrome information needs to be mapped by a small

¹ Note that, however, codes requiring long-range (or 3D) interactions can have much higher fault tolerances, so if the problems concerning long-range interactions between qubits can be solved one day, these codes might gain relevance again.

piece of algorithm onto an ancillary degree of freedom, which is then read out, revealing only syndrome information but not the logical state. The projective nature of the quantum measurement turns analog error amplitudes into probabilities for digital errors and given that the ancilla is entangled with the data qubits also guarantees that their state is projected into a state described by these syndromes. In other words, it is syndrome detection that translates analog errors into digital errors with a probability given by the analog amplitudes, allowing to handle the errors digitally – the projective nature of the measurement performs a test and replaces probability by certainty. It turns out that these resulting digital errors can be completely described by bit-flips and phase flips and arbitrary combinations thereof (Pauli errors). This also applies to over- and underrotation errors (see Section 5.2) where the error probabilities can be calculated from the misrotations according to Born’s rule, with a number of caveats described in section 7.4.1. This is often stated but not exemplified in literature, hence we provide a simple example and contrast in Appendix 22.

This is referred to as a *parity measurement* and the operators being measured are products of an even number (the *weight*) of single-qubit operators which have (degenerate) eigenvalues ± 1 . In fact, in effectively characterizing quantum error correction one tends to describe quantum states not by their state vector, but by their set of *stabilizers*: A set of commuting operators that uniquely define the state up to a global phase by requiring it to be a $+1$ eigenstate of all of them.

7.1.2 Error detection, matching, and correction

An error correcting code is characterized by the number and type of (physical) errors it can handle, with larger codes being in general more powerful. After measurement of the (physical) syndrome operators, it is important to link the syndrome pattern to the errors that have happened. In simple majority-voting this is straightforward, but in complex codes this is difficult and requires complex minimum-weight matching algorithms [Edm65](see Section 7.2.2.4). Once errors are identified, conventional wisdom suggests that corrective action needs to be applied—i.e., gates that correct the error. Experimentally, this requires feed-forward, i.e., conditional application of gates within the time scale of the algorithm. This is enormously challenging—but also understood to be unnecessary. As all the errors, after syndrome detection, are all Pauli gates² (hence form a subgroup of the Clifford group), their effect can be simulated classically with polynomial computational effort, so it is sufficient [Got98] to implement corrections on the final results of the quantum computation. This can be done completely after the quantum computation, as a correction to the classical output bit-string.

What this procedure still requires are initialized ancilla states after each round of error detection, placing a stringent requirement on measurement.

7.1.3 Concatenated codes and the threshold theorem

There are various ways to enlarge the error correction code and correct more errors. A standard and illustrative way are *concatenated codes*: Build first-level logical qubits out of physical qubits and repeat this construction building n th level logical qubits out of $n-1$ st level logical qubits and so on.

Now clearly, adding error correction even to quantum memory alone introduces both more qubits (with more entry points for physical error) as well as extra operations even in the idle state, which naturally gives rise to the question whether this actually is beneficial, or whether the measurement and the extra gates that are also faulty negate the improvement from error correction. This is answered by the *threshold theorem* [ABO99] which exists in various versions. It states that there is a threshold rate for the physical error of a real quantum computer, below which an ideal quantum computer can be simulated arbitrarily close using error correction. Practically this means that below threshold, the logical error rate (measured by the diamond norm, see Section 6.4.1) can be arbitrarily reduced (hence the size of an executable algorithm can be

² In principle, any error can happen to the qubits, but the syndrome measurements are always chosen such that errors are mapped to Pauli (or identity) operators. This can be done by measuring only in Pauli-eigenbases.

arbitrarily prolonged at fixed final logical error rate) by adding more and more error correction (e.g., in the form of more concatenation).

In order to compute thresholds, one needs to analyze each step of an error correction procedure on a given architecture. One designates analytically proven and numerically simulated thresholds, the latter being more generous. Architectural features that need to be discussed include the connectivity (expressed by range of interactions and the dimensionality of the processor), which determine swapping overheads and other features. Proofs of the threshold theorem assume errors that are uncorrelated between qubits. This is a crucial ingredient - logical qubits will be exposed to error when all of its components experience errors at the same time, thus it is only effected if these simultaneous errors are less likely than individual errors. The precise definition of multi qubit errors and the likelihood of them to occur are studied in Appendix 22. The numerical value of the threshold contains the trade-off between the extra operations required for fault tolerance—which are all in themselves assumed to be imperfect—and the protection offered by error correction. Existence of a threshold thus requires to code to be efficient enough for the extra operations to not eat up the protection.

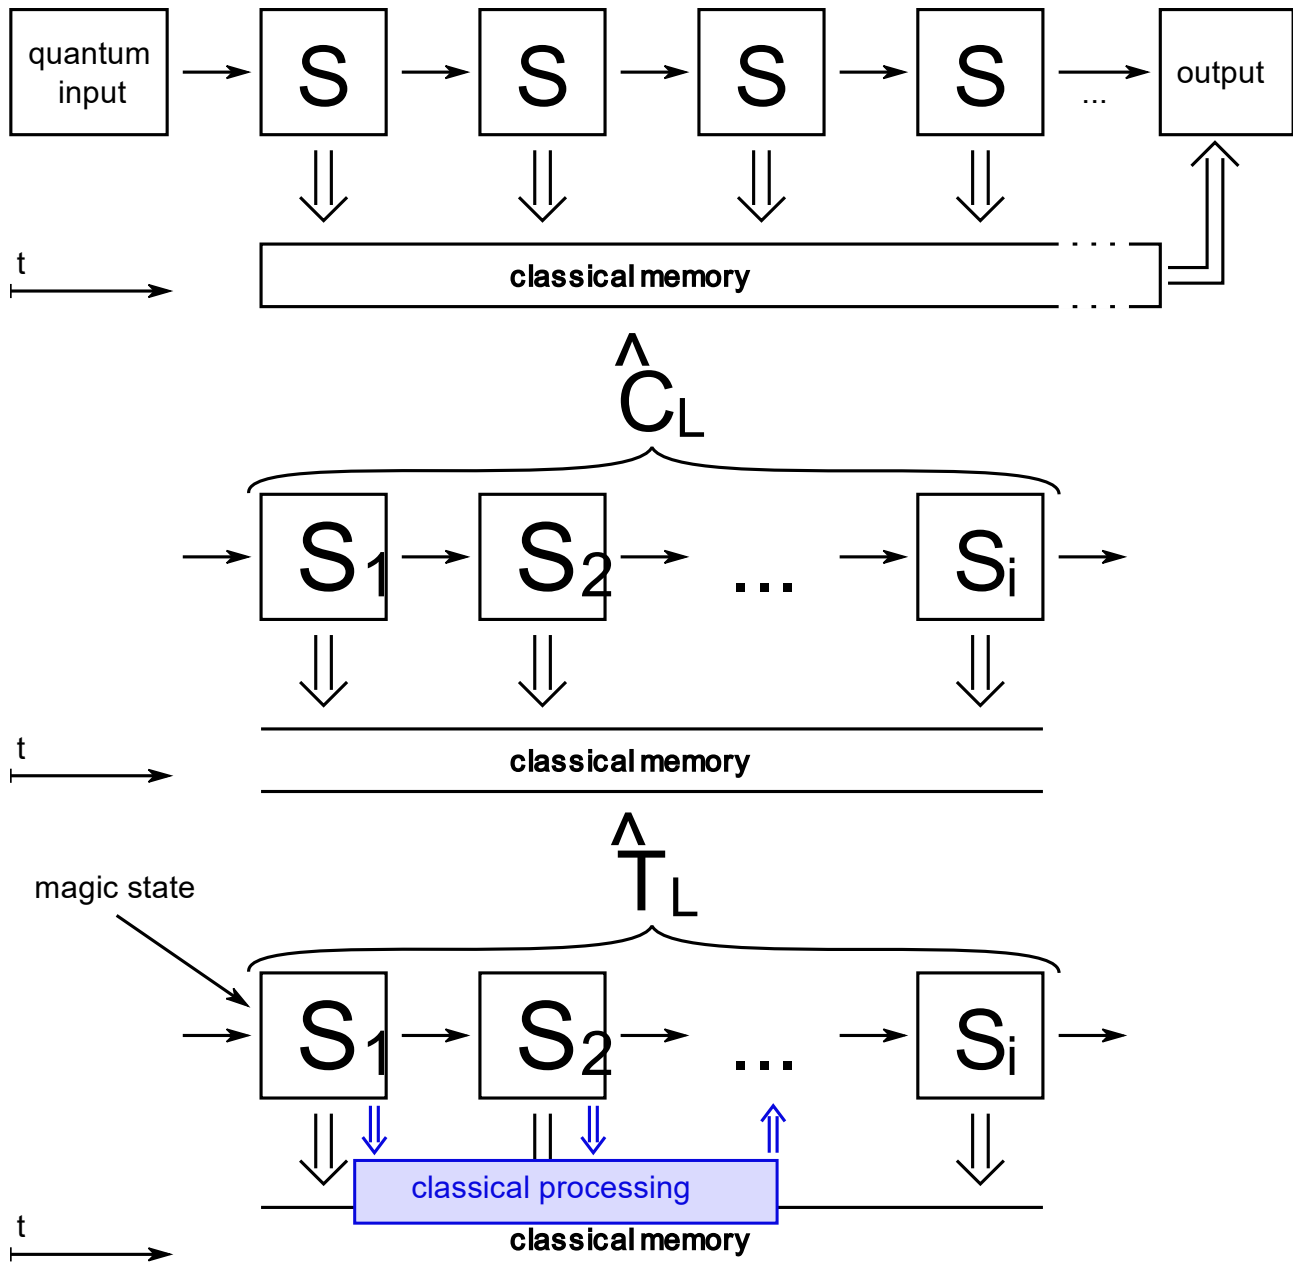


Figure 7.1: Schematic diagram of fault-tolerant quantum computing/storage. Boxes with S represent one cycle of syndrome measurement, in which all stabilizers relevant for the detection of errors are measured. The syndrome measurement consists of several physical operations on the bare physical qubits, which are necessary to get an effective multi-qubit measurement operator. Top: Protection of logical quantum information can be ensured by successive syndrome measurement rounds. The measurement outcomes of each round (double-line arrows) are collected in a classical memory and evaluated at the end of the process to identify and correct the effect of all errors at once. Middle: (Most) Clifford gates can be performed on the logical qubit states by modifying the syndrome measurements. This can be done by adding additional gates on the underlying physical qubits before the usual syndrome measurement circuit (as is the case for simple encodings like the Steane code) or by changing the set of operators that is measured, i.e., including more stabilizer measurements or not measuring some of them. The surface code uses both modification types, depending on the logical gate that shall be performed, see Section 7.2.4. One logical gate can require several rounds of (modified) syndrome measurement rounds. Bottom: Non-Clifford gates like the T gate can not be performed directly with this method. The most common way is to use a (logical) magic state qubit which is produced in a different part of the circuit by state distillation (see Section 7.2.4.3) and combine it using Clifford gates. However, this is a non-deterministic process that needs classical feed-forward (blue box) of intermediate measurement outcomes to correct for the probabilistic effects.

7.1.4 Fault tolerant computation

Error correction narrowly defined talks about stabilizing quantum memory. In order to perform an effective fault-tolerant quantum computation, one needs to be able to implement a complete set of gates between *logical* qubits. The simplest way to do that—decode the logical qubits, operate on the vulnerable physical qubits—recode into logical qubits, would undo most of the benefit of error correction, eliminate the threshold, and is hence not viable. Rather, one would like to perform logic operations on the encoded qubits, so called transversal gates. These gates are usually performed by modifying the syndrome measurement cycles, as shown in Figure 7.1 (middle). This does not necessarily need the application of any additional gates in between the measurement cycle and therefore does not increase the time between two cycles. If additional gates are needed, the number of gates between two syndrome measurement rounds is kept very low in order to not increase the possibility for errors before the next round too much. For an optimal two-dimensional quantum error correction code, the set of transversal gates is the full Clifford group. Now given the Gottesman-Knill theorem, Clifford-only quantum algorithms can be efficiently simulated classically, so this is not sufficient. The most popular non-Clifford gate is the T gate or $\pi/8$ gate, a phase shift of the $|1\rangle$ state by $\pi/4$ relative to the $|0\rangle$ state. The Solovay-Kitaev theorem [Sol00, Kit97b] along with its practical constructions [HRC02] guarantees that with Clifford+T one can efficiently approximate any unitary gate, including small rotations needed in phase estimations. Implementing this efficiently is a major practical challenge in fault tolerance [BK13], as it is not always possible to perform it on logical qubits without turning off the protection. Note that these statements all apply to *logical* gates and would not be true for *physical* gates. Usually, additional resource-demanding codes like magic state distillation (see Section 7.2.4.3) and classical feed-forward are necessary on top of the surface code to perform logical T gates with the same accuracy as Clifford gates (see Figure 7.1, bottom).

7.1.5 Conclusions for the evaluation system

The requirement of fault tolerance makes this a crucial connecting point between algorithmic research and hardware implementations in quantum computing:

1. Given that all cryptographically relevant algorithms are too long to be executed in a non fault-tolerant way, gate counts in quantum algorithms need to be understood as fault-tolerant gates. They should be given in Clifford+T counts. Consequently, gate times and number of qubits required correspond to logical gate times and number of logical qubits.
2. Hardware needs to meet the requirements of fault tolerance. It requires a precise method to evaluate the error rate in order to establish a relation to the threshold. Also, in order for a threshold calculation to be viable, the architecture needs to meet all the requirements of the threshold calculation, specifically the required connectivity of qubits and fast ancilla initialization. As one needs to accept that most of the quantum computer's effort goes into error correction, this is thus a driver for architecture.
3. Careful investigation of the efficiency of error correction below threshold provides an efficient and effective scheme to translate logical qubit requirements into physical qubit requirements. We give concrete formulae in Section 7.5.

The fault tolerance landscape is currently dominated by the surface code, which is the first code that makes full machine-level extrapolations possible. We will thus detail fault tolerance along this example and make it quantitative. We will remark on other codes in the end. Section 7.2 is supposed to give a pedagogical and rather extensive introduction to the surface code and its underlying mechanisms. The parts directly relevant for the assessment of a physical platform are the Basic requirements stated in Section 7.3 and the operational conclusions in Section 7.5. Section 7.4 discusses the performance of the surface code in more detail concerning the underlying error model, trade-offs and different experimental conditions.

Quantum error correction is a large and forward-looking theoretical and mathematical activity that necessarily makes assumptions about the physical world. Also, implementing the far-reaching ideas of this chapter is beyond the current status of experimentation. Still, small instances of error correction have been

experimentally verified. These experimental implementations [WL17, SBM⁺11, KBF⁺15, ARL⁺17] confirm some basic assumptions yet call for refinement of error correction models. Two very specific experiments are described in Appendix 22.4.

7.2 Introduction to surface code quantum error correction

The surface code is the currently leading error correction code for its high threshold and reliance on nearest-neighbor measurements only. Most of the principles described here are analogous to other error correction codes. The basic idea described above that errors are corrected unless they are massively correlated is implemented here in a topological sense—only error patterns that change the topological genus of the state can go undetected, and this is unlikely to happen. We first describe how to correct and stabilize the substrate state in Sections 7.2.1–7.2.2. This strong protection requires extra effort to compute within the states, which is described in Sections 7.2.3–7.2.4 and then continues to implementation of the T gate in Section 7.2.4.3.

7.2.1 Basic setup and measurement cycle: Arrangement of physical qubits

The surface code is a stabilizer code which means that the computational subspace is stabilized by a set of operators (i.e., $+1$ eigenstate of them), and errors up to a certain order lead to changes of the measurement outcome of some of those stabilizers. The stabilizers have to all mutually commute so they can be simultaneously measured, and they need to be complete so their eigenvalues specify the state.




Figure 7.2: (a) Arrangement of physical qubits for the surface code. Data qubits are shown as open circles, measurement qubits as solid circles. The green and yellow crosses denote Z and X stabilizer measurements of the data qubits at the ends of the cross, respectively. At the boundaries, the stabilizer measurements include only three data qubits, represented by truncated crosses. (b) Circuit diagram for the Z stabilizer measurement. Identities are included to compensate for the Hadamards in the (c) X stabilizer measurement. Each step is performed simultaneously for all stabilizers. One round of such circuits for all Z and X stabilizers along the array corresponds to one syndrome measurement Box as shown in Figure 7.1. Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society."

The qubits are arranged in two groups—data qubits and measurement (or syndrome) qubits—as shown in Figure 7.2. The measurement qubits are only used to indirectly measure the operator product $Z_a Z_b Z_c Z_d$ (“measure-Z”) or $X_a X_b X_c X_d$ (“measure-X”) of the four surrounding data qubits (at the boundaries, the operator products only include three qubits). Each data qubit is surrounded by two measure-Z and two measure-X qubits, the boundaries are chosen such that two opposite sides end with measure-Z qubits and the other two with measure-X qubits. The actual measurement of the operator product is performed by initializing the measurement qubit in an eigenstate of either X or Z, successive entangling all surrounding qubits with the measurement qubit using CNOT operations and then measuring it in the corresponding basis, X or Z. One such iteration, including all processes to measure the larger operator product indirectly, is referred as syndrome/stabilizer measurement or surface code cycle. For efficient performance, all steps of the stabilizer measurement (initialization, gates and measurement) need to be done *in parallel* along the whole array. For the CNOT, this means all qubits are arranged in pairs (in each of the four CNOT steps, a measurement qubit is paired with another of its four adjacent data qubits) and a CNOT has to be performed to every pair of qubits simultaneously.

The computational subspace is the set of states that are stabilized by all of those operator products, i.e., their simultaneous +1 eigenstate. For practical reasons, however, one will use the state on which the system is projected after an initial measurement of all stabilizers. This state is random, but sufficiently characterized by the measurement outcomes. This way, initialization of every single data qubit is not necessary.

This stabilizer measurement will be executed consecutively (logical operations are done mostly by adjusting the structure of stabilizer measurements rather than performing additional gates between the stabilizer measurement); whenever an error occurs this can be detected from changes in the measurement outcomes as shown in Section 7.2.2. These error detections are still a crucial part of the error correction as measurement projects the state into a stabilizer eigenstate hence digitizing all continuous errors. The syndromes can be understood as changes in the stabilizer eigenvalues. Detected errors are not directly corrected on the quantum state, they can be tracked through the classical control and corrected all together in the end of the computation process [FMMC12].

7.2.2 Error syndromes

7.2.2.1 Single errors

Given that the stabilizer measurements detect X and Z errors independent of each other, the surface code can also detect Y errors as a combination of an X and a Z error. Single (physical) qubit X or Z errors always lead to a change of the two adjacent stabilizer measurements: An X error on a qubit a will lead to a sign change of the measurement outcome of the operator product $Z_i Z_j Z_k Z_l$ for $a \in \{i, j, k, l\}$, and no change for X stabilizer measurements, analogous an Z error will lead to a sign change only in the corresponding X stabilizer measurements. Thus, each pair of neighboring X or Z stabilizer measurement sign changes can be identified with an Z or X error, and sign changes on all four stabilizer outcomes around one data qubit therefore corresponds to an Y error. There is no need for an extra consideration of errors that are not Pauli operators, since the stabilizer measurements will map all qubit states to either the original state or to a state with a Pauli operator applied to it. Take for example an error of the form $|\psi\rangle \rightarrow (\alpha \cdot 1_i + \beta \cdot X_i)|\psi\rangle$, acting on one qubit of the array. The stabilizer measurement will map the state $(\alpha \cdot 1_i + \beta \cdot X_i)|\psi\rangle$ to either $|\psi\rangle$ (i.e., directly projecting to the error-free state) or to $X_i|\psi\rangle$ (i.e., creating but also detecting an X error at qubit i), depending on the amplitudes α and β . In general, any code that can correct a set of error operators can also correct any linear combination of them [Bac13]. A detailed and more general discussion of this is found in Chapter 22 or reference [Bac13, Section 2.6].

A special case happens at boundaries, when a data qubit is only surrounded by three measurement qubits and therefore an error can lead to a sign change of only one stabilizer measurement. This is the case for an X error next to a missing Z stabilizer or a Z error next to a missing X stabilizer.

7.2.2.2 Error chains

Whenever two identical errors happen to neighboring data qubits, the effect on the measurement qubit in between will cancel out, so one will only detect sign changes at the two outer measurement qubits. Similar, for longer error chains, only the measurement qubits at the ends of the chain will show an effect. This is not a problem as long as the chains are not too long and there are not too many errors along the array. Then, one can deduce the error chain leading to a several measurement outcome by error path with the highest possibility to occur. However, if errors get too dense, the stabilizer results might be misinterpreted and the syndrome extraction algorithm will correct a wrong path. This is not a problem as long as both paths can be topologically deformed into each other (i.e., without crossing any holes - needed for logical qubits, or changing connections to boundaries) since the changes resulting from original error + correction will then only be a closed (empty) loop of sign flips which does not change the logical state.

A worse case are error chains that start and/or end at boundaries (as will be shown, deactivating stabilizer measurements and introducing additional boundaries is a fundamental part of the computation process in

the surface code). If the sign changes at the ends of a chain fall on two deactivated or not existing measurement qubits of the same kind (X or Z), the error chain remains undetected: Along the chain, all sign changes cancel out and at the ends, sign changes are not measured. This is already a problem if only one end lies at a boundary and the other is close to a boundary. In general, whenever more than half of the qubits along a chain connecting two boundaries have an error (all of the same kind that is not measured at the corresponding boundaries), then this will result in a logical error, since the syndrome extraction algorithm will misinterpret the sign changes and correct the wrong qubits (see Section 7.2.3). A more detailed analysis of the impact error chains can be found in [FSG09]. The physical mechanisms leading to error chains are also discussed in the Appendix.

7.2.2.3 Measurement errors

A faulty sign change of a stabilizer measurement can also be caused by an error of the measurement qubit. Since for every surface code cycle the measurement qubits are reinitialized, such an error will probably vanish in the next round (see also [FMMC12], Section V) unless there is a massive correlation. Similar to spatial error chains it is, however, possible that such an error occurs multiple times in a row (with very low error probability). This might look like a real sign change instead of an error on the measurement qubit. Interpreting this the right way is part of the classical software layer of error correction.

In order to distinguish between measurement and data errors better, it is necessary to compare several rounds of syndrome extraction and see if a sign change stays or withdraws in the next round: If the changed sign remains a single event (or very rare), it is probably due to an error on the measurement qubit and can be ignored. So in general, more fault-tolerant implementations will not only need more qubits but also more time-steps during which syndrome measurement is activated. Fault-tolerant computation includes turning syndrome measurements on and off, so the main implication of being able to detect measurement errors is that whenever a syndrome measurement is turned off (i.e., a measurement qubit is not used), after reactivation it needs to stay active for several rounds of error correction—the more, the better.

7.2.2.4 Syndrome extraction

The errors that are most likely to cause a measured syndrome are found by Edmonds' minimum-weight perfect matching algorithm [Edm65], which basically matches all sign change events to pairs (for two sign changes in the same basis) or connections to a boundary with shortest possible chain lengths. In order to include measurement errors, which correspond to pairs of sign changes in time, the algorithm uses a three-dimensional space-time lattice, as shown in Figure 7.3.

The basic approach takes an equal error probability for all possible errors, always assuming the shortest possible path leading to a given error syndrome to be the right one—it involves the lowest number of errors. However, it has been shown [FWMR12] that, taking into account different probabilities for different errors to happen, more accurate handling of errors and less misinterpretation is possible.




Figure 7.3: Three-dimensional space-time lattice of syndrome measurement outcomes. One horizontal layer corresponds to one round of syndrome measurement, where the signs indicate the outcomes. Red lines show where a change of measurement outcome occurs. A single error (X or Z) of a data qubit leads to a neighboring pair of sign changes in a spatial dimension—with the faulty data qubit lying in the middle, a single error on the measurement qubit leads to a pair in temporal dimension—with the error happening between the two changes (M). Error chains lead to pairs of sign changes lying further apart [FMMC12]. "Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society."

7.2.3 Logical qubits and Pauli operations

The computational subspace still has two degrees of freedom, describing a logical qubit and one can find two operator products (linear independent of the stabilizers and commuting with them) to define a logical basis. A logical X operator can be defined by choosing a path connecting two data qubits from the two X-boundaries and performing X operations on all data qubits on the path - this does not change the outcome of any stabilizer measurement. The same can be done for a logical Z.

For a bigger number of qubits implemented in the same lattice, one can create holes in the array, meaning that some measurement qubits are turned off and do not measure the corresponding stabilizer of the surrounding data qubits anymore. These holes act like additional boundaries, and give extra degrees of freedom. Usually, logical qubits are implemented using two holes of the same type. This also leads to two types of qubits: The so-called double X-cut qubit (sometimes called rough or primal qubit) consists of two deactivated measure-X qubits, with the logical Z consists of a path of Z operations connecting both holes, and the logical X of a loop of X operations around one of the holes. An equivalent construction can be done

for the double Z-cut qubit (sometimes also called smooth or dual). Both types of qubits with their corresponding logical operators are shown in Figure 7.4.




Figure 7.4: Implementation of logical qubits: (a) Double Z-cut qubit, (b) double X-cut qubit. The logical operators XL (ZL) consist of X (Z) operations on the physical qubits along the blue (red) lines [FMMC12]. Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society.

Similar to detected errors, in realistic implementations logical X or Z operations do not actually need to be performed by doing all the single-qubit gates, but can rather be carried (and commuted) through the control software.

7.2.3.1 Distance

The distance d of the logical qubit is determined by minimum number of physical operations needed to perform a logical operator—up to $n = \lfloor (d - 1)/2 \rfloor$ errors can be corrected, larger order errors (i.e., including more physical qubits) might be misinterpreted as a logical operator. For even distances, an error chain of half a distance can not clearly be corrected since the decoder has to guess. Thus, one usually uses odd distance codes. The distance of a logical qubit can be increased by putting the holes further apart, although with this approach the maximal distance is limited by the number of qubits around the hole to $d = 4$. However, one can also make the holes bigger by turning off more qubits and thereby reach arbitrary distances. The number of physical qubits thereby is quadratic in the desired distance—not only must the two holes corresponding to the same qubit be separated far enough (for which an increase of qubits in one dimension would suffice), any pair of holes of the same kind (X/Y) must be separated by at least d data qubits to prevent from undetectable errors. However, a Z cut and an X -cut can be close because they have different kinds of boundaries, and an error chain beginning at one kind of boundary can not end at the other without leaving a trace in the stabilizer measurements.

In an array of only one kind of logical qubits, the number of physical qubits per logical qubit is quadratic in the distance, for single-qubit holes it is $n(d) = 8d^2$. For larger distance it is slightly more, according to the extra rows and columns needed due to bigger holes. We can lower-bound the number of required physical qubits by $n(d) = 8(d + d/4)^2$: We need holes of width $\geq d/4$ to ensure that an error chain around a hole has minimum length d and additional d data qubits to separate holes from each other. Additionally, we need a factor 2 due to the arrangement of qubits, a factor 2 for measurement qubits and a factor 2 since one qubit consists of two holes.

Keep in mind, that this is only the spatial distance: One must also preserve the distance in time, so for a distance d code, the temporal spacing of operations that involve turning off measurement qubits (like for example logical initialization or measurements, see Section 7.2.3.2) must also be at least d to distinguish between measurement and data errors. Measurement errors (any error on the measurement qubit, see Section 7.2.2.3) vanish after the next initialization, whereas a data error leads to a sign change that persists. If one wants to detect also the rare cases of measurement errors happening in a row, the number of measurement cycles must be larger than the biggest temporal error chain one wants to detect.

7.2.3.2 Logical initialization and readout

A logical X (Z) operator surrounding one of the holes of a double X (Z)-cut qubit is equal to the stabilizer operator that the deactivated measurement qubit in the hole would measure—if it was not deactivated. Therefore, creation and initialization of a logical qubit in one of its corresponding eigenstates can be done easily by just turning off the measurement qubits. The logical state is then given by the previous stabilizer measurement to be in one of the corresponding eigenstates (depending on the measurement outcome of the hole defining the logical operator).

Implementation in the other basis is more complicated. However, it is still easier than performing a Hadamard gate, as will be seen later. The implementation for a logical X (Z)-cut qubit into a Z (X) eigenstate is done in three steps:

1. Turn off all measure-X (Z) qubits along a path, exclude all data qubits along the edge of the thereby created hole also from Z (X) stabilizer measurements (see Figure 7.5 b,c) and perform Z (X) measurements on all measurement qubits of the hole
2. Initialize all measurement qubits of the hole to the same (desired) eigenstate of Z (X)

Turn all but two measurement qubits on again, leaving two holes at the edges and switch the stabilizer measurements to measure all four surrounding qubits again.

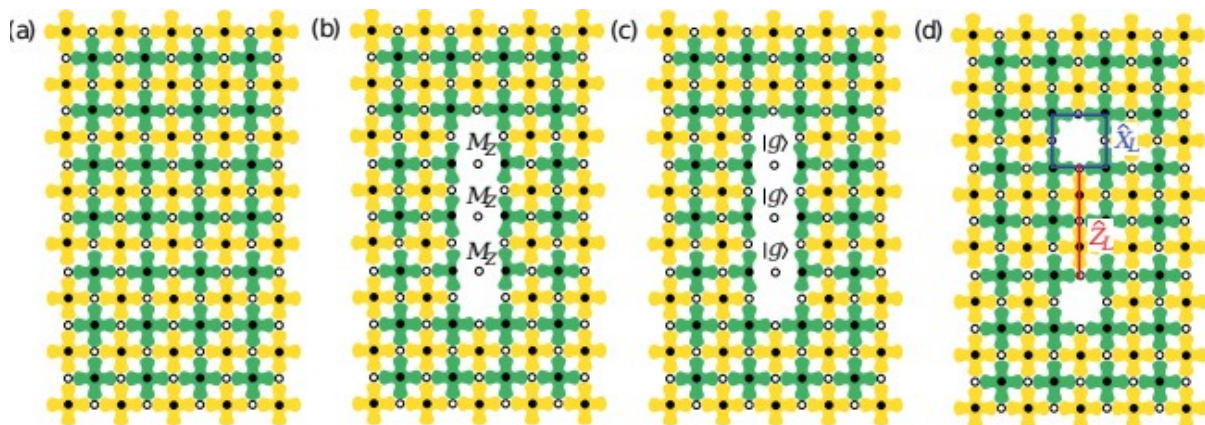


Figure 7.5: Schematic protocol for creating and initializing a double X-cut qubit in a logical Z eigenstate. M_Z denotes measurements in the basis of Z, $|g\rangle$ denotes initialization of the data qubits in the ground state [FMMC12]. "Reprinted figure with permission from. Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society

Measurement works similar to the readout protocol, the easy variant is to measure a double X (Z) cut qubit in the logical X (Z) basis and is performed by turning on the stabilizer measurements for the two holes again—the result of the stabilizer of the hole that defined the logical operator gives the measurement result. For measuring in the respective other basis, stabilizer measurements in between the holes are turned off (or to three-qubit measurements) again, Z (X) is measured for the data qubits along the path—determining the measurement result—before they are reset and all stabilizers are turned on again and the qubit is destroyed.

7.2.4 Logical gates: H, T, CNOT

7.2.4.1 Multi-qubit gates

An interaction between different logical qubits can be achieved by moving holes around the physical lattice. However, the logical CNOT can only be performed between two qubits of differing kind. Usually, this problem is solved by having mainly one kind of qubits, and using the other kind only to transmit the logical CNOTs.

Moving one hole to another position (preserving the logical qubit information) along an arbitrary number of cells in the lattice takes two surface code cycles for the actual move (+ $d-1$ cycles to preserve fault-tolerance), independent of how far the hole was moved: In the first step, the hole is enlarged up to the final position for the move (so that the initial and the desired final position get linked with a chain of deactivated stabilizers), additionally the data qubits that lie fully inside this large hole are measured. In the second step, the hole is shrunk to its original size, but at the new position by activating all other stabilizer measurements again. To assure fault tolerance at distance d one needs to wait another $d-1$ steps to identify measurement error chains, so the whole action takes $d+1$ cycles.

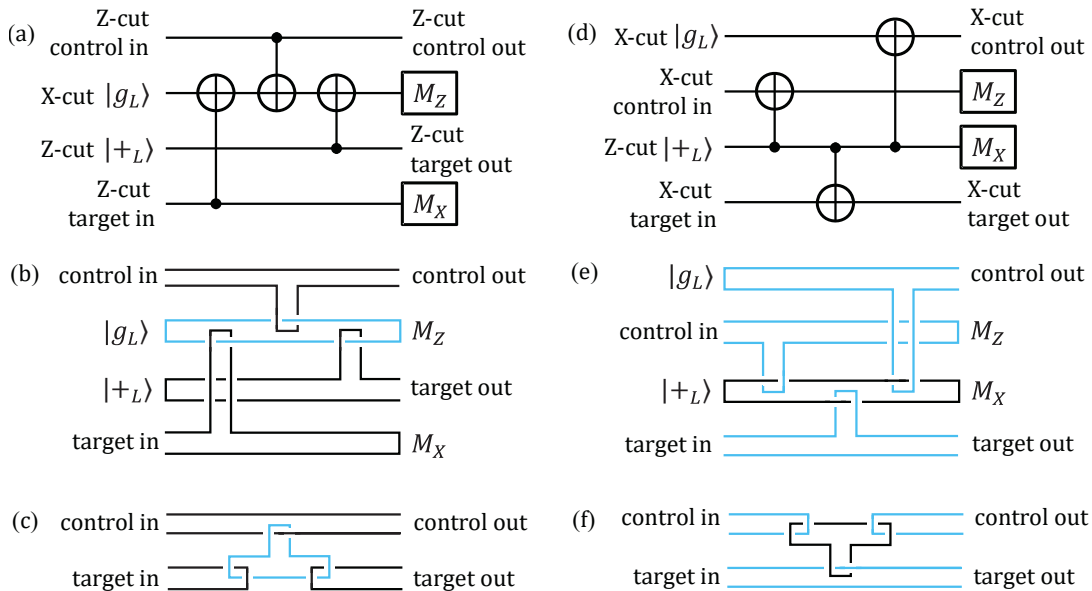


Figure 7.6: (a) Circuit diagram for a logical CNOT operation between two double Z-cut qubits, mediated by a double X-cut qubit. During the process, the target qubit is measured, and a new double Z-cut qubit is initialized in $|+\rangle$ to take the place of the target qubit. (b) Description of the braiding of holes that is done to perform the three CNOT steps: Every double Z(X)-cut qubit is represented by a pair of black (blue) lines, where the movement of the holes in time is shown along the x-axis. Two lines corresponding to two holes of the same qubit join when the qubit is initialized or measured. (c) Simplified representation of the braiding, showing the double X-cut qubit only as an intermediate tool for the gate. In fact, the double Z-cut qubits do not need to be moved at all and the new target qubit can be initialized at the position of the measured old one. (d)-(f) Equivalent representations for an indirect CNOT between two double X-cut qubits. [FMMC12] Reprinted figure with permission from. Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society

In order to perform a logical CNOT, one needs to move a hole of a double X (Z) cut qubit in a closed loop around a double Z (X) cut qubit, ending up in the initial position after two separate moving steps (i.e., $2(d + 1)$ surface code cycles). A direct CNOT between two qubits of the same kind is not possible, but one can use one type of qubit to mediate the gate between two qubits of the other kind, as shown in Figure 7.6. Therefore,

three concurrent logical CNOT operations ($6(d+1)$ cycles), two logical measurements and initializations (only constant time cost) and one additional logical qubit (but with different boundary, so without extra space cost) are required. Measurement and initialization are always in the complicated basis with respect to the qubit type. Given that holes can be moved arbitrary far, there is no fundamental constraint on how far the target and control qubits are away from each other. An application of an CNOT between two qubits of the same kind does not require any physical qubit overhead since the mediating qubit can be created arbitrary close to the other qubits. A schematic representation of the stabilizer measurement pattern for a mediated CNOT between neighboring qubits is shown in Appendix 25.2.

In the same manner, also multi-target CNOTs are possible by just braiding a mediating qubit around more than one target qubit. This can be done in the same amount of steps as the regular CNOT.

7.2.4.2 Hadamard

The Hadamard gate does not need additional qubits but additional physical gates on the underlying physical qubits between the syndrome measurement cycles. Besides multiple ($O(d)$) stabilizer measurements after some of the actions to preserve spatial distance, one needs to

- isolate an area around the targeted qubit by turning off stabilizers on a ring large enough to preserve the distance inside³
- deform the operator loop around one of the holes to an operator chain connecting the two new outer boundaries
- reduce the size of the isolated area to a $d \times d$ (data qubit) array between the two holes by deactivating even more stabilizer measurements
- perform Hadamard gates on all data qubits in the isolated region (simultaneous)
- perform SWAP operations between all data qubits and their neighboring measurement qubits (this happens in two steps operating first between vertical and then horizontal pairings - during each step, all SWAP gates are performed simultaneous)
- turn on some of the stabilizer again to create two holes and deform an operator chain to get a double cut qubit again, but now rotated by 90°
- move the holes to rotate the double cut qubit to its original orientation
- turn the stabilizers on the isolating ring back on.

Thus, the logical Hadamard gate can be performed in $O(d)$ overall time steps, requiring the ability to perform simultaneous physical Hadamard and SWAP gates, respectively. A detailed study of the logical Hadamard gate can be found in [FMCM12], or in a more schematic way in Appendix 25.2.

7.2.4.3 S and T gate: Short qubits and magic state distillation

The T gate, which brings the so far available set of gates out of the Clifford group, thus making computation universal and classical not simulable, is the most challenging one. The T gate, as well as the S gate (with $S = TT$) needs an ancilla qubit prepared in a so-called magic state (which is not an eigenstate of X, Y or Z) and logical Hadamard + CNOT operations between the ancilla and the targeted qubit. The S gate can be done deterministic with one logical Hadamard and two logical CNOTs. The T gate, which requires only one logical CNOT, however, only works in 50% of the attempts and performs the logical operator T^\dagger in the other case. This can be detected by measurement of the non-used qubit (the desired state is on the ancilla qubit in the end - the initial targeted qubit can be measured) and compensated for by a subsequent S gate. Both the S

³ The ring can be placed directly next to the surrounding qubits without the need to include additional physical qubits anywhere. This is no problem in terms of fault tolerance if the surrounding logical qubits are all of the same type since the boundary of the ring is set to be a different one.

and T gate implementations are described in Figure 7.7 as a general circuit diagram—in the surface code implementation, everything will be on a logical level.

The hardest part is the fault-tolerant preparation of the magic state ancilla. There exists no logical gate or initialization protocol to directly create the desired states for a distance d qubit when $d > 1$. Thus, the only way is to use a distance 1 qubit for the time of initialization (for which a physical operation on the data qubit between the two cuts is equivalent to a logical operation on that short qubit) and enlarging it immediately to the desired distance. This procedure will inevitably lead to a reduction of fidelity for this qubit given by the actual physical error rate of the qubit between the cuts, but the states can be made much more precise through distillation, for example with the Steane or Reed-Muller encoding [BK05] or the more recent Bravyi-Haah protocol [BH12]. These codes use multiple faulty (but in our case still logical) qubits to create one or several more precise qubits, a process which can be repeated until the desired fidelity is reached. In the first round, the logical magic states from the distance-1 creation process are used, subsequent rounds use the former purified output states to generate an even more precise state. The distillation codes need only—eventually multi-target—(logical) CNOT operations, (logical) state preparations and (logical) measurements in Pauli eigenstates, so the overhead is mostly due to the high number of logical ancillae and repetitions needed. The exact distillation circuits can be found for example in [FMMC12, Section XVI].

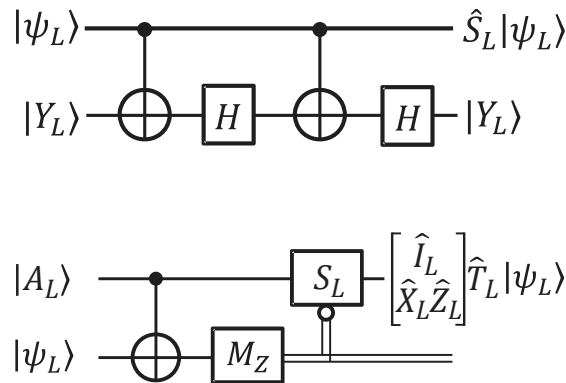


Figure 7.7: Implementation of S (top) and T (bottom) gate on the input state $|\psi\rangle$ with magic states $|Y\rangle$ and $|A\rangle$, respectively. In a more recent version, the S gate can also be performed without the final Hadamard gate, carrying a byproduct operator in the classical control [GF17]. The T gate additionally needs a conditional S gate to correct its non-deterministic nature. The classical process of deciding whether or not to perform the additional S gate after measuring MZ is represented by double lines. When the S gate is needed, the final state will be $XZT|\psi\rangle$, but the X and Z byproducts can be carried in the classical control. Reprinted figure with permission from [FMMC12] Copyright (2012) by the American Physical Society

The Steane code [Ste96], distilling the (accurate) ancilla required for an S gate needs 7 (faulty but logical) ancilla qubits with an error rate p plus one logical qubit in the state $|+\rangle$ (which can be created transversal and therefore fault-tolerant) and creates an output state with error rate $7p^3$, with a probability of $1-7p$. The process is probabilistic, but the measurement outcomes indicate whether the distillation was successful or not - in which case it has to be repeated with new ancillae. A second repetition will lead to an error rate of 7

$\times (7p^3)^3 = 7^4 p^9$, though requiring $7 \times (7+1) = 56$ ancillae⁴. This process can be extended to arbitrary length, the error rate scales with the number of repetitions as $P_1 \sim p^{\wedge}(3^n)$, requiring $\sim 7^n$ (faulty) ancillae. One round of distillation can be fit into a space-time volume of $18 \times (5d/4)^3$ (surface code cycles \times physical qubits) when implemented on distance d logical qubits [FD12].

The Reed-Muller code [BK05] used for the T gate ancillae works similar: It creates an output state with error rate $35p^3$ with a success probability of $1-15p$ from 15 initial ancillae with error rate p and one $|+\rangle$ state. The (physical) space-time volume of one distillation round within a distance d code is $192 \times (5d/4)^3$ [FD12]. Additionally, keep in mind that in 50% of the cases an additional high-fidelity S gate is required to complete the T gate. The volume for the creation of an accurate S ancilla is low compared to the T state and can therefore be neglected in overhead calculations.

In typical applications, two distillation cycles are enough to reach the desired error rate. With that, the creation of one magic state ancilla for the T gate requires $15 \times 16 = 240$ logical ancillae. In the first round of distillation, the overhead can be reduced by using lower distance qubits that match the achievable error rate of the ancillae after only one distillation round.

Alternatives Instead of distilling ancilla states for performing T gates, it is also possible to distill some (not all) other states that might be required, depending on the algorithm to be performed. One possibility is the creation of ancilla states for Toffoli gates, which are for example required in factoring algorithms. A Toffoli gate can be constructed from several T gates + Clifford gates, however, instead of creating an ancilla for each of those T gates, one can also distill a state with which one can perform the whole Toffoli gate. This is more efficient in many cases [OC17, see for example Table I], but can just lead to a constant factor improvement. The main advantage is time, since the Toffoli gate then only needs one time-step with feed-forward, instead of 3 or four steps for each sequential T gate [AMMR13].

7.2.4.4 Ancilla factories

For algorithms using bigger amounts of T gates, the creation can be further optimized by creating the ancillae in parallel (for example using the qubits of the first round already for a new distillation in the second round) and offline (if all magic states created in a separate part of the code to be readily available when needed, state distillation does not necessarily introduce any time overhead to the calculation process). Depending on which distillation protocol is used, one $[n,k]$ ancilla factory will produce k high fidelity states from n input states in one round, with a ratio n/k that can be significantly better than 15. One example is block code state distillation: The Bravyi-Haah code [BH12] distills $n = 3k + 8$ input qubits into k outputs with error reduction $p \rightarrow (3k + 1)p^2$ and a success probability of $1 - (3k + 8)p$. The corresponding space-time volume is $(96k + 216)(5d/4)^3$ [FDJ13]. This code only offers a moderate overhead reduction by a constant factor, compared to Reed-Muller. For bigger systems, the worse scaling of error reduction (quadratic instead of cubic) suggests even less benefit. There are optimization approaches improving the efficiency of block-code state distillation, like for example via module-checking [OC17]. However, any improvement can only lead to a constant factor reduction of space time costs.

Since the ancillae required for state distillation all need to be logical qubits (except at the time of creation), non-Clifford gates are the ones that consume the majority of physical qubits in a computation process, according to an estimation [FMMC12] for a 2000-bit factoring algorithm, state distillation occupies around 90% of the total number of qubits. Furthermore, also the application of all distillation gates embedded in the surface code will introduce an additional error to the final states, depending on the distance that was used.

7.2.4.5 Magic state injection

It has been shown in [Li15] that the fidelity of a raw logical magic state can be higher than that of the unprotected (physical) qubit operations creating it. By increasing the distance of the magic state qubit step by step, the final error rate can be made $p_1 = 2/5p_2 + 2p_1 + 2/3p_1 + O(p^2)$, with $p_{1(2)}$ being the single (two) -qubit

⁴ 7 sets of 7 magic state ancillae and one $|+\rangle$ state per set.

gate error, and p_i the initialization error. In typical setups, two-qubit gate errors are by far the largest ones, so the injection error can be approximated by $\approx 0.4p_2$.

7.2.5 Lattice surgery

An alternative to the described encoding of logical qubits with defects and braids, i.e., in a topologically planar way, is lattice surgery. Lattice surgery here refers to performing operations by cutting and stitching of respective logical qubits as described in [HFDvM12] showing that the storage overhead can be reduced significantly.

The preprint [FG18] reviews how to perform *all* necessary operations for universal fault tolerant quantum computing as well as state distillation using a lattice surgery style qubit encoding. Further they include a section where they talk about state distillation and how to realize it with planar logical qubits using lattice surgery. The main message here is, that it is possible to perform state distillation using lattice surgery in a very intuitive way and again save overhead compared to the usually used two defect logical qubits (mostly because of the intrinsic rotation property of planar qubits).

The authors show for one specific dataset (10^8 T gates and 100 logical qubits on a hardware with gate error rate 10^{-3} and surface code error correction time of $1\mu\text{s}$) a defect and braiding algorithm needs 4.5 h and 1.8 million physical qubits, whereas using lattice surgery the same algorithm would use 5.4 h but only 0.37 million physical qubits. This is a promising development, yet, in order to become part of our evaluation scheme, a scaling analysis beyond this one dataset needs to be performed by the community.

7.2.6 Volume compression and time-optimal computation

We can draw the defect structure (i.e., the positions of deactivated syndrome measurements, see Appendix 25.2 for schematic representation) for any logical gate sequence, including initializations and measurements in a three-dimensional diagram, where the third dimension represents time. A logical qubit is then represented by a pair of tubes of width $d/4$ separated by d , corresponding to the physical qubits needed to encode one logical qubit of distance d (the actual number is 4 times higher due to the syndrome measurement structure). Moving holes to other positions is represented by a tube of length d (i.e., during $d/4$ time-steps, d additional qubits are turned off) in a spatial direction, after which one needs to wait again d steps to prevent measurement errors, i.e., a connection of length d in temporal direction. The basic building block thus has an edge length of $d + d/4 = 5d/4$: d for a waiting or moving tube and $d/4$ for merging them. In such a diagram, topological equivalent structures perform the same operations. Thus, by deforming the structure, it is possible to significantly reduce the space-time overhead of a Clifford gate sequence. A CNOT or Hadamard gate requires an overall space-time volume of $12(5d/4)^3$ in a highly compressed form [FD12]. By increasing the number of qubits, it is also possible to deform any Clifford gate sequence such that increasing the number of gates does not increase the time required for execution (but only the number of qubits). Non-Clifford gates need classical feedback and therefore require a certain time-ordering which must be protected, so they can not be performed in constant time.

In order to reduce the computation time for performing a large quantum algorithm to realistic and useful values, it is most important to optimize a quantum circuit to have lowest possible execution time. It has been shown [Fow12], that arbitrary large Clifford circuits can be performed in constant time by making the 3D structure flat and effectively performing operations in parallel. Furthermore, since the only time-step of a T gate that can not be eliminated is the measurement with classical feedback, a circuit consisting of n T gates can be executed in asymptotic time nt_M with t_M being the physical measurement time + (negligible) classical feedback. Measurements are not necessarily limited to be done during the measurement step of a surface code cycle, so multiple T gates can be performed during one cycle. A typical approximation is $t_M = 0.1t_{SC}$ for a surface-code cycle time t_{SC} .

7.2.7 Performance

It is possible to reach arbitrary low logical error rates by increasing the code distance, as long as all physical error rates per step (*physical* initialization, gates and measurements) are below a certain threshold. The underlying error model assumes [FMMC12, Section VII]

- the probability to initialize a qubit in a state orthogonal to the desired one to be p
- the probability to perform a single-qubit Pauli operator X, Y or Z on a data qubit when intended to do the identity (waiting) to be $p/3$ each
- the probability to perform an additional single-qubit Pauli operator X, Y or Z on a measurement qubit when intended to do a Hadamard to be $p/3$ each
- the probability for performing a tensor product of two Pauli operators, of which maximally one is the identity when intending to perform a CNOT between data and measurement qubit to be $p/15$ each
- the probability for reporting and projecting into the wrong eigenstate after measurement to be p .

Hence, p describes the probability for a Pauli error per step (of the error correction circuit) and per physical qubit, so the overall error probability in one complete error correction round is $\approx 8p$. Appendix 23 gives an analysis on how multiple errors add up during one cycle.

The threshold was found to be $p_{\text{th}} \approx 0.57\%$ for this error model [FMMC12], in the original publication of the surface code it was given as $p_{\text{th}} \approx 0.75\%$ [RH07], the actual value depends on the error model and underlying assumptions. The lower the physical error rate, the less qubit overhead is required to reach the same logical error rate. For physical error rates p much below the threshold and odd distances, the logical error rate was approximated by the empirical formula $P_l \approx 0.1(p/p_{\text{th}})^{(d+1)/2}$ [FDJ13]. Note that in this formula, p_{th} is not necessarily the actual threshold, since for error rates close to the threshold the approximation breaks down and p_{th} is the value where the fit lines (and not the actual curves) for different distances cross, see Figure 7.8. In literature, this value is sometimes referred as threshold under ideal syndrome extraction. Furthermore, the parameters for p_{th} and the prefactor might vary with the error model and the strength of syndrome extraction. [Fow13a] presents a way to find an analytic expression for the asymptotic performance of a code without the need for computational time-consuming simulations.

The important point that we learn here is that the logical error rate scales quadratic with the physical error rate for a distance 3 code, cubic for a distance 5 code, quartic for distance 7 and so on. This dependency can be intuitively understand by associating the logical error rate with the probability of an uncorrectable error chain (which has minimum length $(d+1)/2$). If we ignore non-Clifford gates for now, then the distance scales quadratic with the number of physical qubits: The simplest implementation of a logical qubit needs $(2d-1)^2$ physical qubits [Mar15]⁵. It is important to keep in mind that this is the performance for a surface code only doing Clifford gates. The performance of the S and T gates depends on the distillation process, which has to be matched to the desired error rate (i.e., the error rate that the surface code can reach), requiring more overhead than the logical qubits for the rest of the calculation, see Section 7.2.4.3.

5 For an implementation of multiple double-cut qubits in a lattice it is a bit more, but still quadratic




Figure 7.8: Performance below threshold for the surface code for distances 3, 5, 7, 9, 11, 15, 25, 35, 45 and 55. For distances 3, 5 and 7, quadratic, cubic and quartic fit curves are shown as dashed lines. They only approximate the actual curves for low physical error rates p [FDJ13]. Reprinted by permission from Macmillan Publishers Ltd: Scientific Reports (A. G. Fowler, S. J. Devitt, and C. Jones, *Sci. Rep.*, 3(1), 2013), copyright (2013).

7.3 Basic requirements

In [FMMC12], an exemplary analysis of the quantitative requirements for running a factoring algorithm was done: Assuming the physical error rate to be 10% of the threshold error rate, factoring a 2000 bit number requires $2 \cdot 10^7$ physical qubits for the logical qubits, and $2 \cdot 10^8$ additional physical qubits for the distillation ancillae, i.e., a total number of $2.2 \cdot 10^8$ qubits. With measurement times of 100ns, the computation will run around 27h. In the end of this chapter, section 7.2.7, we will turn this into a general conversion formula for fault-tolerant execution of algorithms, resembling the recent concept of quantum volume brought forward by IBM [BBC⁺17].

The qualitative requirements for the surface code to be implementable can be classified in three levels: The physical qubit level, the experimental setup and the classical control. For the **physical qubits**, the requirements are:

- The ability to initialize and frequently reinitialize the qubit in at least one basis
- The ability to perform single qubit gates, at least Pauli, Hadamard and T gates
- The ability to measure in at least one basis
- The ability to perform two-qubit SWAP and CNOT gates, at least for nearest-neighbors on a two-dimensional square lattice
- Error rates for all possible operations (gate, waiting, initialization and measurement) significantly lower than the threshold of approximately 1% that do not increase with the qubit number / array size. This does not mean that the coherence time of the qubit needs to exceed the whole computation process, but it must be long enough to ensure sufficient low error rates after the time required for one operation. Trade-offs are possible (see Section 7.4.2)

- No qubit losses occurring at all (i.e., photons or atoms disappearing)⁶

Further desirable properties are a low probability of parallel, correlated errors and leakage errors. However, these errors only lead to slightly worse performance or can be corrected with some additional effort (see Section 7.4.1). Requirements on the **experimental setup** are:

- A large number of qubits arranged in a 2D square lattice existing simultaneous and long enough, all fully controllable and all fulfilling the above requirements. This implicates
 - large space at sufficiently low temperature available
 - long timescales to be reached: Even though the coherence time of the qubit can be lower, the qubits must at least exist through the whole computation process (see for example trapping times of ions / neutral atoms as described in 16.1.2.1), also cooling and isolation must be available long enough
- Simultaneous measurement, initialization and gates (Hadamard, SWAP, CNOT)⁷:
 - For a lattice of $2n$ physical qubits (n data and n measurement qubits), n operations at the same time, being either initialization, measurement or CNOT are required. Note that if only one basis is experimentally accessible, measurement and initialization include Hadamard gates.
 - Simultaneous SWAP gates are limited to the region of the logical qubits on which a logical Hadamard shall be performed (which can be several at the same time, though).

The basic requirements on the **classical control** are pretty much met already, using Edmonds' minimum-weight perfect matching algorithm [Edm65] and optimization tools [FWMR12]. It has been shown [Fow15] that, given a 2D array of (parallel) processing elements with communication between nearest neighbors and an external memory, parallel decoding can be performed in constant $O(1)$ time for the surface code, independent of the number of qubits. Thus, the main challenges are large parallel computation, enough memory and most important processors fast enough to preserve the low error rates of the qubits (i.e., significant faster than coherence and gate times) and to retain reasonable computation times.

7.4 Performance discussion

7.4.1 Error Model, implicit assumptions and their limitations

Most large-scale error correction simulations that are used to numerically determine the error correction threshold do, if they are not tailored to a specific architecture, assume a very simple error model: All operations involved (gates, readout, initialization) are assumed to have the same error probability (also called the error rate as in probability per operation). As described in section 6.5, there is more structure to physical error models. In most simulations, it is also assumed that the error is fully stochastic, i.e., described by error operations that occur with a certain classical probability as mathematically outlined in chapter 6.5 and appendix chapter 22. It is furthermore assumed to be described by a depolarizing channel, i.e., by a map of the quantum state onto the fully mixed density matrix with probability p . While taking this model is justified for the sake of modeling the generic performance of error correction, it does not describe all error mechanisms, specifically those discussed in section 6.5. We discuss the relevance of these approximations as much as they are covered in literature.

7.4.1.1 Conversion between Pauli error rates and realistic errors

As the theory of quantum error correction is described in terms of Pauli errors, it is necessary to describe every possible experimental deviation from the ideal case in terms of those Pauli errors happening at

⁶ Quantum error correction with topological cluster states can deal with qubit losses [WF14b]

⁷ This requirement might be softened if gates are fast enough to make partially sequential runs (not scaling with the lattice size) acceptable.

different stages of the computation. Depolarizing errors can be directly modeled by some Pauli error happening with a certain probability in one surface code cycle. However, there are also experimental imperfections that are not Pauli operators, like for example small deviations from the ideal direction of a gate or state (see Appendix 22 for further discussion). In principle, any unitary deviation can be expressed as a rotation in the Bloch sphere of the basis in which a state, gate or measurement is described / performed. Such a rotation (of angle θ around the axis $(\alpha, \beta, \gamma)^T$) can always be written as a linear combination of Pauli matrices $R = \cos(\theta/2)1 - i\sin(\theta/2)(\alpha X + \beta Y + \gamma Z)$, with $\alpha^2 + \beta^2 + \gamma^2 = 1$, $\{\alpha, \beta, \gamma\} \in \mathbb{R}$. The probability for a bit-flip error for example can now be deduced from all operators causing a flip when measuring in the $|0\rangle, |1\rangle$ -basis: $|\langle 1|R|1\rangle|^2 = |\cos(\theta/2) + i\sin(\theta/2)\gamma|^2 = 1 - \sin^2(\theta/2)(1 - \gamma^2)$. Hence, the probability for a bit-flip error to happen (either alone or in combination with a phase-flip, which will have no effect in this case) is $\sin^2(\theta/2)(1 - \gamma^2) = \sin^2(\theta/2)(\alpha^2 + \beta^2)$. The total probability for any error to happen is $\sin^2(\theta/2)$. If gates are performed before the measurement, the operator R can be permuted until it is right before the measurement, for Clifford operators this will only change the Pauli components of R to other Paulis.

In the case of measurement bases that are systematically rotated a small angle away from the exact direction, an error needs to be modeled during initialization *and* measurement, at least if the measurement is used for initialization, too (see Appendix 22.3). If there are gates G_1, \dots, G_n applied between initialization and measurement, the channel of operations acting on the qubit, if assumed to start and end in an ideal basis are $R^\dagger G_n \dots G_1 R$. However, we can also permute until both errors are next to each other and express this product as a new rotation. The same is the case for a gate defined on a wrong basis. Instead of writing $R^\dagger G R$, we can also write $R_{\text{eff}} G$ to model the rotation by a single error following the gate. This way, any noisy state preparation or noisy gate can be modeled by the ideal operation followed by an error operator. When talking about gate or initialization errors in quantum error correction, it always means modeling the faulty process by the ideal one followed by a single error.

Usually, all these error processes do not need to be considered separate, since for error correction, the exact form or origin of an error is not important: What matters is the bare probability for getting the wrong result (and thereby also projecting to the wrong state) during a measurement. Good fidelity measurement protocols are designed such that they include all possible error processes by averaging over a set of different actions (see Section 6.4 for more exact discussion).

For further reading on the treatment of different error channels see for example [LB13, Section 1.4].

7.4.1.2 Simplifications within stochastic errors

The error model on which the standard performance discussions and syndrome extraction algorithms are based on is a rather simple and unrealistic one: Every kind of error, faulty single-qubit gate, two-qubit gate, initialization, measurement and waiting period (identity gate), that can happen is assumed to happen with equal probability. In this section, we discuss more realistic models that are still based on stochastic error models whereas in the next section we discuss the impact of coherent errors. Given the sometimes huge difference already between single-qubit and multi-qubit gate fidelities, this is far from reality and can not reflect a realistic quantum system. Also, as described in section 5.3, incoherent errors are not always depolarizing. More generally, the transversal decay time T_2 is different, often shorter, than the energy decay time T_1 . Also, at low temperature, energy relaxation is asymmetric and the final state the system decays into is not fully mixed. This is not a problem for syndrome extraction per se - if all error rates lie below the threshold (no matter how far below), error correction will still work. The reason of this can be traced back to randomization: A general quantum algorithm is random enough to make the error channel effectively isotropic. This principle is, in fact, used in the randomized benchmarking method, see section 6.4.3. But if one understands the error processes better, we see that it is unnecessary to restrict all errors to the same threshold, since some trade-offs are possible here. It also enhances the syndrome extraction significantly if the exact error rates are known, Autotune makes use of this fact [FWMR12] and reaches highly improved logical error rates. Figure 7.9 shows two plots for the performance of the surface code on a realistic error model, one with the usual syndrome extraction, and one with the help of Autotune. The enhancement

comes from the fact that with realistic error models, the shortest path is not always the most probable one, but the unoptimized algorithms can only find the shortest paths.

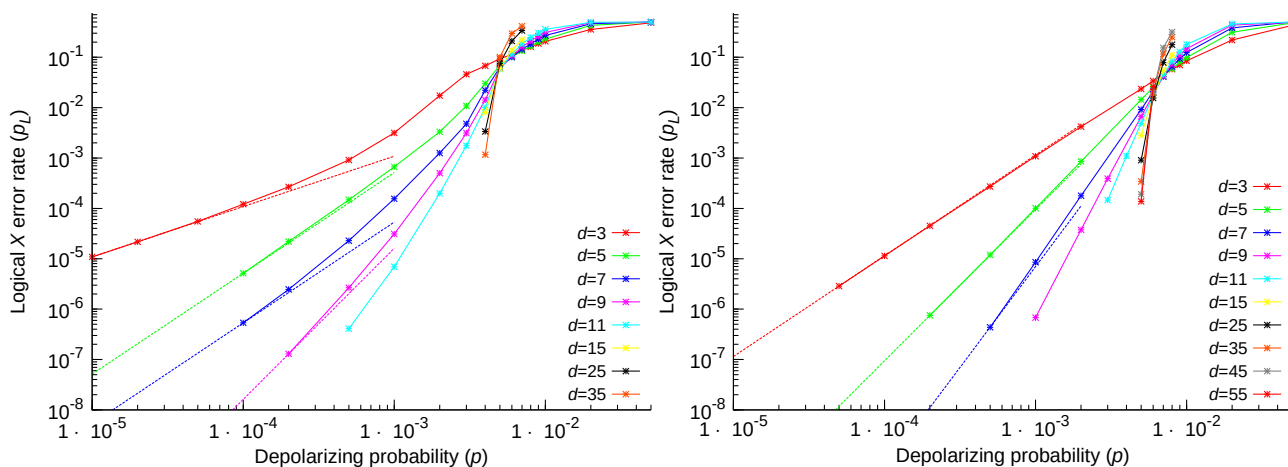


Figure 7.9: Another two threshold plots indicating the threshold at the crossing of the different lines.

With the Autotune library, available on the website of the Topological Quantum Error Correction Group in Melbourne, it is also possible to calculate the actual performance for any given implementation with given error rates. According to [Mar15], measurement errors can be the highest ones, single and two-qubit errors can be around the same range, therefore it is most important to optimize two-qubit errors: With that, the required single-qubit error rates will usually be automatically met. An extreme example of the use of error asymmetries is the result [TBF18], described below under “other error correcting codes”.

The concrete numerical results in this study take these error asymmetries into account if they are known, specifically in chapters 13.2.3 and 14.2.3. In the chapters where we extrapolate on the performance of future devices, we use the generic simplified error model, specifically in chapter 21.

Furthermore, error and overhead calculations are often done for implementations of only one logical qubit along the whole array, therefore not requiring any double-cuts and needing less qubit overhead. This does not drastically change the performance but should at least be kept in mind. All experimental implementations of surface codes so far do the same and only demonstrate single logical qubits. Showing a fault-tolerant two qubit gate is a major program goal. This is a step in the right direction, but does not show the ability to actually perform logical operations on multiple (or even single) qubits. Furthermore, there have never been any implementation shown that can cope with the huge amount of qubits needed for relevant calculations. Thus, we do not know how the fidelities and coherence times of each physical qubit change when embedded in such a huge cluster. Interaction effects, cooling or addressing problems might occur.

Of course, these simulations (and implementations) also do usually not include the ability for non-Clifford gates: The distillation error rates need to be matched to the respective error rate the surface code can perform. Most of the qubit overhead during a calculation is caused by the ancillae needed for state distillation, so the overall qubit overhead does not only depend on the code distance but also strongly on the number of distillation rounds - which has a different dependency on the error rates than the distance (see Section 7.2.4.3).

A further problem are errors that are not at all included in the error model, such as leakage, qubit loss, or correlated errors.

Qubit loss can not be treated in the surface code, but with a variant, the topological cluster states [WF14b].

Leakage errors, that for example appear in superconducting qubits with more than two energy levels, can not be corrected with the standard setup of the surface code but after slightly adjusting it as described in [GF15]. By using SWAP operations between data and measurement qubits after each round of error

correction, each qubit is a measurement qubit every other round and therefore is reset from time to time (during the initialization for the stabilizer measurement). So, even if a qubit leaks out of the computational subspace, it will be brought back during its next round as measurement qubit. This adjustment needs only $2d-1$ additional physical qubits for a distance d code, and additional time cost independent of d . For benchmarking protocols accounting for leakage, see Section 6.4.3.

Correlated and multi-qubit errors were analyzed by Fowler and Martinis in 2014 [FM14]. They found that the surface code is rather robust to such kinds of errors. For local errors, and exponential suppression of large-area errors it is sufficient, typical experimental errors can be compensated for with negligible qubit overhead. Long-range two-qubit errors can be corrected with even less overhead, although there is not always a threshold⁸ depending on how strong correlations are suppressed. However, error rates low enough for reasonable computation can be reached. Experimentally, the important thing is to observe the error rates when a maximum number of qubits is manipulated in parallel, to see whether an implementation is capable of reaching reasonable logical error rates or not. Further discussion on correlated errors can be found in Appendix 22.4.

7.4.1.3 Resilience against coherent errors

Next to stochastic errors - errors in which an unobserved environment also changes states and gets entangled with the qubit (see chapter in the appendix 22.7 for a discussion of the relation between entanglement and incoherent errors) there can also be coherent errors as outlined in 6.5. We describe in section 6.5 in what sense systematic errors have to be coherent. It can be argued that their impact is often comparable to incoherent errors: On the one hand, operating a quantum processor faster reduces the impact of (unstructured) decoherence, on the other hand, it increased the frequency bandwidth for uncontrolled degrees of freedom.

What can we understand about their relation to quantum error correction? Remarkably, as discussed in the appendix chapter 22.2, the probabilistic nature of quantum physics suggests a similarity of these errors using Born's rule and errors are still detected as uncorrelated errors are detected by stabilizer measurements. One needs to ask the question on whether this affects more global conclusions.

The literature situation here consistently points out that the basic notions of error corrections - existence of a threshold and high effectiveness of error correction below threshold - are still intact. On top of this, the paper [BWG⁺18] points out that the measurements involved in error correction turn errors incoherent (consistent with our appendix 22.3) hence expecting that for large code distances both models approach each other. It thus confirms the earlier paper [GD16]. The paper [BWG⁺18] on the other hand shows that error correction for coherent errors is more effective than the diamond-norm of the coherent error suggests. This should, however, not be taken as a reason to dismiss coherent errors, as these are an example for the diamond norm suggesting a much larger error than the average fidelity, with the latter being easier to estimate in experiments.

The article [BEKP18] makes significant advances in tackling the questions on the effect of such coherent errors on the efficacy of surface codes. In particular, it contributes a new highly efficient algorithm to simulate surface codes and the results of corrections, applicable to arbitrary decoders. (data is shown for qubit number up to 37).

The main result of the paper is to approximately identify the error threshold of the surface code under the condition of purely coherent errors. They find the value to be approximately at $0.08\pi < \theta < 0.1\pi$ (where θ is an index for error threshold as it is the amount of rotation around Z the logical qubit undergoes). This result is significant as the threshold for the Pauli-twirl approximation (where coherent noise is replaced by its Pauli-twirled version) is placed at around 0.1π . It is hence shown that the current estimates for thresholds are valid

⁸ No threshold does not mean that no improvement is possible at all. But for a given physical error rate—no matter how low—it is not possible to reach arbitrary low logical error rates.

even for coherent errors. The translation of this error parameter to probabilities follows Born's rule, consistent with our Chapter 22.2.

This confirmation does not hold when considering the performance of the code below-threshold. It is, in fact, evident from the paper, that surface codes will perform worse than expected for $\theta < 0.1\pi$: i.e., the inverse scaling of logical error to code distance d is much slower in reality than estimated using the twirl approximation. Hence, the overhead of the surface code (number of qubits needed to reduce logical error) is heavily effected by coherent noise.

Furthermore, the study only considers coherent errors in the form of unitary rotations about the Z axis (as dictated by limitations of the algorithm for storage). Although it is not clear what the effect would be of considering arbitrary rotations (or even just X,Y, and Z rotations), it is doubtful that the overhead would improve, in fact, it would most probably worsen. Whether the threshold results would change, too is unclear, however the positive results of the paper call for optimism.

In conclusion, the article brings confidence in the threshold estimated using the Pauli approximation of coherent errors but warns of the lowered scaling performance on the fault-tolerant side.

As a conclusion to this section, while coherent errors do not invalidate the general results of error correction, they can affect our quantitative results, with no consensus in the literature on how far, as only few and restricted analyses have been performed. In order to gather more complete understanding of this problem, further work is needed. This should on the one hand analyze the error budget of experiments, on the other hand use theoretical tools to refine the different predictions about correction of unitary errors by larger scale theoretical studies.

7.4.2 Possible Trade-offs

There are several trade-offs between physical error rates, qubit overhead and computation time, making different approaches possible.

For both the bare surface code and the distillation, the relations between qubit overhead and error rate improvement are exponential. The exact relations differ, but both processes can be done with less overhead if lower initial error rates are available; but also with higher initial error rates allowed for the cost of more qubit overhead, depending on the underlying physical conditions.

A significant effect on the threshold can be reached with particularizing the error model to realistic values. Already early calculations [WFH11] have shown that using a model with error rates typical for ion trap implementations, i.e., $p_2 = p$, $p_1 = p/1000$ and $p_M = p/100$ for two-qubit, single-qubit and measurement errors, respectively, the threshold condition could be raised to $p < 1.4\%$. Using Autotune [FWMR12], acceptable measurement errors of up to 10% were reported, provided that two-qubit gates can be performed with a rate of $p_2 = 0.1\%$.

Knowing that logical T gates are the most demanding elements of fault tolerant computation, it also makes sense to think of algorithms that, even if requiring more qubits or computation steps, need less total T gates. This is a problem of finding trade-offs on the code side. It has been shown [BK13] that in general, 2D codes can not implement non-Clifford gates without turning off the topological protection at some points and thereby reducing the fault-tolerance. For 3D geometries, however, some gates, including the T gate, can be implemented without the large overhead caused by state distillation. Although some of these three-dimensional codes can be mapped onto a two-dimensional lattice [JOB16, BC15], they do not reach the same high level of error threshold as the surface code and thus can only be realized with physical error rates that are one magnitude lower.

A trade-off between space and time can be reached by deforming the topological space-time structure of defects (see Section 7.2.5), but also in the production of the magic state ancillae. When created in ancilla-factories in an independent part of the circuit, they can be directly used in the actual computation circuit whenever they are needed (without waiting for them to be created), so the actual T gate can be performed in the time required for the CNOT, measurement and feed-forward. Further time-optimization [Fow12]

proposes that the only relevant time scale is the time needed for the classical feed-forward of non-deterministic gates—one measurement time per T gate—as long as enough additional qubits are introduced (as further explained in Section 7.2.5). Even in time-optimized models, typical times for relevant factoring algorithms are still in the order of days or hours. For this reason, as long as no faster algorithms (in terms of T gate rounds) or faster measurements are available, a qubit-optimized but time-consuming run is not worth consideration.

In general, determining the effort for a given task consists of determining its logical volume—number of qubits and gates and translate it into physical volume.

7.4.3 Other error correction codes

Another possibility for finding trade-offs are using other implementations or code variants. These often have special requirements on the physical system, but therefore come with additional benefits in terms of fault-tolerance or overhead. Hence, for some specific quantum computing platforms, they might still be a potential alternative. A short discussion of several code variants is given in [CTV16].

Multi-level system codes Implementations using multi-level systems, qudits, have been shown to be capable of reaching threshold error rates of more than 8% for high enough qudit dimensions [WAB15]. However, handling of such high-dimension multilevel systems is complicated and in most physical implementations not even possible. A possible platform would be molecules [TdVR02].

Long-range and high-dimensional interactions Error correction codes using interactions of qubits that lie arbitrary far away can have much better threshold and performance than 2D nearest-neighbor codes like the 2D surface code. Often, they attract only few interest because in the leading platforms they are more than hard to implement. However, in distributed implementations using photonic interconnects (or any other flying qubit), like NV centers, quantum dots or trapped ions [MK13, MRR+14], arbitrary qubit interactions are not a big problem.

A rather old proposal of Knill [Kni05] reports a threshold of 3% for a code using non-local interactions and post-selection. Besides a change of threshold, codes using more than two dimensions are also capable of implementing transversal non-Clifford gates [BMD07], which significantly reduces the qubit overhead due to magic state distillation. However, one needs to be careful here: The ability to perform T gates often comes with a reduction of the set of easily realizable Clifford gates. The Hadamard gate for example can not be performed in the three-dimensional color code without an extra logical ancilla.

Tailored codes If the error model of the underlying is known, e.g., if it is dominated by phase errors but has negligible bit flip, the error correction code can be tailored to that mechanism. [TBF18] details a slight modification to the well known surface code that leads to significant improvement in the error threshold for precisely that code. The modification consists of changing the plaquette stabilizers from Z to Y operators and using a specific decoder [BSV14]. The results of this alteration to the surface code are made explicit in the error threshold values of 43.7% for pure dephasing noise, and 28.2% for a bias 10 (bias is the fraction of Z noise over X and Y noise).

While these results are promising, a full fault-tolerant analysis has not been included, i.e., the code behavior is only studied for ideal measurement and gate operations, so the threshold should not be compared to the results used in this study. It is unclear whether the positive results listed above will remain relevant after the code is adapted to provide fault-tolerance, and, even if it will show an improvement, it most probably will not be as substantial as the one in the paper.

Topological cluster states There are physical systems, for which an initial collective interaction of the qubits is easier to implement than specific multi-qubit gates along the computation process. One-way quantum computation using topological cluster states is a possible alternative for these systems, also accounting for qubit loss, for which a fault-tolerant protocol similar to the surface code has been proposed [RHG06, WF14b]. However, this scheme requires a three-dimensional cluster state. Even if implemented on a 2D lattice, this needs more connectivity than the usual square lattice used for the surface

code and usually also fast entangling gates on the fly (which is in contradiction to the cluster state idea). Besides, this approach also does not allow transversal non-Clifford gates.

Color codes The color code is defined on a three-valent, three-colorable lattice with qubits on its vertices and stabilizers as operator chains around plaquettes. The two-dimensional color code can be mapped to a slight variation of the surface code by folding it along its diagonal [KYP15]. It has a higher connectivity than the surface code, but also comes with lower noise threshold and can be harder to implement in experiment.

The gauge color code [BNB16] is a three-dimensional variation of the color code, which uses gauge-fixing to perform non-Clifford gates in constant time, but since the code itself needs more physical qubits the overall space-time overhead can only be reduced by a constant fraction at most.

A rather promising approach is the doubled [BC15] or stacked [JOB16] color code, which can be mapped to a two-dimensional implementation. Computational universality is reached by switching between two encodings: One 2D code that can implement transversal Clifford gates on a 2D lattice, and one 3D code in which the T gate is transversal. Although the three-dimensional part can be mapped to a two-dimensional lattice, it requires some non-local interactions (but in a very limited way with only a small set of qubits involved).

7.5 Operational conclusions

Gate	Volume (building blocks)
H	12
CNOT	12
S	180
T	$120 + v_{\text{dist}}$

Table 7.1: Space-time volume in fundamental surface code building blocks for Clifford and T gates. One building block represents $4(d+d/4)^2$ physical qubits undergoing $d+d/4$ error correction cycles. The count for the S gate is an approximation, including the costs for a two-level distillation with constant distance d . The distillation cost for the T gate v_{dist} must be optimized separately for every problem.

So why is this extensive description of error correction and the surface code key to any evaluation? Because error correction is a central ingredient of operating a quantum computer that largely dominates its resource requirements - error correction is complicated and needs to bridge the gap between expected physical precision and algorithmic requirements, see the overview in Figure 2.1. In contrast to an execution of a quantum algorithm directly on the hardware, in error-corrected quantum computation one can not simply count the number of qubits and time-steps from the algorithm. Different gates can be performed in different ways and under different circumstances: The Hadamard gate for example, does not need any additional qubits if performed on a single logical qubit. However, it occupies all physical qubits around so that the neighboring logical qubits are limited in their action (The occupied space can for example not be used for braiding a CNOT, or for another Hadamard gate). On the other hand, the topological nature of the surface code brings more freedom to deform circuits to make them time optimized, or to avoid qubits being idle for too long. Hence our calculations will not be based directly on logical qubit and gate counts but on fundamental space-time building blocks (called plumbing piece in some works, e.g., [FDJ13]) of $4(d + d/4)^2$ physical qubits (required for one hole in the surface code lattice including isolation from neighboring holes) and $d + d/4$ time steps (to preserve temporal distance after moving holes). In this representation, the logical error rate per building block can be written as $P_b = d(p/p_{\text{th}})^{(d+1)/2}$. The number of building blocks required for a certain gate is shown in Table 7.1. For the sake of clarity, capital letters have been chosen for logical counts and rates, small letters indicate an actual physical parameter.

The application of our evaluation scheme can be staggered as follows

1. We assume that the requirements of lower-level criteria are met: Low-level functionality as well as clear insight on architectural elements such as connectivity and layout. We also assume that performance has been quantified at least for average (two-qubit-) gate error p and readout fidelity.
2. We verify that the architectural requirements of Section 7.3 are met for the error correction code of interest.
3. We choose an algorithm with gate counts from Chapter 9. We record the total numbers of Clifford gates N_{Cl} and T gates (including T^{-1}) N_T , ideally also further detail on the distribution of Clifford gates to CNOT, H and S gates, and the maximal number of non-parallelizable T gates, the T-depth $N_{T,s}$.
4. We calculate the minimal time for execution of the whole computation from the number of sequential T gates $N_{T,s}$:

$$t = N_{T,s} t_M \quad ,$$

where t_M is the time required for physical measurement + classical feed-forward. The number of surface code cycles that can be done in that time is approximated by $N_{SC} = 0.1N_{T,s}$ (see Section 7.2.5). Depending on the platform, this factor can be adjusted (if the ratio of measurement time to gate times differs too much).

5. We calculate the overhead consumed by state distillation, as described in the next list and calculate the corresponding number of physical qubits n_{dist} in a time-optimal realization

$$\eta_{dist} = v_{dist} / N_{SC} \quad .$$

6. From the total number of logical Clifford gates N_{Cl} and logical T gates N_T , we calculate the target logical error rate per building block for the Clifford-part of the circuit (i.e., everything apart from the T gate state distillation). A CNOT or Hadamard gate can be fit into 12 blocks [FD12, FMMC12], the S gate needs a volume of 36 building blocks (for three Clifford operations) [GF17] plus approximately 144 for distillation⁹ and a T gate—consisting of a CNOT and a corrective S gate, conditional on measurement—needs $12 + 36 + 144/2 = 120$ building blocks additional to the state distillation. So, the total number of building blocks is $N_b = 12N_{H,CNOT} + 180N_S + 120N_T$ and we choose a target error rate per block $P_b = 1/N_b$. Note that the space-time cost calculated with these numbers is only a lower bound to the actual cost since deforming the structure or remaining idle qubits might lead to a slight increase. On the other hand, the gate counts from the underlying algorithm are an upper bound since gates might be combined via bridge compression [FD12, PF13] to smaller structures. Nevertheless, the bare space-time cost still gives a good estimate for the total resource costs.
7. We use the formula for the logical error rate per building block $P_b \approx d(p/p_{th})^{(d+1)/2}$ [FDJ13] to calculate the distance required for the target error rate.
8. With d , we can calculate the total space-time overhead for the Clifford-part

$$v_{Cl} = 4(d + d/4)^3 N_b \quad .$$

9. Having the number of surface code rounds N_{SC} fixed in 4, we calculate the number of physical qubits for the Clifford part

$$\eta_{Cl} = v_{Cl} / N_{SC} \quad .$$

10. We get the total number of qubits

$$\eta = \eta_{Cl} + \eta_{dist} \quad .$$

⁹ Since the S gate state distillation is much less resource demanding than the T gate distillation, we will not aim to optimize the distillation process but just approximate the volume by the costs of two successful full-distance distillation levels with a volume of 18 per single round [FD12]. For improved physical error rates, one level will be sufficient, leading to an overall volume of 18.

The optimal overhead for the state distillation circuit can not be calculated directly, simulations are necessary to determine the best values and codes. We explain the general approach in the following, however, we will use exemplary values of $\epsilon = 1$ and only Reed-Muller distillation for simplicity and note that improvements from optimizations will only give rise to constant factor improvements [OC17, FDJ13]. Furthermore, we will omit the overhead due to corrective S gate distillation, since it is negligible compared to the T gate (already ten times smaller for only one round of distillation).

1. We record the number of total T gates N_T and set the target error rate per magic state to $P_T = 1/N_T$
2. We choose a factor ϵ between the ideal distillation output error P_{dist} and the error due to the surface code distance so that the overall output error of one distillation round is $P_{\text{out}} = (1 + \epsilon)P_{\text{dist}}$. We additionally choose the code we want to use for distillation. Usually, Bravyi-Haah block-code distillation makes only sense for small error improvements, for example if one round of Reed-Muller is slightly not enough and only a small amount of additional error reduction is needed. We calculate the required distance for the top round of state distillation from $V d(p/p_{\text{th}})^{(d+1)/2} < \epsilon P_T / (1 + \epsilon)$ where V is the space-time cost of the distillation protocol in units of fundamental building blocks. The Volume for the Reed-Muller code is $V_{\text{RM}} = 192$ and for Bravyi-Haah block-code $V_{\text{BH}}(k) = 96k + 216$ [FDJ13].
3. We determine the required input error rate for this round from $P_{\text{dist}}(p_{\text{in}}) = P_{\text{out}} / (1 + \epsilon)$. We get $p_{\text{in}} = (P_{\text{out}} / (35(1 + \epsilon)))^{1/3}$ for the Reed-Muller code and $p_{\text{in}}(k) = (P_{\text{out}} / ((3k + 1)(1 + \epsilon)))^{1/2}$ for Bravyi-Haah.
4. To compensate for unsuccessful distillation rounds, we choose to produce $1/P_{\text{succ}}(p_{\text{in}})$ times as much magic states as originally needed, so the actual volumes needed per successful output state for Reed-Muller and Bravyi-Haah are $V_{\text{RM}} = 192 / (1 - 15p_{\text{in}})$ and $V_{\text{BH}} = (96k + 216) / (1 - (3k + 8)p_{\text{in}})$, respectively. Since the actual size of the blocks depends on the distance, we calculate the real volume in terms of physical qubits times surface code cycles directly: $v = 4(d + d/4)^3 V$
5. If the initial error rate of the bare injected magic state $0.4p_2$ is below the required input error rate p_{in} , one distillation round suffices, if not: Take the required input error rate as new target error rate for a second, preceding distillation round with a new choice of distance and code (and k). Iterate steps 2–5 until error rates match.
6. We calculate the total qubits-time volume of the whole distillation by summing up the required distillations of all rounds i : $v_{\text{dist}} = \sum_i N_{\text{RM(BH),i}} v_{\text{RM(BH),i}}$. For the Reed-Muller code we use the formula counting physical qubits times surface code cycles.

$$v_{\text{dist}} = \sum_i 4 (d_i + d_i/4)^3 15^{(i-1)} N_T V_{\text{RM}} / (1 - 15 p_{\text{in},i}) .$$

As discussed earlier, a surface code cycle consists of the gates necessary to map an error syndrome to a syndrome qubit as well as syndrome readout and reset. It is practically in most cases dominated by the readout time.

7.6 Groups

Austin Fowler, Google Google Staff Researcher Dr. Austin Fowler (also <http://www.topqec.com.au/about>) has acted as a trailblazer for practical fault tolerance. For a set of error correcting code he has worked out all details of error correction and computing, including the appropriate minimum matching software *Autotune*. He has collaborated with **Thaddaeus Ladd** of HRL Laboratories in applying these ideas to semiconductor platforms. Furthermore, he is working on decreasing the resource costs for implementing both Clifford and non-Clifford gate sequences on the surface code.

Robert Raussendorf, UBC Vancouver Besides his contribution to measurement-based quantum computation, Robert Raussendorf also studies fault-tolerance of low-dimensional systems. Together with **Jim Harrington** (HRL), he first proposed the 2D surface code, and earlier also a fault-tolerant one-way quantum computer. Much of his work has recently developed towards more general quantum information theory rather than practical error correction.

Dorit Aharonov, Hebrew University Dorit Aharonov is a distinguished quantum computer scientist at the Hebrew University in Jerusalem. She has proven a very rigorous version of the threshold theorem.

John Preskill, Caltech John Preskill has a distinguished background in cosmology and particle physics. In recent years, as director of the Institute for Quantum Information at Caltech, he changed to studies of the mathematical aspects of quantum information theory and new schemes for error correction. Together with Daniel Gottesman, he presented proofs of the threshold theorem for different code variants [AL06, AGP08].

Daniel Gottesman, Waterloo Daniel Gottesman did his PhD under supervision of John Preskill at Caltech. At the Perimeter Institute in Waterloo, he currently works on quantum error correction, fault-tolerant quantum computation but also quantum cryptography and complexity. He is known for developing the stabilizer formalism which is able to describe a lot of quantum codes, including the surface code.

Andrew Steane, Oxford An experimentalist in ion traps, Steane has co-developed the most common concatenated code scheme, in particular the 7-qubit version of a code that completely corrects single-qubit errors.

Peter Shor, MIT A leader and pioneer in quantum algorithms, has also made early contributions to quantum error correction.

Alexei Kitaev, Caltech Alexei Kitaev works together with John Preskill at Caltech. He developed the concept of topological quantum computation and is currently interested in the mathematical structures of topological phases. Although his work is more on the conceptual level, he still made several contributions to fault-tolerant quantum computation.

Barbara Terhal, Delft / Sergey Bravyi, IBM Barbara Terhal is professor and moved from the Quantum Fault-Tolerance and Error Correction group at RWTH Aachen to TU Delft in the Netherlands. She has made contributions in many different areas of quantum information and published a comprehensive review article on quantum error correction for quantum memories. Besides, the group is also interested in 4D and concatenated codes and their implementation. Much of this has been achieved in collaboration with Sergey Bravyi, who supports the IBM experimental effort with error correction know-how.

Simon Devitt, Macquarie University, Sidney A PhD-colleague of Fowler, he has adapted error correction and fault tolerance schemes to atomic and optical architectures similar to what Fowler has done for solid state.

Microsoft QUARC group, Redmond While Microsoft's main bet is on topological qubits, they are active in error correction. They do provide error correction support for Leo DiCarlo's IARPA-funded Josephson qubit project.

Earl T. Campbell, Sheffield Campbell leads a small research group working on fault-tolerant quantum computing. His current work focuses on magic state distillation and possible ways of resource and overhead reduction.

James Wootton, Basel James Wootton, postdoc in the Quantum Computing Group of the University of Basel, works on quantum error correction with a focus on topological quantum computing. He has presented several different schemes for error correction and decoding.

7.7 Glossary for error correction

Coherent error – Error represented by a quantum-coherent operation; here used synonymously with unitary errors, see section 6.5, with a caveat in chapter 7.4.1.3

Clifford gate / Clifford gate - Normalizer of the Pauli group, chapter 6.4.3

Concatenated code - error correction code consisting of connected layers of error correction codes, see chapter 7.1.3

Depolarizing channel - decoherence model describing the symmetric randomization of a qubit state, see chapter 6.4.3

Distance - size measure for error correcting code related to the number of correctable errors, see chapter 7.2.3.1

Error correction cycle - sequence of initialization of syndrome qubits, mapping of error information onto the syndrome qubit by quantum gates, syndrome readout, processing of errors and corrective operations, see chapters 7.1.4 and 7.2.1

Error rate - probability of an error per operation (see e.g. chapter 6.4.3)

Error syndrome - bit value containing information about the location and nature of errors, see chapters 7.1.1 and 7.2.2

Gottesman - Knill theorem - theorem stating that a Clifford-only quantum computer can be simulated efficiently on a classical computer see chapter 6.4.3

Logical qubits - an error corrected qubit that is used in an algorithm, chapters 7.1.1 and 7.2.3

Magic state distillation - leading procedure to implement the T gate, see 7.2.4.3

Pauli Error - an error described by phase and bit flips and combinations thereof

Physical qubits - the physical devices whose errors create the need for error correction, see chapter 7.1.1

$\pi/8$ gate - see T-gate

Stabilizer (of a state) - a set of commuting operators to which the state is an eigenstate with eigenvalue 1 and that uniquely determine the state, see chapter 7.1.1

Stochastic errors - errors, that occur by applying an operation with a given classical probability, see chapter 7.4

Surface code - currently leading code for quantum error correction, see chapter 7.2

Surface code cycle - Error correction cycle of the surface code (see there)

Syndrome qubit - a qubit containing information about an error syndrome (see there)

Syndrome measurement cycle - Error correction cycle (see there) without the last two steps

Systematic error - errors that occur with certainty (but can be small by another measure), see chapter 22

Threshold - numerical value of a physical error below which error correction is effective in reducing the logical error, see chapter 7.1.3

T-gate - a single-qubits non-Clifford gate used to go beyond the limits of the Gottesman-Knill theorem, see chapters 7.1.4 and 7.2.4.3

Unitary error – an error described by a unitary operation, see chapters 6.5 and 7.4.1.3

8 Benchmarking and fault-tolerance on non-standard architectures

Some implementation platforms are not well suited for application of the surface code or other standard error correction models. This can either be because they are not based on the gate model (as in quantum annealing) or because the resource inventory is vastly different from that of most platforms, such as in cluster-state quantum computing. The benchmarking scheme for them deserves separate evaluation.

8.1 Quantum annealing

As described from a hardware perspective in 13.1.4, Quantum annealing/adiabatic quantum computing is based on slow global control of the qubits rather than on delicate and fast local control. Quantum annealing can efficiently simulate gate-based quantum computing if many-body interactions are available (n -local with $n \geq 3$ can be implemented) which so far has only been proposed but not implemented, and it is not known how to benchmark it. On the other hand, quantum annealing for optimization problems (for which no quantum circuit is known and likely does not exist) has been implemented by Canadian company D-Wave Systems, using rather incoherent qubits and 2-local couplers. So while this platform *fails* a fundamental resource requirement, it still gives an outlook on how to evaluate quantum annealing.

8.1.1 Coherence and control

As quantum annealing strives to use the lowest energy eigenstate of the system, relaxation due to contact with a cold heat bath, i.e., a directed T_1 process can in principle assist the annealing process. Details about ideal values are not currently known. Also, it is known that strong decoherence will suppress any quantum properties, which is why for stronger coupled systems, shorter annealing times are often advantageous. However, the annealing time might still be orders of magnitude higher than the coherence time of the qubits without leading to failure. Realistic devices are also limited by nonuniformity of the qubits and fabrication defects.

8.1.2 Benchmarking quantum annealing

Annealing does not rely on the application of accurate quantum gates, measurements during the computation or exact initializations (in the initial Hamiltonian, the ground state will be the easiest to reach, and if not reached, relaxation can always help), so the typical error rates known from circuit-based quantum computers do not play a big role anymore. It is not even important to end up in the final state with probability close to 100%, since the computation can be repeated and the right result found by comparing energies. Even probabilities smaller than 50% are no problem, as long as the computation is performed enough times.

The only really important criterion for a quantum annealer is thus the time until the ground state (which is the solution of the computation task) is found. This time is dominated by two variables: The running time of one repetition and the number of repetitions. Running the annealer faster results in a lower probability for ending up in the ground state, and thus needs a higher number of repetitions. There exists an optimal balance between both, leading to the lowest overall computation time.

A common measure reflecting this trade-off is the time-to-solution (TSS) metric, as explained in [AL17] and [RWJ⁺14] (Supplementary Materials): It gives the overall time until the ground state is found at least in one of the repetitions with probability p (usually 99%). It is calculated as $TTS(t_j) = t_j R(t_j) \alpha$ where t_j is the time for one repetition, $R(t_j) = \ln(1-p) / \ln(1-p_s(t_j))$ is the corresponding required number of repetitions with a per-run success probability of $p_s(t_j)$ and α is the number of parallel runs that can be performed by devices

providing more qubits than required. In some cases it might be necessary to include initialization and readout times or other time costs that might occur when running the annealer multiple times in series. For current architectures as for example D-Wave these costs are much smaller than the running time and can be neglected. The performance of a quantum annealer is usually compared to classical algorithms by considering a particular quantile q regarding a set of problems, for example the median of $TTS(t_f)$ for a set of different problem instances and searching for the optimal run time t_q^* , minimizing this quantile. The optimized quantile is denoted by $\langle TTS(t_q^*) \rangle_q$. High quantiles are usually more informative in terms of scaling, since they include also the hardest problem instances [RWJ⁺14].

8.1.2.1 Quantum speedup

Quantum speedup is defined to compare quantum devices to classical devices solving the same problem and to find if a quantum computer can beat the performance of a classical algorithm. Especially in quantum annealing this is an important question, since annealers often have the same scaling as classical algorithms, but with different prefactors. Quantum speedup in general is defined as the ratio of the (overall) run time of a quantum annealer $Q(N)$ to the run time of a specific classical algorithm $C(N)$ in the limit of large problem sizes N : $S(N) = C(N)/Q(N), N \rightarrow \infty$ [RWJ⁺14]. The problem with this approach is the definition of the classical algorithm. It is not always known if a certain algorithm is optimal, so one can only compare the quantum device with the best available classical algorithm.

Another way of comparison is to define *limited quantum speedup*. It compares a quantum computer with a classical computer following the same algorithmic approach. For a quantum annealer, the corresponding classical algorithm is for example simulated annealing or simulated quantum annealing: Algorithms that run on classical hardware using Monte Carlo simulations. Limited quantum speedup does not prove that a quantum computer is an improvement to classical computers, but shows that quantum effects actually show up and help improve the computation process in a quantum device.

Apart from the choice of classical algorithm it also makes a big difference when the quantum algorithm is not optimized, for example when the optimal per-round run time is shorter than the smallest available time of an annealing device (this time is given by the underlying hardware of the annealing device—even if a solution could be found faster, the annealer will run at least its minimum run time). Then it becomes impossible to determine the optimal overall annealing time $\langle TTS(t_q^*) \rangle_q$ which is crucial for making estimations for big N . The run time needs to be optimized for each N . Fixed values for t_f might mislead the conclusion, as for small N the total computation time scales only slow with the problem size until N gets too big for the chosen t_f . Then, the slope increases to higher than optimal. Figure 8.1 illustrates this behavior.




Figure 8.1: Sketch of total time until the ground state is found with desired probability as a function of the problem size. The dotted lines show the performance for several fixed values of per-round run time t_f . The blue line shows the optimal result, reached if the run times t_f were optimized individually for each problem size. When measured with a fixed t_f (for example because of limitations of the annealing device), the slope of the measured curve (red) might indicate a wrong behavior: For small N , the slope is lower than optimal (possibly faking speedup where there is none), for large N , the slope is higher than optimal (which might mask possibly existing speedup) [Ami15]. "Reprinted figure with permission from [M. H. Amin. *Phys. Rev. A*, 92(5):052323, 2015.] Copyright (2015) by the American Physical Society."

Although measurements with fixed (too high) t_f do not represent the large- N behavior of an annealing device, they can still be used to eliminate the possibility for quantum speedup: The slope of the measured curve gives a lower bound to the optimal slope, therefore it is sufficient to show that the measured slope is higher than that of a classical algorithm. Conclusions in the other direction are not possible.

The concept of TTS and quantum speedup can easily be applied to annealers using many-body interactions. However, no such annealers have been built yet and thus no such study has been done.

8.1.2.2 Typical causes for misinterpretation and overestimation

There are many different ways of presenting speedup comparisons depending on what the authors want to tell. This section shall give an overview of common situations that might tempt the reader to overrate the performance of an annealer.

Fixed run times As already mentioned above, although the TTS values for fixed run times lie above optimal, the curves have a smaller slope than the optimal curve for small problem sizes (so, for the problem sizes that can be tested). Hence, if extrapolated to bigger values of N , it wrongly indicates better scaling than the actual optimal curve, as can be seen in Figure 8.1.

Crafted problems There exist problem instances that are far more suitable for quantum annealing than others. Especially potential landscapes with thin but high barriers are relatively easy for quantum devices. In contrast to thermal hopping (crossing potential barriers classically using thermal energy), for which the

probability scales exponentially with the height of the barrier, quantum tunneling depends on the size of the tunneling domain, i.e., its probability scales not only with the height but also with the width of the barrier. Across thin barriers, quantum tunneling is more likely to occur and thereby helps the annealing process to find the ground state faster. However, not all problems can be implemented with such Hamiltonians. So, although comparisons for this class of problems show stunning results [DBI⁺16], they do not prove that speedup can be observed for other problem instances, especially not that universal adiabatic quantum computing is possible faster than with classical computers. Furthermore, in many problems tunneling is fast to bring the system in an approximate ground state, but if the task is to find the global energy minimum (which is often separated by a broader energy barrier [AKR10]), also the presence of tunneling can not bring a significant time reduction.

Low quantiles An other issue is the choice of the right quantile. Even if the performance of an annealer is determined by implementing a broad set of problem instances, one can not show all of the results. Giving a several quantile of the distribution is a good thing, but can also mislead. Especially low quantiles give information of only the easy problem instances, but a quantum computer should be able to solve all sorts of problems. So, in order to get relevant information on the scaling of an annealer, one should consider the scaling of the highest quantiles, which also include harder problem instances [RWJ⁺14].

Omitting efficient classical algorithms In the ideal case, the performance of a quantum annealer should be compared to that of the best known classical algorithm solving the same problem. However, sometimes speedup is only detected in comparison with a certain algorithm, but not with all. One example is the definition of limited quantum speedup: It gives important information on the quantum properties of a device, but not on its computational value. Usually it is clearly stated what kind of speedup is considered, nonetheless one should always be careful here. Furthermore, it is known that some of the algorithms that outperform annealers for current architectures will soon get ineffective as the devices are improved in terms of connectivity.

8.1.2.3 Further evaluation criteria

Besides the benchmarking available for current architectures, which basically only focus on the time until the solution is found, future generations of annealers should also be evaluated in terms of their connectivity and control possibilities. Although three-local interactions are said to be sufficient for universal quantum computing with annealers, higher-weight interactions are favorable in order to perform efficient error correction and make the computation scalable. Furthermore, architectures providing more connections between the qubits, for example on a three-dimensional lattice can also overcome the limitation of two-locality since three-local interactions can be mapped to such implementations [LHZ15] with only two-body interactions.

8.1.3 Quantum annealing for the shortest vector problem

A promising modern cryptosystem, Lattice Based Cryptography (LBC), is currently viewed as quantum-safe. In LBC the security relies on hardness of the Shortest Vector Problem (SVP) in both exact and approximate form. In [JGLM19] an analysis of the potential of Adiabatic Quantum Computation (AQC) for attacks on LBC is given. There is no proof of speedup yet a noteworthy numerical evidence in the sub-adiabatic regime.

SVP is conjectured to be hard at this moment but there is no proof that quantum computers cannot solve them in polynomial time. Despite the fact that the time complexity of AQC algorithms is in general hard to estimate, AQC is a valid candidate for the attack on LBC for two reasons : (1) LBC can be formulated as an optimization problem, and (2) while AQC in general has a prohibitive time cost of achieving adiabaticity, for approximate SVP up to a threshold, approximate solutions are also admissible.

8.1.3.1 SVP attack on LBC

Attacks on SVP are based on taking the underlying lattice and finding the shortest nonzero vector in that lattice expressed in a given basis or at least finding a vector of given length beyond that shortest vector. The paper [JGLM19] proposes an embedding into an adiabatic quantum computer achieving that result.

The embedding proceeds in multiple steps. It utilizes the Bose Hubbard model. This is a model of quantum particles embedded in a lattice that can move in the lattice and that can repel each other both on the same lattice site as well as across the lattice. Tailoring these interactions defines the lattice for the lattice-based algorithms and minimizing the interaction energy corresponds to solving the SVP problem. Quantum annealing is proposed to solve this model. The quantum tunneling term that is used to initialize the state is in this case the kinetic energy of particles hopping in the lattice. It is adiabatically switched off in order to settle the particles in a state that minimizes the interaction in order to solve the problem.

This being based on the Bose (not the Fermi) Hubbard model means that more than two states are allowed per lattice site, i.e., more than one qubit. This overhead does not change the observation that the embedding is efficient in the number of qubits. As an important technicality, rather than the ground state (which is the zero vector), one is looking for the first excited state. This complication is elegantly circumvented by using a separate state to represent the zero vector.

8.1.3.2 Potential for speed-up

As described above, a clear proof of quantum speed-up would require to extract the energy gap across the anneal, which is not achieved in Ref. [JGLM19]. Rather, they rely on numerical simulations on classical computers covering small examples (dimensions 2,3, and 4). These experiments confirm the value of adiabaticity, i.e., that if sweeps get slower, the distribution of output values clearly clusters at low energies. This is not made quantitative into a scaling analysis with, proven or extrapolated, speedup. It is pointed on new discussions of continuous but not fully adiabatic algorithms with no clear conclusion given.

8.1.3.3 Analysis

The lack of a speedup analysis but the expression of hope is rather common in adiabatic quantum computing, see above. The faster sweeps that are proposed are an interesting metaheuristic the potential of which is not fully understood. Given the hardness of gap extraction, in order to evaluate these heuristics in interesting size, benchmarking on actual hardware would be the most important way forward and should be closely monitored. There is some indication that fast sweeps (that can be repeated more often) reduce the time to solution [CFLLS14] but those have been done for generic cases and not for this specific model.

Also, the analogy to the Bose Hubbard model is noteworthy. This model can be directly simulated in the sense of analogue, single-purpose quantum simulation, specifically cold atoms in optical lattices, see chapter 16. Design of appropriate programmable interactions as it is, e.g., done in the EU Flagship project PASQuaS [PAS18] would allow to scale rather quickly and is actively pursued and certification of quantum supremacy in these systems is an active field of research [EHWR¹⁹, HKEG19]. The Bose-Hubbard model can also be studied in superconducting circuits [FZ01, LH10], combining the ease of programming and design of these systems (boosted by the tools developed around the quantum supremacy experiment) with a compact native application of this model. An aggressive scaling project of those simulators is not known.

8.1.4 Fault tolerance for quantum annealing

Up to now, there is no scheme known to provide arbitrary fault-tolerance to a quantum annealer, especially, there is no evidence for a threshold. However, there exist several approaches towards error-suppression and

some simple error detection and correction ideas. Most protocols aim to increasing the minimal energy gap between the ground state and the first excited state and suppressing coupling to the environment, leading to a higher probability for ending up in the ground state. This energy gap scales as an inverse polynomial in the problem size, so without protection of the gap height, the annealing time would raise polynomially with the problem size, too. Although necessary, these methods only suppress errors rather than actively correcting them. Using simple repetition codes, i.e., encoding the qubits in multiple copies and introducing majority-votes can also provide some ability to explicitly check and correct low-weight errors (i.e., error with only one or very few qubits involved).

None of these methods is a satisfying solution for error correction, since they either do not provide enough error correction ability, or are not scalable, or need many-body interaction and controls that are not feasible with current hardware. This section will discuss various approaches, to analyze where they fail and which technological developments might bring them to relevance again. A good overview on the topic can be found in [YSBK13].

8.1.4.1 Error suppression

Energy gap protection The energy gap protection protocol, as realized in [PAL14] relies on quantum stabilizer code. By introducing extra qubits, the original qubits can be mapped to logical qubits consisting of several physical qubits. Usually a simple repetition code is used: Operators are replaced by the sum of equal operators acting on multiple qubits and an extra term in the Hamiltonian, with some extra ancillary qubits is introduced which gives an energy penalty to single qubit flips out of the code space. A code using n times as much qubits as for the original problem can penalize up to $\lfloor n/2 \rfloor$ qubits. The encoding itself already increases the energy scale and thereby also the ground state energy gap by a factor of n . The penalty term additionally lowers the probability of undesired excitations out of the code space. Errors that commute with the penalty term (usually these are phase-flips) are not suppressed. There are stabilizer codes that can correct all errors, however, these codes require high-weight terms in the Hamiltonian which are experimentally challenging to implement in this scheme. In principle, it is also possible to manually correct errors by measurement of the stabilizer operators and applying corrective gates, however this is not a technique that is usually available in quantum annealing devices.

Further progress with energy gap protection schemes has been made for minor embedding [VAPS⁺15]. By introducing penalties that vary with each qubit, corresponding to the respective problem Hamiltonian, the performance could be significantly improved. In the same work, a scalable square code is introduced, which makes concatenated encoding and thereby high error-tolerance (at the cost of an increasing number of qubits) possible.

Dynamical decoupling In a rotating frame, energy gap protection can also be viewed as modulating the term of the Hamiltonian responsible for coupling to the environment by a fast (depending on the penalty energy) oscillating term to make sure it cancels out for sufficient time scales. Dynamical decoupling takes the direct path to this oscillation, applying a sequence of stabilizer control pulses in time [QL12]. This technique is a well-known method for suppressing errors due to any spurious terms in the system Hamiltonian. However, it does not create an energy difference between code and non-code space, so the code space is not energetically preferred. A big advantage of dynamical decoupling is that also high-weight Hamiltonians can be implemented using many different single-qubit Hamiltonians, as long as the pulses are significant stronger than than the encoded adiabatic Hamiltonian.

Since there are codes that use many two-body interactions and only very few high-weight operators, a combination of dynamical decoupling for the high-weight terms and energy gap protection for all other control might work in some cases. A problem that both energy gap protection and dynamical decoupling struggle with is that they can only rescale the system-bath coupling, a complete suppression would only be possible with control Hamiltonians of infinite energy, or pulses of infinitely high frequency. As the problem size and annealing time increase, the inevitably increasing demands for fault-tolerance will require at least increasing control energies which at some point will not be possible to fulfill anymore.

8.1.4.2 Error correction

The easiest way to include error correction to error suppression schemes is to measure the stabilizer operators at the end of computation together with the qubit information and if necessary, to correct the outcome classical. However, besides the fact that errors amplify during long computations, errors happening during the annealing process might evolve differently in the adiabatic sweeping and become uncorrectable quite soon. This is because the annealing Hamiltonian acts differently on the different subspaces of the code. Correction during the process would require fast quantum gates and measurements, not only in the end. This is not compatible with the general idea of annealing.

Protected Hamiltonians It is possible to create Hamiltonians that act similar in the non-code space as in the code space, for example if every error maps the ground state of the system Hamiltonian to another eigenstate. This way, it is sufficient to measure only in the end and track the errors back. Again, this is not possible without high-weight terms in the Hamiltonian. A distance d code must have at least d -local interactions. Furthermore, the projected operators must be a sum of an extremely high $O(2^d)$ number of Pauli operators. It might be possible to lower this number by factoring terms, however this is still an open problem.

Local cooling Additional resistance against local excitations can be reached by coupling each physical qubit to a low-temperature bath that pulls entropy out of the system. The coupling should be implemented in a way such that the bath can absorb the energy penalty of an unwanted excitation of a qubit out of the code space and put it back to the original state. This only works for local excitations, if one wants to correct higher-weight errors one needs very special Hamiltonian structures that common stabilizer codes do not have, or one needs high-weight Hamiltonians or high-dimensional interactions. However, local cooling can still help to protect the code space, at least in some way, which, in combination with other error correction techniques might be of some effort. It is possible that novel cooling mechanisms can act on multiple qubits to enact higher-weight protection.

It turns out that the key point for all error corrected quantum annealing to work is the ability to either implement high-weight Hamiltonians or to include circuit-model techniques like fast measurements and gates into the computations. The latter raises the question if there is an advantage of annealing over circuit-based quantum computing at all. High-weight Hamiltonians might sound like a problem that can be easily solved since high-weight unitary gates are rather easy to implement in circuit-based quantum computation by performing many low-weight gates in parallel. However, there is no comparable way known for Hamiltonians. One approach is using perturbative gadgets [JF08] which create high-weight operators using only weight-two terms. Besides introducing additional qubits, this requires coupling strengths to scale exponentially with the desired weight.

8.2 One-way quantum computing

One-way, or measurement-based quantum computing [RBB03, RB01] is a third approach to universal quantum computers using no gates during the computation, but—other than annealing—measurements. However, multi-qubit gates are still required in the preparation of the cluster state that is used as a resource for the computation: In a first step, all qubits are initialized, each in the state $|+\rangle$. Subsequent, CZ gates are applied to pairs of neighboring qubits on a usually two-dimensional lattice. In photonic systems with flying qubits, also higher dimensions are realizable. Note that since this is still a preparation step, the CZ gate may also be performed non-deterministic. Once such a cluster state is created, logic gates are implemented by applying measurements in combination with classical feed-forward. Clifford gates do not need feed-forward as the corrective Paulis can also be accounted for in the end of the calculation. Therefore all measurements that represent Clifford gates can be performed in a single step at the beginning of the calculation. If we are not restricted to a two-dimensional lattice structure, the state resulting from this first measurement round can also be created directly as a graph state (a state with entanglement connections between arbitrary pairs of qubits). The solution of the encoded problem is found by measuring a certain set of qubits in the very end of the computation, until all other qubits have been measured. Each of these steps

is susceptible to errors. Hence, although one-way quantum computing is strictly measurement-based during computation, also initialization and gate errors, and not to forget storage errors play a role. On a lowest lever, initialization and measurement (in an arbitrary basis, so this might include some rotations depending on architecture) accuracies can be assessed in the same way as for circuit-based quantum computing. For gate errors¹⁰ it is sufficient to know the fidelity of the entangling process, since no other explicit gates are applied.

8.2.1 Benchmarking one-way quantum computers

It is important to note that these physical error rates do not directly represent the logical errors in this scheme, since gates are only used in the initialization process and only measurements are used to perform logical gates. Hence, it is desirable to also have a fidelity measure for this logical construct. This can be done by RB [ATB16], in analogy to circuit-based QC as described in Chapter 6.4. The protocol can be even simplified using the intrinsic randomness of the measurement processes¹¹. Leakage errors (for example photon loss) can not be characterized with the standard benchmarking protocol. Thus, it is important to take this error source into account if it is present. Benchmarking schemes that include leakage errors exist, but have not yet been adapted to one-way quantum computing.

8.2.2 Error correction in one-way quantum computing

Possible error sources in one-way quantum computation lie in the preparation of resource states and in the measurements. Imperfections in both processes can be modeled by single-qubit depolarizing noise, i.e., Pauli errors acting on single qubits independently. For faulty Bell measurement (or CZ gates in preparation), both qubits can suffer from Pauli errors, before and after the measurement. Furthermore, a Pauli error can occur while storage and before single-qubit measurements. An additional error source occurring often in photonic implementations is photon loss. This cannot be modeled by Pauli channels.

It can be intuitively understood that there exists a threshold for one-way quantum computation from the fact that it can efficiently simulate any circuit-based computation, including an error-corrected circuit. With the use of a hybrid scheme, as introduced in [ZBD14], connecting small algorithm-specific resource states—in the form of graph states—in a circuit-based manner, a threshold of 13.6% local depolarizing noise can be found for Clifford-only circuits using Shor-type codes [SS07]. With magic state distillation, this can be expanded to universal computation without decreasing the threshold. Clifford-error-correction can be done with moderate overhead: The encoding of a distance d code only requires a $d+1$ qubit resource state. These resource states might have a complicated underlying graph structure, but any graph state can be either created directly by applying the appropriate CZ gates, or by initially using a (larger) 2D cluster state and performing a round of Pauli measurements to transform the graph. Each elementary building block can already contain a fault-tolerant encoding in the graph structure (error correction works with Clifford-only gates and thus can be fully implemented in the graph without need for additional gates). The blocks are combined sequentially using bell measurements, which at the same time act as syndrome measurements. It is sufficient to create every block right before it is needed. This reduces storage time and also the number of required qubits, since qubits that have been measured can be reused in the next block.

10 Here, gate error means the error of the CZ gates during preparation and not the logical gates that are accomplished by measurements.

11 Each measurement can lead to two different logical gates, depending on the measurement outcome. Usually this is corrected by adjusting the bases of successive measurements or, in the case of Clifford gates by simply calculating the consequence for the final measurement outcome of the computation. However, this can also be used as an additional source for randomness. As it turns out, fixed measurement patterns are sufficient to reach effective randomized gate sequence. Using this fact, it is even possible to characterize logical non-Clifford gates, as long as the measurement outcomes are all equally probable.

Depending on the CZ fidelity, creation of the resource states might require additional entanglement purification, for which protocols exist [KMBD06, GKV06]. Even for the modular approach, probabilistic entanglement creation is sufficient (although very resource-consuming).

As an alternative to magic state distillation, non-Clifford gates can also be performed by switching to another encoding and using transversal single-qubit rotations. Switching means connecting a block of standard encoding to a block with different encoding, for example the 15 qubit CSS code (Reed-Muller code, see Section 7.2.4.3), in which T gates can be implemented transversally, i.e., by single-qubit rotations. The rotations can be in principle also performed by measurements, but they are usually very easy to perform directly, so there is no need for a measurement-based implementation. The next building block can then be in the original encoding again, if no other non-Clifford gate follows. The threshold for this method is with 0.64% much lower than for the rest of the code. However, single-qubit gates can usually be performed with very high fidelity; in these cases universal fault-tolerant computation is still possible with per-qubit error rates of $\sim 1\%$.

The magic state distillation approach allows higher error rates, but for the cost of qubit overhead for the distillation. As long as the magic states have a fidelity of at least 0.8, they can be purified with Clifford operations and Pauli measurements. Thus, with acceptable depolarizing error probabilities of 13.6% and magic state error probabilities of 20%, the threshold of 13.6% is still valid.

A threshold and error-correction for photon loss exists [DHN06b, TB05], but has not been evaluated for the scheme explained above. There are also other protocols for combining smaller graphs, which do not rely on Bell measurement but on using parallel fusion (i.e., compensating the probabilistic character of entangling operations by making several attempts in parallel) [DHN06b, DHN06a].

8.2.3 Resource calculations

The overhead calculation for error-corrected one-way quantum computation strongly depends on the physical platform and its possibilities. In atom setups for example it is easy to create a big 2D cluster state in a single constant time step. Hence, the bottom-up approach to start with a 2D cluster and create the desired graph state by a round of Pauli measurements suggests itself. On the other hand, flying-qubit implementations like in optics benefit from the freedom to entangle whatever pair of qubit is desired for a certain graph. Here, it would take long times to implement a 2D cluster which will inevitably contain much more qubits and connections.

For each setup, the corresponding method for the creation of graph states needs to be considered, and in a next step, the total resources for running the hybrid code consisting of multiple such graphs can be calculated.

2D cluster states

(suited for next-neighbor-restricted qubits): The number of qubits required for a computation depends on the number of all logical gates applied—including Cliffords—and also on the structure of the computation. We can compare the physical size of the cluster to the size of a picture showing the (logical) circuit diagram that shall be implemented on the cluster. One dimension of the lattice corresponds to consecutive gates: Every logical gate needs a certain number of qubits to be measured along a chain of qubits. The second dimension represents the number of logical qubits operated in parallel: Every chain encodes one logical qubit, multiple chains are arranged in parallel (and multiqubit gates are performed by connecting chains). The chain structure can be achieved by either leaving some qubit positions of the cluster empty, or by disentangling them from the rest of the qubits via measurement. The time for the creation of the cluster state can be constant, independent of its size, in some implementations. After performing all Clifford gates to the cluster (can be done in parallel—feed-forward is only needed for non-Cliffords), it has the structure of a graph state. Now all other operations can be performed, with each non-Clifford gate (apart from parallel ones) requiring one measurement + classical processing step.

Algorithm-specific graph states

(suited for flying/distributed qubits): If any graph structure can be created directly by entangling the right qubits, a Clifford-only circuit of n input and m output qubits and arbitrary length or complexity can be simulated with $n+m$ qubits. Each non-Clifford gate needs at least one extra qubit. The creation time for the graph again highly depends on the underlying physical platform, for non-deterministic CZ gates as often used in optics, it scales exponentially with the number of qubits. The graph is designed such that only non-Clifford gates need to be performed, each taking one measurement step. All other gates are encoded in the graph structure.

When using a hybrid (module-based) error correction scheme, the single blocks of the computation are used one after each other and also in parallel. If qubits are to be reused in later blocks, the computation time for performing calculations on a block and the coupling Bell measurements for each block in a (temporal) row add up. The required number of qubits will be higher than (but in the same order as) the maximal number of qubits in parallel operated blocks, since at the same time, the next blocks already need to be created.

8.2.3.1 What is often not said

The whole graph-state computation itself seems like a very resource-efficient method taken into account that the calculation can be executed with only measurements. However, many steps need extra resources due to experimental limitations, especially whenever it comes to multi-qubit operations, which can not be avoided completely (at least in the beginning, and also in the merging of blocks):

Depending on the fidelity of the CZ gates, extra qubits and time for entanglement purification needs to be taken into account to meet the threshold. For a non-deterministic creation, the success probability scales exponential with the number of involved qubits, so a high number of attempts to create the desired state is necessary. Furthermore, the Bell measurement might require significant effort (especially in linear optics [Gri11]), since it can not always be directly implemented. It can for example be performed by a CNOT operation followed by X and Z measurements on the two qubits, respectively. This again requires on-demand deterministic two-qubit gates, which we originally wanted to avoid with the one-way approach. If the complete graph is to be created at the beginning without adding blocks during calculation, this will result in a huge space overhead and with that increasing error rates due to storage, especially if only 2D cluster states can be created.

8.2.4 Topological cluster states

The topological cluster state [FG09, WF14b] is an encoding similar to the surface code in a special three-dimensional graph state (see Chapter 7.4.3). The resource state is consumed along one dimension in time by stabilizer measurements. It can in principle be created on-the-fly, however, this requires many parallel successful CZ gates. If this is possible, then only a few layers need to be operated and thus existing in parallel. Otherwise, a huge entangled 3D lattice is required, which very fast goes beyond the scope of experimental realizable implementations, either in terms of space / connectivity (atoms, solid state) or success probability (linear optics) and especially in terms of error rates due to long storage times.

8.3 Oscillator encodings

There are several approaches making use of the high coherence time of harmonic oscillators. States that are encoded in an oscillator can have for example the form of (minimum uncertainty) coherent states, squeezed states, i.e., eigenstates of a quadrature operator, or Fock states, i.e., eigenstates of the photon number operator. A detailed introduction to oscillator states can be found in [Fox06]. All these states are not necessarily restricted to two levels, they can be defined in d -dimensional or even continuous Hilbert spaces.

8.3.1 Continuous variables

In the continuous variables approach, qubits are replaced with a continuous set of squeezed states, each qubit corresponding to the state of one oscillator mode (frequency). The analog of the $|1\rangle, |0\rangle$ basis and the $|+\rangle, |-\rangle$ basis are often chosen to be eigenstates of the two quadrature operators (for example position and momentum, respectively). Similar to the Hadamard gate, an eigenstate of one quadrature can be transformed to an eigenstate of the other quadrature by a Fourier transform. In the same manner, analogies for all Clifford gates can be found for continuous variables by Gaussian operations (non-Clifford gates require photon-number resolving operations/measurements or other higher-order nonlinear effects), with the addition of extra variables, creating a continuous set of CV operators out of one qubit operator. Especially, continuous variable entanglement, as required for cluster state creation, can be implemented by an (ideally infinite strong) multi-mode squeezing process as happening automatically in nonlinear media. Furthermore, all modes are stored in the same cavity or waveguide, so even big systems do not necessarily need much space—but good frequency-resolving devices.

A big disadvantage, however, is that the ideal quadrature eigenstates would require infinite squeezing to create and also other ideal operators are unphysical in their implementation. Hence, only approximate eigenstates can be created and approximate gates are performed, which already shows the strong need for error correction in this protocol.

A first level to estimate the accuracy of states and gates is measuring the amount and direction of squeezing that can be reached, or the general distribution of a state in phase space, which can be done analogously to process tomography with homodyne detectors.

8.3.1.1 Error correction for continuous variables

Error correction protocols for Gaussian states that only use Gaussian operations can not correct Gaussian errors¹². This statement, proved in [NFC09] restricts fault-tolerant continuous-variable quantum computing to codes using either non-Gaussian computational states [GKP01], non-Gaussian operations [RGM⁺03], or codes that are only tolerant to non-Gaussian errors [vL10]. However, typical error sources in experiment include also loss and thermal noise, which are both Gaussian. The first code introduced for continuous variables [Bra98] was a Shor code, which was obtained by mixing ancilla modes via beam splitters. This code was only using Gaussian methods and thus could only correct non-Gaussian errors. More promising approaches are the GKP [GKP01] encoding proposed by Gottesman, Kitaev and Preskill, which encodes information in a discrete subspace of non-Gaussian states. Besides a rather high threshold condition and a preference for small errors occurring continually (having low efficiency for rare but large errors), it needs highly nonlinear interaction for the state preparation like for example the cross-Kerr effect in nonlinear media. To be able to correct large errors, too, the code must be concatenated with other codes. The states needed for GKP can also be approximated by highly squeezed states, and the deviation from a perfectly squeezed state can be modeled as a gate error. This way, for a cluster-state implementation, a threshold for the squeezing strength was found to be 20.5 dB [Men14], when using the GKP scheme for error correction.

8.3.2 Cat code

The cat code [OPH⁺16, LKV⁺13] is a way of encoding a qubit into a superposition of coherent states of an oscillator, designed for correcting photon loss errors. The cavity is coupled to a single qubit for control with which in principle, universal computation is possible [HRO⁺16, MLA⁺14]. Photon loss is detected due to a change in parity. Although this method increases the lifetime of such encoded qubits, it is not possible to

¹² A Gaussian state is any state with a Gaussian characteristic function, i.e., with a Gaussian distribution in phase space. Typical Gaussian states include vacuum, coherent and squeezed states. Gaussian operations are operations due to Hamiltonians which are at most quadratic in the quadrature/ladder operators. They map Gaussian states to Gaussian states.

reach arbitrary low error rates by concatenation or increasing the ancilla space. There are proposals on using $2n$ superpositions to create a n th order encoding which can correct multiple photon losses, but the higher average photon numbers required for these superpositions introduce even more errors [Mir16]. One idea that might lead to a threshold consideration is to use multiple coupled modes of the resonator, but there are no results in that direction yet. Thus, it seems like the cat code may only be used as a first layer of error correction, if combined with other, standard error correction protocols. In this case, the cat code just gives a different qubit platform with a better start-point for the fidelity. However, this platform is still outperformed by other, conventional ones in terms of state and gate fidelities, gate times and space cost.

Extremely recently, a first path towards concatenation was proposed [AND⁺17], but not fully analyzed.

9 Algorithmic innovations

Subsequently we provide a description of algorithmic innovations in the area of quantum computing which are relevant for judging the security of currently prevalent cryptographic solutions. On the asymmetric side, we restrict, with one exception, to discussing algorithmic innovations that have been identified as useful for decomposing integers into prime factors and for solving the discrete logarithm problem in a suitable finite cyclic group. In particular, we consider the discrete logarithm problem on elliptic curves over a finite prime field.

9.1 Quantum circuit versus adiabatic quantum computation model

The quantum cryptanalytic literature focuses on the quantum circuit model and on minimizing the number of qubits, quantum gates, and the circuit depth as critical parameters. Depth and gate counts for T -gates are often considered separately, to facilitate accounting for the implementation cost of this non-Clifford gate. *Adiabatic quantum computation* [AL16] offers an alternative approach, but already in 2001, van Dam et al. showed, that adiabatic quantum computation can be simulated efficiently by quantum circuits [vDMV01], and Aharonov et al. established that the converse direction holds, i.e., the quantum circuit model is in fact polynomially equivalent to adiabatic quantum computation [AvDK⁺07]. If one is willing to consider 5-local operations, a reduction polynomial of degree five (in the number of two-qubit gates) is essentially sufficient in Aharonov et al.'s result. While being satisfactory from a theoretical point of view, it is not clear how to actually implement such operations. Aharonov et al. show that it is in principle possible to restrict to 2-local nearest neighbor Hamiltonians, but this comes at the cost of invoking six-state particles on a two dimensional grid. The experimental realization of the latter is unclear.

At this point, the cryptanalytic significance of the adiabatic approach for realistic cryptographic parameters remains unclear. While there is some interesting work on toy parameters available, a reliable way to extrapolate from these results to genuine cryptographic parameters is lacking. For instance, Dattani and Bryans' work on factoring 56,153 with only 4 qubits [DB15] in the adiabatic regime shows interesting potential for the operand size that can be considered when translating the integer factorization problem to an optimization problem, but it is not clear how their finding can be leveraged to factor a realistic RSA modulus. Similarly, Schaller and Schützhold's work [SS10] evidences that one can solve the factoring problem for an RSA modulus more efficiently than a generic NP problem with the adiabatic approach, but a quantifiable impact for realistic RSA parameters is unclear. Differing from the literature on quantum circuits, there is at the moment no obvious road map or implementation strategy available on how to apply adiabatic quantum computation for computing a discrete logarithm or for factoring a real RSA modulus.

On the symmetric side, the potential improvement over a Grover-based exhaustive search suggested in the discussion of the Tiny Encryption Algorithm in [SS10] deserves mentioning. However, also in this case there is no clear statement about the expected running time available, and for established block ciphers (including AES) no non-trivial resource analysis of the adiabatic approach is available in the literature. Despite the polynomial equivalence with the circuit model one could hope for an improvement in the exponent, but the current literature does not offer a sufficient foundation to make reliable quantitative estimates. Consequently, in our discussion we focus on the quantum circuit model.

The preprint [JBM⁺18a] shows a method to perform prime factorization on a quantum annealer [AL16]. It does not follow the description from the old version of the study, namely, to take Shor's algorithm and map it onto the annealer by including a clock register [AvDK⁺07]. That method is very indirect, but it has the advantage of proving that the speedup gained from a quantum circuit also occurs in adiabatic quantum computing, at least at zero temperature, at the expense of space overhead.

The new method is a more direct implementation of factoring as an optimization problem. It is qubit efficient in principle, but there is significant overhead created by the need to simulate the required interactions on the restricted d-Wave architecture. This could be reduced on the improved coupling architectures currently developed.

The main challenge in investigating claims based on adiabatic quantum computing / quantum annealing is the careful benchmarking of speedup. The paper explicitly refrains from any statement to this end, see second paragraph of its conclusion. We describe this challenge along this paper in two ways:

First, orthodox adiabatic quantum computing requires the quantum computer to remain in the ground state at all times thus mandating a duration of the annealing schedule proportional to the inverse minimal spectral gap of the problem. This is a sufficient condition at zero temperature. An annealing schedule thus has polynomial time scaling, if that minimal gap drops polynomially in system size. For general constrained optimization problems, this gap is at least conjectured to drop exponentially. Again, an analysis of scaling of the gap is not given in the paper. While previous work of Aharonov et al. [AvDK⁺07] states that any gate-based algorithm can be mapped onto an adiabatic algorithm without changing time complexity, this construction is not used in Jiang et al.

While this is a zero-temperature argument, one needs to observe that for a large problem, the spectral gap even if only polynomial in problem size will dive below the experimental temperature. This requires either error correction or genuine quantum annealing, i.e., quantum-assisted relaxation from the low-lying excited states to the ground states. This is what is typically happening in d-Wave machines.

Quantifying complexity and speedup in quantum annealing in a reliable mathematical way has not been achieved in the literature. No state-of-the-art analysis has been presented in [JBM⁺18a]

In summary, while this paper shows an approach to factor on an annealer, it does not give any indication of quantum speedup. While not impossible, speedup is unlikely. Full investigation of speedup would require a full research project with major access to hardware.

9.2 Quantum circuits without error correction?

For low-depth circuits with a small number of gates, one could in principle consider a scenario without error correction, based on the idea that imperfections in the experimental realization might not significantly reduce the success probability. For cryptanalytic applications, the literature does at this point not offer quantum circuits that have been shown to meet these criteria. Low-depth solutions have been considered for solving the discrete logarithm problem on particular elliptic curves [RS14], but this comes at the cost of a large number of gates and qubits. In a similar vain, Cleve and Watrous [CW00] show that the Quantum Fourier Transform (QFT)—which is at the heart of Shor’s algorithms—can be realized in logarithmic depth, but for the number of gates needed, only a polynomial bound is available.

Still, in view of the overhead incurred by error correction (cf. Section 7.5), one may ask if we might be able to tolerate errors at the gate-level without impeding the logical correctness of a cryptanalytic algorithm. Indeed, Nam and Blümel (see [Nam17, NB15b, NB15a]) make a case that a QFT implementation can perform very well even in the presence of noise and gate defects—thus suggesting that if the QFT is performed at the end of Shor’s algorithm, one could try to be lenient with error correction. One may also hope to simplify the QFT by passing to an approximate QFT (see [Cop94, NSM20]), but for state-of-the-art implementations of Shor’s algorithms the logical gate cost is dominated by the arithmetic portion (see Section 9.7). State-of-the-art implementations of Shor such as [RNSL17c] save qubits by using a semi-classical QFT variant, with repeated (single qubit) measurements, where the rotations needed are chosen adaptively (in dependence on measurement outcomes so far), and savings/avoidance of error correction in the arithmetic would be particularly valuable. One can expect that any “accidental error tolerance” of arithmetic operations will depend on specific algorithmic choices (e.g., how exactly is a modular multiplication implemented, or how exactly is a point addition on an elliptic curve realized?). A common approach for the arithmetic tasks is to start with a reversible circuit which is then further decomposed into Clifford and T -gates—resulting in various options, e.g., to decompose a Toffoli gate (see [AMMR13, Jon13]). Having said this, there is very limited literature on error tolerance of arithmetic in Shor’s algorithm. Notably, in [Nam17], Nam considers an implementation of Shor’s algorithm for factoring in the presence of errors in the angles occurring in elementary gates used. Due to resource constraints, the simulations he reports are restricted to very small examples (Chapter 9 discusses a factorization of 21), and making extrapolations for the arithmetic in

cryptographically relevant factorization problems from this limited data set seems problematic. In recent work [NB17] on working with imperfect gates, the question to what extent errors can be tolerated in a large-scale (cryptanalytic) computation still remains open. In [NB15a], one particular adder design is considered and identified as quite robust against gate errors, but it remains open to what extent this can simplify a full-scale implementation of Shor's algorithm. Taking into account debugging considerations, implementing a Toffoli-based arithmetic (cf. [HRS17, RNSL17c]) may in fact be considered as preferable over a (QFT-based) adder design as considered in [NB15a].

Work predating Nam and Blümel's on the robustness of Shor's algorithm in the presence of errors is due to Devitt et al. [SJD06]. They consider specifically the quantum period finding (QPF) subroutine of Shor's algorithm and explore if a more lax error bound than imposing a precision of about $1/(\text{depth} \times \text{qubits})$ can be achieved. To test this, they apply three different discrete errors (bit flip, phase flip, both) randomly to the QPF portion of Shor's algorithm. Each number of errors was simulated 50 times for specific factorable numbers with a binary length L ranging from 5 to 10 (invoking $2L + 4$ qubits) to determine how many errors were allowable until the result was no longer useful. Their results suggest that for larger L , more errors were acceptable. For example, when $L = 5$, at most 15 errors were acceptable before the result was unrecognizable from random, but with $L = 8$, up to 40 errors could be allowed. However, even a single error for $L = 5$ reduces the probability of success to 0.34. These results suggest that the precision of $1/(\text{depth} \times \text{qubits})$ can be reduced to $p(L)/(\text{depth} \times \text{qubits})$ where $p(L)$ is monotonically increasing and at least linear in L . Devitt et al. note that the greatest benefit of these results is for small simulations of QPF where observing the quantum process is the goal and extensive quantum error correction may not be feasible. However, for large factoring problems (such as attacking cryptographically relevant RSA parameters) extensive error correction will still be required since the overall size of the quantum algorithm grows much faster ($O(L^4)$) than this error rate.

Recent work on variational quantum factoring (VQF) [AOGC18a] offers an alternative to Shor's algorithm to find the prime factorization of an integer using a hybrid quantum-classical algorithm. These hybrid algorithms like the Quantum Approximate Optimization Algorithm (QAOA) [FGG14] allow in some cases to benefit from quantum advantage with short segments of algorithms hence compatible with a rather large logical error rate. They employ the variational principle to find approximate solutions to a given problem by encoding the problem in a Hamiltonian whose ground state corresponds to the approximate answer one is seeking. They depend on a heuristically chosen Ansatz to probe the Hilbert space around an initial guess for the ground state.

The paper [AOGC18a] maps the problem of finding the prime factors to an Ising Hamiltonian, whose ground state measurement statistics guarantee to render its factors while wrong answers incur energy penalties through the mapping and thus lie in higher energy states. By using efficient classical preprocessing, they are able to greatly reduce the number of qubits needed. They provide empirical data claiming that using their preprocessing methods, they require about only 50 qubits to factorize a number of size 10^5 , a threefold improvement compared to the qubit requirements without their pre-processing method. They simulate their algorithm under the assumption of noise by a Pauli error channel. They check their algorithm for integers 35, 77, 1207, 33667, 56153, and 291311, and find that VQF can in principle find the prime factors, even though in some cases it performs rather poorly if certain symmetries are violated. As all variational quantum algorithms, it remains to be seen whether their approach scales asymptotically under realistic conditions and growing system size.

[AOGC18a] gives numerical evidence that prime factoring using VQF is possible. The power and limitations of VQF are those of QAOA and much work lies ahead for numerically and analytically studying its capabilities and potential speedups over classical approaches.

Overall, the question of intrinsically fault-tolerant cryptanalytic algorithms is still open. The existing literature does not provide a roadmap to avoid quantum error correction in a serious cryptanalytic application. With the current state of the literature, explicitly budgeting overhead for error correction in quantum circuits for cryptanalysis appears adequate.

9.3 Quantum circuits for cryptanalysis—Key players

The number of groups that make a dedicated effort to pass from asymptotic discussions of quantum cryptanalytic algorithms to explicit circuits is still small, and collaborations among the groups are very common. Important players (in alphabetic order) are

Dalhousie University’s Dept. of Math. & Stat., Halifax, Canada. Peter Selinger is a key person for the development and maintenance of the Quipper language [Sel16], arguably the most popular open source tool for experimenting with quantum algorithms and quantum circuits. Peter Selinger is an expert in optimizing circuits in the Clifford+ T model and in quantum resource estimation.

École polytechnique fédérale de Lausanne, Switzerland. Here, a research team including Mathias Soeken explores techniques to compile a higher-level operation into quantum circuits comprised of elementary gates as supported by quantum processing units from IBM and Rigetti [SMS+19]. Soeken now also works at Microsoft.

Eidgenössische Technische Hochschule Zurich, Switzerland. Here, a Python-based open-source software framework for quantum computing was initiated, known as ProjectQ [SHT18a] Founders of this project are Thomas Häner, Damian Steiger, and Matthias Troyer.

Florida Atlantic University’s CCIS, Boca Raton, FL, USA. Research at FAU’s Center for Cryptology and Information Security includes results on quantum circuits to attack block ciphers with Grover’s algorithm (including AES, MARS, and SERPENT) and with Simon’s algorithm, quantum circuits for elliptic curve arithmetic and for parallelizing Shor’s algorithm for the discrete logarithm problem on elliptic curves. The university also put in place an MoU with South Korea Telecom, a commercial player in the quantum technology space.

IBM Research – Ireland. With QISKit [IBM18], IBM offers an open-source framework to experiment with quantum circuits, including an interface to an actual quantum computing device. The team includes Adi Botea, Akihiro Kishimoto, and Radu Marinescu who contributed to understanding the complexity of quantum circuit compilation [BKM18]. Dmitri Maslov, Chief Software Architect, is one of the leading experts in designing optimized quantum circuits, and he co-authors work on implementing Shor’s algorithm for the discrete logarithm problem on elliptic curves.

Inria Paris, France. Gaëtan Leurent, Anthony Leverrier, and María Naya-Plasencia explore in joint work with Marc Kaplan (University of Edinburgh’s School of Informatics) quantum cryptanalytic attacks against symmetric primitives other than Grover’s algorithm. Their work includes in particular work on quantum versions of differential and linear cryptanalysis. They also explored the use of Simon’s algorithm for cryptanalysis in connection with quantum queries to a device storing a secret key.

Institute for Quantum Computing (IQC) in Waterloo, Canada. Work in this group includes efficient implementations of the QFT—both Richard Cleve and John Watrous, the authors of the highly relevant paper [CW00]—are members of this group. Michele Mosca, a co-founder of IQC, is one of the most prominent researchers in quantum computing, and he is part of a team that designed a quantum circuit for a preimage attack against state-of-the-art hash functions. The group is also very active in identifying techniques to optimize quantum circuits.

Microsoft Research’s QuArC group in Redmond, WA, USA. Microsoft’s Quantum Architectures and Computation Group developed the LIQ*Ui* toolsuite [Res16], a software tool to design and experiment with quantum circuits. Martin Roetteler, one of the group members, was in particular involved in the design of quantum cryptanalytic circuits for AES, SHA, the discrete logarithm problem on elliptic curves, and a quantum-related key attack. Krysta Svore designed in collaboration with Paul Pham (University of Washington) an efficient 2D-implementation of Shor’s algorithm for factoring. More recently, Microsoft introduced with Q# a dedicated programming language to express quantum algorithms [Mic18].

QuICS in College Park, MD, USA. The Joint Center for Quantum Information and Computer Science is formed through a partnership of the University of Maryland with NIST. Yi-Kai Liu worked with Eddie Schoute on quantum circuits for applying Grover’s algorithm against lightweight block ciphers.

Technische Universiteit Eindhoven, The Netherlands. Daniel J. Bernstein (also affiliated with the University of Illinois at Chicago, IL, USA) and Gustavo Banegas have been working on *low-communication parallel quantum multi-target preimage search* [BB17]. Among various other contributions, Daniel J. Bernstein also co-authors work on a *low-resource quantum factoring algorithm* [BBM17]. He and Tanja Lange are experts in optimizing cryptographic implementations, and work by this group can be expected to offer further improvements in the exact cost of quantum attacks in terms of the number of qubits and gate counts.

University of Kentucky’s College of Engineering. Himanshu Thapliyal co-authors several publications on the design and synthesis of efficient quantum circuits. He enters the field from the angle of reversible computing and contributed in particular to circuits for basic arithmetic tasks. The latter is relevant both for factoring and for computing discrete logarithms with Shor’s algorithm.

9.4 Minimizing quantum circuits

When discussing quantum circuits, different elementary gate sets are possible, and it is not uncommon in the literature to start out with classical reversible circuits, which are then translated (without further low-level optimization) to a particular universal gate set. It is reasonable to assume that such “naïvely compiled” circuits are in general not optimized yet. Minimizing quantum circuits at the lower level is an active research area, and much emphasis is currently placed on the Clifford+ T gate set. The latter can be implemented in a fault-tolerant manner, e. g., by means of surface codes [FFSG09, FMMC12]. Using number-theoretic tools, Kliuchnikov (now a member of Microsoft Research’s QuArC group) showed how an arbitrary unitary transformation on n qubits can be approximated with precision ϵ using a Clifford+ T circuit of size $O(4^n n(\log(1/\epsilon) + n))$ and two ancillas [Kli13]. This approach is optimal, if the number of qubits is fixed. For 4-bit circuits, the problem of finding optimal reversible decompositions has been solved already [GFM10], and from a cryptanalytic angle this is rather useful. For instance, the block cipher Serpent’s S-boxes operate on four bits, and in an exhaustive key search with Grover’s algorithm this result can be leveraged to derive efficient quantum implementations of Serpent’s non-linear part.

Current reality in the quantum cryptanalytic literature is that circuit minimization relies heavily on heuristic techniques and manual optimization. A popular technique is the use of algorithmic tools from permutation group theory on the reversible level (cf. [GLRS16]). However, with powerful software tools like ProjectQ, Q#, or Quipper being available now, automated tools get broader use in the quantum cryptanalytic community, leading to improvements in the derivation of efficient quantum circuits for cryptanalysis. In recent quantum cryptanalytic work, measurement-based uncomputation has gained popularity to implement pertinent arithmetic (e.g., an AES S-box [JNRV20] or computations on an elliptic curve [HJNRS20]). Instead of translating Toffoli gates in a classical reversible circuit directly into a Clifford+ T circuit, AND-gates are used. To realize such an AND gate (and thereby multiplication in $GF(2)$), an ancilla qubit is used, and uncomputation involves the execution of gates conditioned on the outcome of a measurement. Key feature of this approach is that – at the cost of a measurement and conditioned operations – T -gates can be avoided. For instance, an AND-gate implementation described by Gidney [Gid18, Figure 3] involves 4 T -gates (along with some Clifford gates), and the uncomputation can be done without any T -gates (but involves a measurement).

9.5 Algorithmic innovations with relevance for symmetric cryptography

The most prominent quantum algorithm that is known to be applicable to the analysis of symmetric cryptographic primitives is Grover’s algorithm [Gro96], and we start by taking a look at this search procedure.

9.5.1 Grover's algorithm

Even though the running time improvement of this algorithm over a classical solution is only in the order of a square root—and therefore an exponential time bound still remains exponential—the speed-up is relevant when quantifying security margins. In general, Grover's algorithm is a versatile tool for hybrid attack strategies: one tries to rephrase a(n exhaustive) search inside some classical cryptanalytic approach in such a way that the requirements of Grover's algorithm are met. Ideally, one can in this way expedite a time-critical component of the classical attack with a “quantum subroutine,” possibly even with an asymptotic gain. Arguably the two most prominent cryptanalytic applications of Grover's algorithm are

- Speeding up an exhaustive key search against a block cipher.
- Speeding up a preimage search against a hash function.

Loosely speaking, Grover's Algorithm can complete a search of a space of size 2^n in $2^{n/2}$ steps (with very high probability), therewith offering a substantial speed-up over a classical search that would take on average 2^{n-1} steps. At the core of the algorithm is a Grover operator which encodes a predicate that decides if a candidate element meets our desired search criteria. If there are M elements satisfying this predicate in the search space of size N , the Grover operator needs to be applied $O((NM)^{1/2})$ times. In case of a uniquely characterized secret key (through a collection of plaintext-ciphertext pairs), the total number of times the Grover operation would need to be run can be calculated easily based on the key size. For a key of size n , precisely $\lceil (\pi/4) \cdot 2^{n/2} \rceil$, or approximately $2^{n/2}$, but this can only give a lower bound on the attack costs (qubits, gates, and depth), as the implementation cost for the encryption scheme itself plays an essential role—the details of the Grover operator depend on the targeted primitive.

So while using a high-level description of Grover's algorithm to compute the cost of breaking symmetric cryptographic systems such as AES- k ($k = 128, 192, 256$), MARS, SERPENT, SIMON, SPECK, etc. is the right approach, the details of the cost can vary greatly and rely heavily on the key size as well as the implementation complexity of the cryptographic system.

Note: Grover's search algorithm was proved optimal for quantum searching, and it allows no non-trivial parallelization [Zal99]; improvements would require an attack on the targeted cryptographic scheme itself.

In cryptographic terms, suppose we have a symmetric encryption scheme F that takes a 128-bit key k as input to encrypt a plaintext P into a ciphertext $C = F_k(P)$. In order for Grover's Algorithm to work, we would need a plaintext-ciphertext pair (P, C) and a quantum realization of the symmetric key encryption scheme F . The result of the algorithm will be the appropriate key k^* , which when used in the encryption scheme yields the correct ciphertext with high probability. To characterize the target key uniquely (or at least reduce the number of candidates to a small set) multiple plaintext-ciphertext pairs may be needed—a typical estimate being 2 or 3. This causes no fundamental difference for mounting the attack, but impacts the number of qubits or gates.

Grover's algorithm creates a superposition of all candidate keys, so that each key has equal probability. The algorithm runs the superpositioned key through F^* , which is a Boolean function that returns 1 if and only if the key is the correct key and 0 otherwise. A key k being correct translates into the condition $F_k(P) = C$. Owing to the superposition, each possible key is in effect tried simultaneously and the one correct key (here we are assuming there is only one correct key) will be “tagged.” Once the correct key is “tagged,” the second phase of the Grover algorithm, the transform, is run. This transform increases the likelihood of the correct key being produced when measured. These two phases represent one iteration of the Grover algorithm. If the quantum key was measured after just one iteration, it would not only fall out of superposition and thus end the ability to proceed, but also the probability of the correct key being produced is only minimally more than that of producing a random key. However, since each time the two phases are run, the probability increases, if measured after the correct number of iterations, the correct key would be produced with high probability.

Note that to evaluate the Boolean function F^* , the full encryption process must be implemented on the quantum hardware and then its result can be compared to the known ciphertext. While the final

comparison is a simple and short quantum operation, the depth and cost of implementing the encryption scheme can vary drastically and is needed at least once in each iteration of the algorithm. This means not only will, say AES-256, take more iterations of Grover than AES-128, each iteration will probably require more quantum gates and qubits, increasing the overall cost further. This additional cost may or may not be negligible in comparison to the Grover operations, but for a system such as AES-128 which would take approximately 2^{64} iterations of Grover, it is a pertinent factor to consider in a quantitative analysis.

Case study: the AES family. The Advanced Encryption Standard (AES), designed in 1998 by Rijmen and Daemen and accepted by NIST in 2001 [NIS01] as the replacement for DES (Data Encryption Standard) [NIS99] is a subset of the Rijndael cipher [DR99]. AES encrypts with three different key sizes, 128, 192, and 256 and all three have been adopted world-wide and are of cryptographic interest.

The resource estimates for a quantum implementation of AES [NIS01] found in [GLRS16] describe the costs of the three variants of AES, AES-128, AES-192, and AES-256 at the logical level. Based on this analysis, a quantum implementation of AES-128 would require 1,380,420 Clifford gates and 1,060,864 T gates with a T -depth of 50,688 and overall depth of 110,799 using 984 qubits. For AES-192 the number of Clifford gates is 1,567,296 with 1,204,224 T gates and a T -depth of 44,352 and overall depth of 96,956 on 1,112 qubits. Finally, AES-256 required 1,956,099 Clifford gates and 1,505,280 T gates with a T -depth of 59,904 and overall depth of 130,929. The total number of qubits for AES-256 is 1,336. It should be noted that these are upper bounds in the sense that further optimization is possible. For instance, Almazrooie et al. [ASA⁺18] show that AES-128 can be fit on 928 logical qubits already (instead of the 984 used by Grassl et al.). Making use of a different approach to representing the S-box in AES, [LPS20] present a more efficient quantum circuit to implement AES, requiring only 864 qubits for AES.

While these are not the final Grover numbers, there is also an additional computation, explained in [GLRS16] as well as in [RS15]. Since AES operates on 128-bit plaintexts, it is plausible that a single plaintext-ciphertext pair may not produce a unique secret key. An argument is made that in order to have uniqueness, 3 plaintext-ciphertext pairs are sufficient in AES-128, four in AES-192 and five in AES-256. This means the numbers computed above would be tripled, quadrupled or quintupled and then Grover applied to compute the final numbers. Using a different line of reasoning (cf. [AGL+18]), by now the estimated number of plaintext-ciphertext pairs needed to ensure uniqueness of the target key is lower, and when facing multiple plaintext-ciphertext pairs, depth can be traded for the number of qubits. The table below gives resource estimates taken from [LPS20] for a key search with Grover's algorithm.

For comparison, we include in the table also the resources for a preimage attack against SHA based on Grover's algorithm, which we discuss next. It is clear that in all cases the circuit depth/the running time is the limiting factor—the number of qubits is by no means extreme.

	Clifford	T-gates	T-depth	Depth	Qubits
AES-128	7.06×10^{24}	3.55×10^{24}	2.18×10^{23}	8.40×10^{23}	865
AES-192	7.10×10^{34}	3.49×10^{34}	8.18×10^{32}	3.19×10^{33}	1,793
AES-256	3.68×10^{44}	1.82×10^{44}	4.63×10^{42}	1.80×10^{43}	2,465
SHA-256	3.33×10^{45}	1.27×10^{44}	3.76×10^{43}	4.44×10^{44}	2,402
SHA3-256	1.84×10^{46}	2.71×10^{44}	2.31×10^{41}	5.90×10^{42}	3,200

Table 9.1: Clifford+T gate counts for Grover-based key searches according to [LPS20] and [AMG+16a].

In very recent work, Kim et al. [KHJ18] presented a framework to explore time-space tradeoffs for quantum cryptanalytic attacks like a key search in AES. They focus in their analysis on the cost of Toffoli gates rather than on a model at the Clifford + T level, and explore different design choices in parallelizing a Grover-based attack or ensuring uniqueness of the target key. Moreover, [JNRV20] show how AND-gates and measurement-based uncomputation can be leveraged to reduce the T -depth and overall depth in a key

search for AES – at the cost of increasing the number of qubits and introducing measurements. Still, the exponential scaling of Grover’s algorithm remains a formidable hurdle, and no feasible quantum attack against AES has been identified so far. The quantum security analysis of AES in [BNPS19] (which still builds on the gate counts in [GLRS16]) concurs with the positive view on the post-quantum security of AES-256 in regard to the infeasibility of a key search, but – referencing [CNPS17] – points out that in moving forward the 128-bit size of the internal state may offer an avenue for quantum cryptanalytic progress for certain modes of operation.

9.5.2 Pre-image search for a hash function—the case of SHA

The Secure Hash Algorithm (SHA) is a family of hash functions standardized by the National Institute of Standards and Technology (NIST) starting in 1992, with the SHA-3 versions being added in 2015. While a collision was recently found in SHA-1 [SBK⁺17], first published in 1995, there have been no collisions found in SHA-2, which was published in 2001 and designed by the National Security Agency (NSA) or SHA-3, published in 2015 [NIS15a]. The latter was designed by Bertoni et al. and is a subset of the Keccak family of cryptographic functions [BDPA11]. Both SHA-2 and SHA-3 are still in use today and are of more cryptographic interest.

Mosca et al. discuss in [AMG⁺16a] the costs of a generic pre-image attack on two specific variants of the SHA family of hash functions. Specifically, SHA-256 [NIS15a] and SHA3-256 [NIS15b]. The paper computes these costs and then uses a quantum circuit optimization tool “*T-par*” [AMMR13] to reduce the number of *T* gates and *T*-depth. This *T-par* optimization tool has the ability to introduce the *P* gate. This gate is equivalent to a T^2 gate, and we count it for gate counts and depth like a *T* gate.

A quantum implementation of SHA-256 would require 648,640 Clifford gates, 401,584 *T* gates, a *T*-depth of 144,786, and an overall depth of 528,768 [AMG⁺16a]. When optimized using *T-par* these numbers change to 6,129,072 Clifford gates and 301,968 *T* gates with a *T*-depth of 70,400 and overall depth of 830,720. Both of these computations require 2,402 qubits. In comparison, an optimized quantum implementation of SHA3-256 would need 34,429,525 Clifford gates and 499,200 *T* gates with a *T*-depth of 432 and overall depth of 11,040 using 3,200 qubits. For computing $(\pi 4) \cdot (2^{128})$ total Grover iterations, the total numbers for SHA-256 and SHA3-256 can be found in the table above—not surprisingly, the running time remains a limiting factor of the attack.

Note: Recent work by Banegas and Bernstein [BB17] considers a multi-preimage search with a quantum algorithm. They combine Grover’s technique with a reversible parallel rho-algorithm and make a case that *quantum preimage search benefits asymptotically from having multiple targets*. At this point, their analysis focuses on asymptotics, and gate-level resources counts are not available yet.

9.5.3 Collision search with a quantum algorithm

The standard use of Grover’s algorithm is to determine a secret key that matches a plaintext-ciphertext pair (or several of them). However, one can redefine the algorithm to find a hash collision. By default, Grover encrypts a known plaintext using candidate keys in superposition and returns a ‘1’ if the computed ciphertext is equal to the known ciphertext and a ‘0’ otherwise, and (for a unique solution) after $O(N^{1/2})$ iterations we have a good chance to measure the correct key.

Now, imagine instead, the plaintext is in quantum superposition and the algorithm only returns a ‘1’ when the ciphertexts are equal but the plaintexts are not, thus finding a collision. After the same $O(N^{1/2})$ iterations, the algorithm would return a plaintext distinct from the given plaintext that produces the same hash value. If we take that idea and combine it with the birthday attack, we can speed-up the computational time of the algorithm using the traditional time/space trade-off of finding a random collision (as opposed to a specific solution).

Before proceeding, recall that Grover's Algorithm has two similar but different forms, depending on whether the number of solutions is known or not which is directly related to the number of expected collisions here. If the specific number of collisions is known, the more simple form of Grover can be applied while an unknown number of solutions requires the use of the more generic form of the algorithm found in [BBHT98].

While hash-based functions like SHA-256 [NIS15a] and FORK-256 [HCS⁺06] take an input of (for practical purposes) arbitrary size and map it to an output size of 256 bits, it might be beneficial to explain the quantum collision algorithm assuming the number of collisions is known and finite.

Assuming the hash function is r -to-one. As explained in [BHT98], assume there exists some random hash function H such that $H : X \rightarrow Y$ is an r -to-one function, meaning exactly r inputs produce each output where $r \geq 2$. Thus, if $|X| = N = 2^l$ then $|Y| = Nr$. If space is available, the best solution requires the computation of a random subset K of X of cardinality $k = (Nr)^{1/3}$ and each tuple stored in a table. This table can be computed on a classical computer and would take k evaluations of H . This list would then need to be sorted and if any collisions are found such that $H(x_i) = H(x_j)$ then $\{x_i, x_j\}$ can be output and the search is over, however this probability is quite low.

While this list can be computed on a classical computer, the table would need to be stored in quantum bits so the modified version of Grover can reference this list of values in the table each iteration. Thus, if $O((Nr)^{1/3})$ storage qubits are unavailable or too costly, the list would need to be reduced which would increase the running time of the algorithm.

The algorithm would compare the computed hash value with all the values in the second column of the table and return a '1' if the output value is found in the second column of the table and the input value is not found in the first column. The algorithm would return a '0' otherwise. After a specific number of iterations a collision would be found with probability $1/2$ and the result would be a plaintext $x \in X \setminus K$ such that $H(x)$ is a value in the stored table.

To complete the process, $H(x)$ would need to be computed and found in the table. If $H(x) = H(x_0)$ for some $(x_0, H(x_0))$ tuple in the table, then $\{x, x_0\}$ is a collision which can be output.

Since k distinct input values are stored for comparison and each output value has r distinct input values that hash to it, the probability of a collision is $r \cdot k/N$. Thus, the expected number of Grover iterations would be about $(Nr/rk)^{1/2} = (Nr/rN^{1/3})^{1/2} = (Nr)^{1/3}$. Since the number of classical computations of H is $k+1 = (Nr)^{1/3} + 1$ we get the expected run time of the algorithm to be $O((Nr)^{1/3})$ times the time it takes to compute the hash function.

Generic hash functions. When less is known about the hash function or even when we just know it is not specifically r -to-one for any $r \geq 2$ the argument above must be slightly modified. Changes must be made to how $K \subseteq X$ is chosen, but the more general version of Grover can be used. Obviously, the smaller the chosen $K \subseteq X$ the longer it will take to find a collision and while a larger K will reduce the number of Grover iterations, the storage and classical computations of the hash will increase.

However, if the input size is known to be a specific finite number or at most some finite number, then $|K|$ can be determined based on the probability of each output being repeated [FHZ14], but it is still $O(N^{1/3})$ where N is the size of the hash space. When searching for a collision in SHA-256 or FORK-256 or other hash functions, this is all that is necessary since the searched input size can simply be fixed to be anything bigger than 256 to guarantee a collision.

Searching for a claw. Another result in [BHT98] is that of finding a claw. A claw is similar to a collision in a hash function, but is a collision among two hash functions [OK91]. Specifically, if F and G are two distinct hash functions such that $F : X \rightarrow Z$ and $G : Y \rightarrow Z$ with $|X| = |Y| = |Z| = N$, then one can find a pair $x \in X$ and $y \in Y$ such that $F(x) = G(y)$ in an expected number of $O(N^{1/3})$ by applying the same algorithm as above expect picking the K subset from before from X and applying the Grover search to Y .

Even though these functions more closely resemble permutations, this algorithm can be again extended to more general r -to-one hash functions in the same way as before. And while it is not expressly stated in

[BHT98], it would seem to still hold for hash functions that are not specifically r -to-one either. For Grover to run most efficiently when looking for a claw, any collisions completely in K should be removed and replaced before continuing. This is because the Grover search is most efficient when the exact number of solutions is known and a collision inside K reduces the probability of finding a collision outside of K . However, since we already assume this probability to be extremely low and are already sorting K , this is a minor additional step.

As a consequence, for any generic hash function where $N = 2^m$ is the size of the hash space, picking a random $K \subseteq X$ such that $|K| = O(N^{1/3})$ yields an expected collision with probability greater than $1/2$ after a run time of $O(N^{1/3})$. The exact run time depends on the size of K , the number of Grover iterations, the cost of computing the hash function and searching for a collision in K . Also, the assumption is that there are $O(N^{1/3})$ quantum bits of storage to run Grover's Algorithm which is non-trivial. Less than this would increase the overall search time which would max out at $O(N^{1/2})$ (the standard Grover run time).

Note: Similar as for AES, Kim et al. [KHJ18] explore trade-offs for attacks against hash functions with Grover's algorithm, including the case of SHA-256 and SHA-384. However, again the exponential scaling of Grover remains a serious obstacle, and no feasible attacks against SHA-256 or SHA-384 are identified.

9.5.4 Simon's algorithm and superposition queries to keyed primitives

Grover's algorithm is by no means the only quantum cryptanalytic tool available to attack symmetric primitives—see, for instance, [RS15, KLNP16, KLLNP16, SS17, BNP18]. Notwithstanding this, it is fair to say that for attacking today's implementations, which are entirely classical, Grover's approach is currently the most relevant tool. Conceptually, *Simon's algorithm* enables an interesting and different type of attack, but from a practical point of view, it is important to pay close attention to how an application/oracle interface is accessed. What is Simon's algorithm? It is a quantum algorithm which can solve the following problem in expected polynomial time, provided that the involved function $f: \{0,1\}^k \rightarrow \{0,1\}^k$ can be evaluated in polynomial time on a superposition of inputs. For the function f it is assumed that $k' \geq k$ and one of the following conditions holds:

1. f is injective, or
2. there is a bitstring s , not entirely zero, such that for every $x \neq x'$ we have $f(x) = f(x')$ if and only if $x = x' \oplus s$.

The task is to decide for a given f which of the two cases holds, and in the second case to determine the "hidden shift" s . From a cryptanalytic point of view, involving Simon's efficient solution to this problem is attractive, when the adversary can actually implement and evaluate f at a superposition. A common problematic assumption in attacks based on Simon's algorithm is that the function f depends on the attacked secret key, so that *quantum access to an implementation of the attacked cipher which stores the attacked secret key* becomes necessary. It is fair to say that for today's classical implementations this assumption is not met.

Related-key attack. To illustrate this point, let us take a brief look at a quantum version of a related-key attack in [RS15], which relies on Simon's algorithm and in principle enables the recovery of the secret key of a large class of block ciphers in polynomial time (measured in the key length). So if quantum access to the keyed primitive is/were indeed possible, the attack is highly potent against symmetric encryption schemes. The setting considered in [RS15] is a related-key attack, where the function f depends on the attacked block cipher (which can reasonably be assumed to be known), but also on the attacked secret key. Access to the latter is in a related-key attack in principle available—commonly modelled through a suitable encryption oracle. However, to bring Simon's algorithm to use, the oracle must accept a superposition of inputs, which for classical implementations is not the case. This is very different from (and less threatening than) Grover's algorithm, where the attacker needs only the specification of the block cipher (plus plaintext-ciphertext pairs) to mount an attack against the secret key. However, if the quantum access were available, the resulting attack would be polynomial time—unlike a key search with Grover.

Modes of operation. Similar as in the case of the related-key attack just mentioned, in [KLN16] Kaplan et al. show how Simon’s algorithm can be leveraged to invalidate the security of popular modes of operation for achieving authenticated encryption or for constructing a MAC from a block cipher. The same paper uses Simon’s algorithm—with the same limitation—to expedite a slide attack. Conceptually, these attacks are interesting, but from a pragmatic point of view they are not an imminent threat for today’s implementations, as the assumptions of the attack model are not met.

Differential and linear cryptanalysis. In [KLLNP16], the study of differential and linear cryptanalysis in connection with quantum attacks is initiated. On the one hand, a scenario with quantum queries to the attacked block cipher is considered, which for today’s implementations on classical platforms may be considered an unrealistic model. However, the paper also makes the point that even when restricting to classical queries, a quantum algorithm in combination with differential and linear cryptanalysis can sometimes yield a more efficient attack than a key search with Grover. In particular for settings where the key size exceeds the message-block size, a quantitative comparison of these different attack strategies can be worthwhile. However, other than the examples in [KLLNP16], so far the literature does not appear to offer use cases, where this strategy has been brought to fruition.

Relating to Simon’s algorithm and superposition access to the attacked cipher, recent progress in [BHNP⁺19] deserves to be monitored. This paper shows that in specific cases, Simon’s algorithm can be leveraged for a key recovery with fewer or no superposition queries to the attacked cipher than was reported before.

9.6 Algorithmic innovations with relevance for asymmetric cryptography

For cryptographic algorithms that rely on the computational hardness of factoring “large” integers or of computing discrete logarithms in a suitably chosen cyclic group, the impact of quantum algorithms appears at this point much more severe than for symmetric cryptography. Leaving aside a possible performance penalty and careful review of quantitative implications for linear and differential cryptanalysis [KLLNP16], doubling the key length appears at the moment a viable approach to counter the speed-up from Grover’s algorithm over a classical exhaustive key search. Similarly, the available quantum speed-up in finding a collision for a hash function is only polynomial. For factoring and computing discrete logarithms, Shor’s seminal work in [Sho94, Sho97] reveals a very different picture: he presented polynomial time solutions for factoring and for computing discrete logarithms, which is very different from the best known classical algorithms, which exhibit, at best, a subexponential run time. Before quantifying quantum resources for factoring and computing discrete logarithms for common cryptographic problem instances, it may be helpful to start out with an asymptotic perspective as presented, e.g., in [BBM17] and [RNSL17c].

Factoring integers. The predominant classical approach to factoring in cryptanalytic contexts is the Number Field Sieve (NFS). For factoring a composite n -bit number, the running time of this algorithm is estimated to be subexponential of the form $(\exp(n^{1/3} \times \log^{2/3} n))^{c+o(1)}$ with a constant c of about 1.902. Bernstein et al. in [BBM17] present a quantum algorithm with a better – but still subexponential – running time: They reduce the exponent c to about 1.387, and at the same time they ensure that the number of qubits needed by their method grows with $n^{2/3+o(1)}$ only, i.e., the growth is sublinear. This is conceptually different from Shor’s algorithm, where the number of qubits needed is linear in n . However, the expected running time of Shor’s algorithm is only cubic in n , i.e., polynomial in the bit length. While there has been progress in constant optimization in the complexity of Shor’s algorithm, achieving simultaneously a polynomial running time and using only a sublinear number of qubits remains a challenge.

Discrete logarithms. The classical literature on computing discrete logarithms in finite fields is currently quite active. From a cryptographic perspective, the multiplicative group of a prime field is arguably the most interesting case, and again an NFS-based technique is available in this scenario. Also in this case, [BBM17] offers a way to speed-up at least one of two phases of the classical algorithm—in the running time of $(\exp(n^{1/3} \times \log^{2/3} n))^{c+o(1)}$ (now n represents the bit size of the field), the constant c can again be reduced from about 1.902 to about 1.387, involving only a sublinear number of qubits. From what we know so far, for

elliptic curves over prime fields, the techniques in [BBM17] do not apply, and for adequately chosen curves, the expected running time of the best available classical algorithm (a parallel version of Pollard’s rho method) is exponential in the bit length n of the group size: $((\pi/2)^{1/2} + o(1)) \times N^{1/2}$, where N is the group size. Shor’s algorithm offers here an exponential speed-up: Shor’s algorithm has an expected running time that is cubic in the bit length of the group size, i.e., it is a polynomial time solution. For implementing this method, an—in the bit length of the group size—linear number of qubits is used, however. Similar as for factoring, being able to restrict the number of qubits to a sublinear range while preserving polynomial running time remains a challenge.

9.6.1 Factoring Integers

Shor’s solution can be expected to find a factor of a composite n -bit number in time $O(n^3)$ —following Proos and Zalka [PZ03], we can estimate the constant to be about 4. Essentially, there are two phases to the algorithm, the second of which is entirely classical. This portion is not hard to implement, but it is worthwhile to put some thought into the implementation of this portion, so that unnecessary repetitions of the first phase, which relies on a quantum computer, can be avoided [Joh17]. The first phase is also referred to as *order finding*: we have to find the order of a randomly chosen $\alpha \in \mathbb{Z}/N\mathbb{Z}^*$, where N is the number to be factored. No efficient classical algorithm is known for this problem, but from Shor’s work we know that the *Quantum Fourier Transformation* (QFT) can be invoked to solve this problem efficiently on a quantum computer. Unlike a classical Fast Fourier Transform (FFT), a QFT can be implemented in polylogarithmic time (namely $O(\log^2 N)$). Cleve and Watrous [CW00] offer a theorem which neatly separates a classical and a quantum portion of factoring an integer—they argue that a polynomial size classical pre-processing and a polynomial-size classical post-processing can be combined with an $O(\log n)$ -depth quantum circuit of polynomial size.

The most expensive operation, and the bottleneck in running Shor’s algorithm, is the computation of a modular exponentiation on a quantum computer: we have to be able to compute $\alpha^k \bmod N$, where k is in superposition. There is ample classical work available on implementing this type of arithmetic, but we do need this arithmetic as a quantum circuit. There is a strong connection with Shor’s algorithm for the discrete logarithm problem in a prime field, as in the latter case we also face an exponentiation task with a known modulus. So the question of implementing modular arithmetic efficiently as a quantum circuit is of key importance to actually use Shor’s algorithm.

The number of qubits needed to factor with Shor’s algorithm is quite modest—Beauregard [Bea03] showed that for an n -bit number a circuit with $2n + 3$ qubits and $O(n^3 \log(n))$ elementary quantum gates (and cubic depth) is available. Zalka [Zal08] argues that $1.5n$ logical qubits suffice. Work by Ekerå and Håstad [EH17] gives additional hope, that factoring might be feasible with a substantially smaller (though still linear) number of qubits. Having the number of qubits proportional to the bitlength n of the number to be factored appears essential with the current state-of-the-art, unless one is willing to sacrifice the polynomial running time (see the discussion of Bernstein et al.’s approach from [BBM17] below.) One of the most recent detailed cost analyses of factoring an RSA modulus is due to Gidney and Ekerå [GiEk19], and also here the number of logical qubits is essentially chosen to be linear (actually, slightly worse): $3n + 0.002n \times \lg n$.

Resource counts and going beyond Shor’s algorithm. It is fairly common that available research papers do not elaborate on how to pass from a (correct) high-level algorithm to an actual quantum circuit. An interesting and elaborate design to factor an RSA modulus with Shor’s algorithm is due to Pham and Svore [PS13]. They employ a 2D *nearest-neighbor* quantum architecture with the following resource counts: depth $O(\log^3 n)$, size $O(n^4 \log n)$, and $O(n^4)$ qubits. Pham and Svore offer explicit bounds, and the constants hidden are non-trivial. E. g., for the number of qubits in a modular exponentiation the multiplicative constant in front of the n^4 is about 95,000 and the number of gates hides a term in the magnitude of $3.5 \cdot 10^6 \cdot n^4$. Still, this result offers an exponential improvement in circuit depth over prior *nearest-neighbor* solutions at the cost of a polynomial increase of gate count and number of qubits.

If we give up the nearest-neighbor restriction and would like to keep the number of qubits small, a design proposed by Häner et al. in 2016 can at the moment be seen as leading contender for a cryptographically relevant implementation of Shor's algorithm [HRS17]. This proposal has been analyzed at a quite detailed level, and the available algorithm analysis is backed by serious LIQUi| > simulation (with input sizes as large as 8,192 bit). In addition, the authors make a case that their Toffoli-Network based arithmetic facilitates debugging when being implemented on a quantum hardware, which from an experimental point of view is indeed a valuable feature. With $2n + 2$ qubits, the width of Häner et al.'s solution is quite moderate, the depth is $O(n^3)$, and the number of gates calculates to $64n^3 \log_2(n) + O(n^3) = O(n^3 \log n)$. If one is willing to invest a larger number of logical qubits, Gidney and Ekerå's [GiEk19] offers an attractive alternative. The paper uses a number of techniques to reduce the cost for the arithmetic. In addition to reducing the circuit depth by factoring through reduction to a short discrete logarithm computation (see below), the arithmetic is optimized in several ways. For instance, instead of a traditional representation of a modular integer as a computational basis state, a so-called coset representation (cf. [Zal08]) is invoked and the exponentiation makes use of windowing. While the number of logical qubits is larger than in Häner et al.'s approach, Gidney and Ekerå's approach offers savings in the number of non-Clifford gates (see Table I in [GiEk19]).

If we give up the nearest-neighbor restriction and accept some uncertainty about the asymptotic cross-over, one of the currently most promising algorithms for factoring integers is from April 2017 and due to Bernstein et al. [BBM17]. This algorithm comes with a heuristic complexity analysis, emphasizes the saving of qubits and restricts to using $(\log N)^{2.3+o(1)}$ logical qubits to factor an RSA modulus N . The running time is about $\exp((\log N)^{1/3}(\log \log N)^{2.3+o(1)})$. Key contribution of this algorithm is that the number of qubits grows *sublinear* in the size of the number to be factored—this is different from Shor's algorithm. Bernstein et al.'s approach starts out with a fast classical factoring algorithm (the Number Field Sieve) and then leverages Grover's algorithm to speed-up the sieving step. Remarkably, the predicate used in Grover's algorithm relies on an implementation of Shor's algorithm—basically, Shor's algorithm is deployed as smoothness test. At the current status of the literature, the exact cost of this elaborate algorithm for a fixed key size (like a 2048-bit modulus) is not clear yet, even though asymptotically this new approach plausibly saves qubits over Shor's algorithm.

Reducing the circuit depth. Several authors proposed modifications of Shor's algorithm with the goal of reducing the requirements on the underlying quantum hardware. At the cost of a (more expensive) classical post-processing phase, the use of a simpler quantum hardware becomes possible. The simpler quantum device may then have to be used multiple times to collect sufficient data for a successful factorization.

Specifically, Ekerå and Håstad [EH17], building on earlier work by Ekerå [Eke16], reduce the task of factoring an RSA modulus $N = pq$ to computing a short logarithm inside the multiplicative group modulo N . This reduction is entirely classical – a discrete logarithm problem is derived, where the discrete logarithm is known to be $p+q$ (with high probability) – which for practical choices of p and q is small compared to the order of the multiplicative group modulo N . The fact that the exponent is small can then be leveraged in the quantum portion of the algorithm to work with smaller exponents. The resulting setting is pretty much the same as for Shor's original approach, but one can reduce the number of T -gates by approximately a factor 4 (cf. [RNSL17c, Remark 3]). When using a semi-classical implementation of the QFT, which is standard, there is not really a saving in the number of qubits, however.

Seifert's earlier work [Sei00] is similar in nature to Ekerå and Håstad's approach in that a more elaborate classical post-processing is used, with the goal of having a simpler quantum hardware that is used multiple times to collect sufficient information to ensure a successful factorization. However, the key obstacle in implementing Shor's algorithm – implementing the modular arithmetic – remains. At this point, it seems fair to say, that the techniques introduced by Seifert, Ekerå, and Håstad reduce the number of gates (and circuit depth) by roughly a factor 4, but do not have a relevant impact on the number of qubits needed. Still, Gidney and Ekerå's recent work [GiEk19] nicely illustrates that for other resource counts (such as gate count and circuit depth), these observations are valuable.

9.6.2 Computing discrete logarithms

Shor’s solution to the discrete logarithm problem is remarkably generic and affects any finite cyclic group of cryptographic interest. Again, the algorithm includes a classical phase, which is easy to implement, and a phase which requires a quantum computer. The pertinent quantum circuit begins with a simple application of Hadamard gates, finishes with a QFT, and in between relies on the availability of efficient group arithmetic: to find the discrete logarithm of $h \in \langle g \rangle$ we need to be able to compute (in multiplicative notation) $g^k \cdot h^{k'}$ where the exponents k and k' are in superposition. The exact cost of this operation will depend on the complexity of the underlying group arithmetic. Just as in the case of factoring, this (double) exponentiation is the bottleneck of Shor’s algorithm—for the QFT portion we can rely on the results of Cleve and Watrous again. As noted by Mosca and Zalka [MZ04], for the discrete logarithm setting, we could in principle even use an exact Quantum Fourier Transform instead of an approximate version (whose dimension is a power of two), but there appears no obvious practical gain in this.

In a preprint from August 2018, Ekerå extends earlier work on trading more extensive classical post-processing for the (repeated) use of a simpler quantum device [Eke18]. Similar as for the case of factoring integers discussed above, this approach does not really impact the number of qubits needed. However, again the number of operations needed in the underlying group can be reduced. This has the potential to cut the number of gates and circuit depth in half. In [Eke19], Ekerå gives a fairly detailed analysis of the success probability of Shor’s algorithm in the case of solving a discrete logarithm problem – suggesting that with moderate overhead, a success probability of more than 0.99 for one run can be achieved.

If the target subgroup is embedded in a finite prime field, we face the task to implement simple modular arithmetic, resulting in a situation pretty much identical to the one for factoring. The fact that our modulus is now a prime number is, as far as the implementation complexity goes, without significance. We can (re-)use the modular exponentiation circuits from Shor’s algorithm for factoring to implement the needed group arithmetic. Perhaps unsurprisingly, detailing circuits for this scenario has not been a topic significant interest in the research literature so far. More recently, [GiEk19] gives some explicit cost estimates, taking into account, e.g., if a discrete logarithm is known to be small. Overall, the cost estimates provided by [GiEk19, Table IV] confirm the similarity in complexity for factoring and discrete logarithms in a prime field.

For elliptic curves, which are at the moment arguably the most popular platform for discrete-logarithm based cryptographic solutions, the situation is different. There is a large body of work on implementing elliptic curve arithmetic efficiently on classical hardware architectures. For efficiency reason, it is therefore tempting to invoke a carefully drafted curve representation to minimize the circuit cost for the (double) exponentiation in Shor’s algorithm. This opens seemingly an interesting degree of design freedom, but an important technicality needs to be considered: For Shor’s algorithm to work, a *unique representation of group elements* must be ensured before the QFT step—this implies that a naïve use of projective coordinates is not adequate. The problem is somewhat similar to the uniqueness requirement for distinguished points in parallel classical implementations of the Pollard-rho algorithm.

A popular algorithm layout to implement the double exponentiation is to juxtapose two sequential variants of a double-and-add procedure [PZ03, KZ04, MMCP09]. Each addition circuit adds a precomputed point, i. e., one operand can always be “hard-coded,” to the current intermediate result if and only if the appropriate bit in the exponent k respectively k' is set. To reduce circuit complexity, this addition is commonly only synthesized for the “generic case” (doubling, addition of the inverse, and identity argument are ignored). High-level modifications can be applied to reduce the depth of such a circuit: Roetteler and Steinwandt [RS14] suggest to parallelize the double exponentiation $g^k \cdot h^{k'}$ with a tree structure that substantially reduces the circuit depth at the expense of additional qubits. The extreme parallelization considered in [RS14] to achieve low depth exploits a uniform addition law on the underlying elliptic curve. In principle this is not needed, but unlike the sequential solution, a full implementation of the addition law—including all “exceptional cases”—is assumed. So far no gate-level analysis of this parallel approach has been published for the prime field case. For binary fields, a depth $O(\log^2 n)$ implementation of Shor’s algorithm is possible, but this comes at the cost of investing many additional qubits, so that a tree structure can be realized. For the prime field case one would have to cope with the same issue—to implement the parallel tree structure, a

large number of qubits would be needed, and implementing the general case of the addition law appears quite costly. It is fair to say, that at this point the “obvious” sequential approach to implement a double-and-add is the most promising approach for realizing the scalar multiplications/exponentiations in Shor’s algorithm for elliptic curves.

As far as the curve representation and ground field arithmetic are concerned, various papers looked at the case of binary fields, but for the prime field case only a very basic variant has been explored more thoroughly:

Direct use of a (short) Weierstraß form with affine coordinates. This is the approach taken by Proos and Zalka [PZ03] and avoids any problems with a non-unique representation of group elements. For performing the necessary ground field inversion, a quantum implementation of the Euclidean algorithm is employed. The resulting implementation of Shor’s algorithm over an n -bit elliptic curve group over a prime field requires about $6n$ logical qubits and a running time of about $360kn^3$, where $k(\approx 1)$ is a constant to reflect the cost $k \cdot n$ of an addition between a classical and a quantum argument. Work by Roetteler et al. [RNSL17c] on the discrete logarithm problem on elliptic curves over prime fields opts for a (short) Weierstraß form with affine coordinates, too, and from the discussion given there, it is indeed not clear that alternative point representations enable a relevant reduction in quantum resources. A recent paper by Häner et al. [HJNRS20] uses again affine Weierstraß arithmetic, but introduces various improvements, including the use of measurement-based uncomputation. The number of logical qubits is about $8n+10.2[\log n]-1$, using a depth of about $509n^3 - 1.84 \times 2^{27}$. A different trade-off reduces the depth to about $2523n^2 + 1.10 \times 2^{20}$, but needs $11n + 3.9[\log n] + 16.5$ qubits. Häner et al. offer optimization for the number of T -gates, too.

Temporary use of projective coordinates. It is a natural idea to work with projective coordinates and convert them back to affine coordinates before the QFT step, so that the number of divisions can be reduced to one. Cheung et al. [MMCP09] used this approach to find a depth $O(n^2)$ circuit for the discrete logarithm problem on ordinary curves over $GF(2^n)$. Subsequent work by Amento et al. [ARS13] and Budhathoki et al. [BS15] goes a similar line and tries to employ a carefully chosen curve representation to reduce the number of T -gates with mixed projective-affine additions. However, no such approach has been explored at the gate level for prime fields. At this point the seemingly obvious approach with a short affine Weierstraß equation can indeed be considered the state-of-the-art.

9.6.3 The Quantum Linear System Algorithm (HHL)

In [HHL09], Harrow et al. proposed a quantum algorithm to efficiently solve a type of linear algebra problem. This quantum linear system algorithm is now commonly referred to as HHL – after the initials of its inventors. The algorithm can provide an exponential speed-up over the best available classical algorithm (based on the conjugate gradient method). An aspect that makes the HHL algorithm potentially interesting from a cryptanalytic point of view is the possibility to use it to solve polynomial systems of equations over $GF(2)$. We give a brief overview of key aspects of HHL, mostly following [SVMA⁺17,DHMS⁺18], and then look into a potential cryptanalytic use, mostly following [Pla19]. The details of the algorithm are fairly involved, and we refer to [SVMA⁺17,DHMS⁺18] for a more elaborate discussion.

Given a Hermitian $N \times N$ matrix A with unit determinant (through a suitable embedding, this condition on A can be relaxed) and a vector b , the HHL algorithm in essence finds a solution to the quantum linear systems problem $A|x\rangle = |b\rangle$, making use of a spectral decomposition. More specifically, the quantum linear systems problem asks that given (oracle access to) the matrix A and given the state $|b\rangle$, to find a state $|x'\rangle$ such that the distance between $|x\rangle$ and $|x'\rangle$ is below or equal some error ϵ with probability greater than 0.5. The running time of HHL to solve this problem is given by $O(\log(N)s^2\kappa^2/\epsilon)$, where s is a parameter characterizing the sparsity of A , and κ is the so-called condition number of A , i.e., the ratio between A ’s largest and A ’s smallest eigenvalue.

The original HHL algorithm has been improved by a number of authors. In [Amb10], the running time is reduced to $O(\log(N)s^2\kappa \text{polylog}(\kappa)/\epsilon)$, and Childs et al. [CKS17] achieve a running time of $O(s\kappa \text{polylog}(s\kappa/\epsilon))$;

further improvements have been reported by Wossnig et al. [WZP18]. Regrettably, the literature does not offer much work on gate-level discussions of HHL and its successors. Work by Scherer et al. [SVMA+17] is somewhat of an exception – their paper offers a case study with a quantum resource analysis of HHL in the Clifford+T model. Using a problem size of $N=332020680$ – which was hoped to be near the cross-over point with the best classical algorithm for an accuracy of $\epsilon=1/100$, the resulting circuit depth has been found to be in the order of 10^{25} or more – not yet taking into account overhead for fault tolerance. Scherer et al. make use of a HHL generalization from [CJS13], which has two key components: (1) quantum phase estimation, involving the QFT and Hamiltonian simulation and (2) quantum amplitude estimation which involves Grover’s algorithm. The first part extracts information about the eigenvalues of A , and according to Scherer et al.’s analysis, the Hamiltonian simulation in this part, which implements an operator of the form $\exp(iAr)$ with $r=O(\kappa/\epsilon)$, turns out to be very costly in terms of gate complexity. In fact, taking this cost into account changes the circuit depth by more than three orders of magnitude, and the estimate for the number of qubits by more than five orders of magnitude. With these high resource counts, the cryptanalytic value of HHL over classical linear algebra tools remains fairly unclear.

Building on HHL, Chen and Gao [ChGa18] established an interesting result that a polynomial system of equations in n variables with a total of t non-zero terms over $GF(2)$ can be solved with gate complexity $\sim O((n^{3.5} + t^{3.5})\kappa^2 \log(1/\epsilon))$ and success probability at least $1-\epsilon$, where $0<\epsilon<1$. Main idea is to translate a (sparse) Boolean system of polynomial equations into a (sparse) system of polynomial equations over the complex numbers, and then use a Macaulay matrix approach to solve this system with HHL. The κ -value in the cost estimate is the condition number of the involved Macaulay matrix, so the obvious question is to determine/estimate such condition numbers. Regrettably, determining actual values for κ in systems of interest appears so far highly non-trivial; even collecting experimental data appears already computationally demanding [Pla19]. Conceptually, Chen and Gao’s result offers a pathway for attacking proposals in multivariate cryptography or to mount algebraic attacks on symmetric ciphers. However, without good estimates for κ , and in the absence of efficient quantum circuits to actually implement HHL, this line of work seems at the moment merely of theoretical interest, and it is not clear that the use of quantum algorithms offers a relevant cryptanalytic benefit here.

9.7 The cost of factoring and computing discrete logarithms in the Clifford+T model

Owing to the superpolynomial efficiency gain over the best classical solution, Shor’s algorithm is particularly influential from the cryptanalytic point of view. Based on Roetteler et al.’s work from June 2017. [RNSL17c], elliptic curves over prime fields appear indeed a particularly attractive goal for Shor’s (discrete logarithm) algorithm—at comparable security level (cf. [Gir17], which includes the BSI recommendations) they appear to require less resources than factoring an RSA modulus with Shor’s approach. The advances put forward in January 2020 [HJNRS20] further manifest this impression of elliptic curves over prime fields as an attractive quantum cryptanalytic target.

9.7.1 Elliptic curves over prime fields

Applying Shor’s algorithm against an elliptic curve requires in particular the implementation of arithmetic on this curve. According to Proos and Zalka [PZ03], it is not necessary to implement the complete addition law on the elliptic curve, and it suffices to restrict to implement a generic case of a point addition—doubling and adding the inverse can be ignored. Roetteler et al. [RNSL17c] adopt this saving technique. An alternative would be to consider complete addition laws, but for odd characteristic no quantum circuits for such an approach appear to be available in the literature at this point. To implement the pertinent prime field arithmetic, Roetteler et al. invoke Montgomery multiplication and the extended binary Euclidean algorithm as a quantum circuit—not surprisingly, the inversion operation requires particular care, as the running time of a straightforward Euclidean algorithm depends on the inputs. The elliptic curve arithmetic itself relies on the familiar affine representation with a short Weierstraß equation $y^2 = x^3 + ax + b$. This may at first glance

seem surprising, but it is indeed not obvious that some version of projective coordinates would be more efficient when having to express the curve arithmetic as a quantum circuit. Roetteler et al. explicitly point out the question of handling temporary variables, which are commonly used in classic approaches to efficient elliptic curve arithmetic, and taking care of the necessary uncomputations. The most recent work [HNRSJ20] on discrete logarithm calculation on elliptic curves over prime fields, again relies on a short Weierstraß equation.

A completely exact resource count would have to take the bit structure of the underlying prime field (and the constants defining the curve) into account, but based on the discussion in [RNSL17c] it seems reasonable to assume that the variation caused by this is not really significant. Overall, [RNSL17c] obtain the circuit characteristics shown in Table 9.2 for an elliptic curve over a prime field $\text{GF}(p)$ with $n = \lceil \log_2(p) \rceil$, which are backed by simulation results in software. The table given here already takes into account a correction from [Roe17] for the case $n = 160$. The approach in [HNRSJ20] leverages AND gates and measurement-based uncomputation, and various trade-offs are offered. The number of qubits in their low-width designs is comparable to the values in Table 9.2. Specifically, for a 256-bit modulus a design with 2124 qubits, for a 384-bit modulus a design with 3151 qubits, and for a 521-bit modulus a design with 4258 qubits is reported. However, as shown in [HNRSJ20, Table 1] other trade-offs are possible, which at the expense of additional qubits reduce the circuit depth. For a fair comparison, it should be pointed out, that unlike [RNSL17c], in [HNRSJ20] measurements are part of the algorithm – e.g., for a 256-bit curve with the above-mentioned 2124 qubit implementation, about 1.76×2^{26} measurements are involved.

In general, [RNSL17c] offers an upper bound of about $9n + 2\lceil \log_2(n) \rceil + 10$ qubits and $448n^3 \cdot (\log_2(n) + 4090)$ Toffoli gates to be sufficient to implement Shor’s algorithm. Additionally, $8n^2$ T -gates are required for small rotations, and at least $290n^3 \log_2(n)$ CNOT and $71n^3 \log_2(n)$ NOT gates [RNSL17c]. The results in [Eke18] do not change the qualitative picture – the number of qubits is not affected, but the number of gates and depth may be halved.

Using a circuit from [AMMR13, fig. 7a], each Toffoli gate translates into a circuit on three qubits comprised of seven T - (resp. T^\dagger) gates plus two Hadamard and six CNOT gates with a T -depth of 4. An alternative circuit [AMMR13, fig. 13] reduces the T -depth to 3, using one more CNOT and a slightly larger overall depth. When implementing in the surface code in a time-optimized manner, the latter is the preferable circuit since the computation time only depends on the T -depth, and an additional CNOT does not produce much overhead compared to the T -gates (see Chapter 7, or section 7.2.5 for time optimal computation). Based on [HNRSJ20, Table 1], an alternative approach with measurement-based uncomputation (rather than a direct decomposition of Toffolis) can reduce the T -gate count by more than two magnitudes, but this comes at the cost of introducing numerous measurements.

Remarks

- Roetteler et al. exploit a qubit-saving technique which avoids a dedicated Quantum Fourier Transform step in the last part of Shor’s algorithm. This is different from the “textbook description” of Shor’s algorithm. Instead of the QFT step at the end of the algorithm, $2n$ (single qubit) measurements are conducted *during* the execution of Shor’s algorithm and phase shift gates that are to be implemented depend on the outcome of these measurements.
- Earlier work by Proos and Zalka [PZ03] suggests that resource savings are possible compared to the above approach. No actual simulation results for the latter approach have been reported so far, neither have exact Clifford+ T counts been documented. However, if implemented as predicted, the number of qubits could potentially be reduced to about $5n + 8n^{1/2} + 4 \log_2(n)$. Proos and Zalka’s more optimistic estimate suggests that for a 256-bit curve, 1500 logical qubits and a Toffoli depth of about $1.8 \cdot 10^{10}$ suffice. For a 512-bit curve, Proos and Zalka’s more optimistic estimate suggests 2800 logical qubits and a Toffoli depth of about $1.5 \cdot 10^{11}$ to be sufficient.

n	Number of qubits	Number of Toffoli gates	Toffoli depth
160	1466	$2.97 \cdot 10^{10}$	$2.73 \cdot 10^{10}$
224	2042	$8.43 \cdot 10^{10}$	$7.73 \cdot 10^{10}$
256	2330	$1.26 \cdot 10^{11}$	$1.16 \cdot 10^{11}$
384	3484	$4.52 \cdot 10^{11}$	$4.15 \cdot 10^{11}$
521	4719	$1.14 \cdot 10^{12}$	$1.05 \cdot 10^{12}$

Table 9.2: Toffoli gate counts for a dlog computation over an elliptic curve over a prime field $\text{GF}(p)$ with $n = \lceil \log_2(p) \rceil$, according to [RNSL17c, Table 2], [Roe17]. With recent work by Ekerå [Eke18] one may hope to halve the gate and depth counts, the number of qubits is not affected.

9.7.2 Factoring an RSA modulus

Similar as for the case of a discrete logarithm, it is possible to avoid the QFT step at the end of Shor's algorithm and replace it with (single qubit) measurements along the way (cf. [Bea03]), and the critical operation that needs to be implemented as a quantum circuit is a modular multiplication with a constant – the modulus being the number to be factored (say an RSA modulus). Again, various choices are possible how to implement the pertinent arithmetic. Arguably, currently the most practical proposed circuits are due to Häner et al. [HRS17] and the approach by GidneyEkerå and in [GiEk19], using approximate arithmetic. Similarly as in the above-discussed elliptic curve discrete logarithm, Häner et al.'s arithmetic builds on a (classical) reversible Toffoli network. Each of these Toffoli gates can be decomposed into Clifford and T -gates without introducing additional synthesis overhead. From [HRS17,RNSL17c] we obtain the resource estimates shown in Table 9.3, where n denotes the bit length of the number to be factored.

n	Number of qubits	Number of Toffoli gates
1024	2050	$5.81 \cdot 10^{11}$
2048	4098	$5.20 \cdot 10^{12}$
3072	6146	$1.86 \cdot 10^{13}$
7680	15362	$3.30 \cdot 10^{14}$
15360	30722	$2.87 \cdot 10^{15}$

Table 9.3: Toffoli gate counts for factoring an n -bit number according to [HRS17], [RNSL17c, Table 2]. With work by Ekerå and Håstad [EH17] one may hope to reduce the gate count by a small factor (4), the number of qubits is not affected.

Similar to the previous case, there is only few parallelization in the circuit, so that the Toffoli depth is in the same order of magnitude as the total number of Toffoli gates.

In general, [HRS17, RNSL17c] suggests that for factoring an n -bit RSA modulus, $2n + 2$ qubits and $n^3 \cdot (64 \cdot (\log_2(n) - 2) + 29.46)$ Toffoli gates suffice. Following [GiEk19], the number of logical qubits is higher (scaling with $3n + 0.002n \times \log n$ instead of $2n + 2$ in the bit-length n), but the Toffoli count is lower – the dominating term is $0.0005n^3 \log n$ (see [GiEk19, Table I]). In terms of the number of qubits, Proos and Zalka [PZ03] also estimate around $2n$ qubits, but they offer a slightly more optimistic estimate for the number of Toffoli (and therewith T) gates. For factoring a 3072 bit number, they expect only $3.6 \cdot 10^{11}$ Toffoli gates. The results in [EH17] do not change this picture drastically – the number of qubits is not affected, and the number of gates may be reduced by approximately a factor 4.

9.7.3 Discrete logarithms in $GF(p)^*$

Interestingly, for discrete logarithms in $GF(p)^*$, there appears to be no detailed analysis of Shor’s discrete logarithm available in the literature predating [GiEk19]. The essential change compared to the elliptic curve situation is the pertinent group arithmetic that needs to be implemented. As group elements are larger, we would accordingly invest more qubits to represent group elements, and again modular arithmetic would need to be implemented. As the basic structure of the algorithm itself would not change, taking the estimates for factoring a modulus of comparable bit size as p appears a reasonable lower bound on the required complexity (cf. Table 9.3). Indeed, the estimates provided in [GiEk19, Table V] – which include overhead for fault-tolerance – are comparable to the cost of factoring, e.g., for finding a discrete logarithm with a 2048-bit modulus (a safe prime) with Shor’s algorithm, the use of 26 million qubits is estimated for one (seven hour) run of the algorithm. For comparison: To factor a 2048-bit RSA modulus, [GiEk19] estimate that about 20 million qubits – including the overhead for fault tolerance – are used for one (5.1 hour) run. With a (more traditional) Toffoli-based arithmetic, we obtain an estimate of about one million qubits as being needed to factor 2048-bit RSA in 100 days, taking fault-tolerance into account.

9.8 Translating algorithmic gate counts into fault-tolerant building blocks

The first step in calculating the physical qubit overhead of a surface code implementation can be done hardware independent. For the algorithms given in chapter 9, we use the instructions of section 7.5 to calculate the number of fundamental space-time building blocks N_b of the Clifford-part (i.e., everything except of magic state injection and distillation) and the target logical error rate per block P_b that is required to perform the whole algorithm with maintainable success. Furthermore, we extract the number of total and sequential T gates N_T and $N_{T,s}$ respectively and calculate the target error rate per distilled magic (T gate) state P_T . Combining the errors of both distillation and Clifford circuits, we reach an overall success probability of at least 13.5%. Table 9.4 and Table 9.5 show all logical block and gate counts and their corresponding target logical error rates.

n	N_b	P_b	N_T	P_T	$N_{T,s}$
160	$3 \cdot 10^{13}$	$3.3 \cdot 10^{-14}$	$2.1 \cdot 10^{11}$	$4.8 \cdot 10^{-12}$	$8.2 \cdot 10^{10}$
224	$8.4 \cdot 10^{13}$	$1.2 \cdot 10^{-14}$	$5.9 \cdot 10^{11}$	$1.7 \cdot 10^{-12}$	$2.3 \cdot 10^{11}$
256	$1.3 \cdot 10^{14}$	$7.7 \cdot 10^{-15}$	$8.8 \cdot 10^{11}$	$1.1 \cdot 10^{-12}$	$3.5 \cdot 10^{11}$
384	$4.5 \cdot 10^{14}$	$2.2 \cdot 10^{-15}$	$3.2 \cdot 10^{12}$	$3.1 \cdot 10^{-13}$	$1.2 \cdot 10^{12}$
521	$1.1 \cdot 10^{15}$	$9.1 \cdot 10^{-16}$	$7.7 \cdot 10^{12}$	$1.3 \cdot 10^{-13}$	$3.2 \cdot 10^{12}$

Table 9.4: Logical gate counts for elliptic curve Shor attack using dlog. With recent work by Ekerå [Eke18] one may hope to halve N_b , N_T , and $N_{T,s}$.

n	N_b	P_b	N_T	P_T	$N_{T,s}$
1024	$6.2 \cdot 10^{14}$	$1.6 \cdot 10^{-15}$	$4.1 \cdot 10^{12}$	$2.4 \cdot 10^{-13}$	$1.6 \cdot 10^{12}$
2048	$5.5 \cdot 10^{15}$	$1.8 \cdot 10^{-16}$	$3.6 \cdot 10^{13}$	$2.8 \cdot 10^{-14}$	$1.4 \cdot 10^{13}$
3072	$2 \cdot 10^{16}$	$5 \cdot 10^{-17}$	$1.3 \cdot 10^{14}$	$7.7 \cdot 10^{-15}$	$5 \cdot 10^{13}$
7680	$3.5 \cdot 10^{17}$	$2.9 \cdot 10^{-18}$	$2.3 \cdot 10^{15}$	$4.3 \cdot 10^{-16}$	$8.4 \cdot 10^{14}$
15360	$3.1 \cdot 10^{18}$	$3.2 \cdot 10^{-19}$	$2 \cdot 10^{16}$	$5 \cdot 10^{-17}$	$7.7 \cdot 10^{15}$

Table 9.5: Logical gate counts for factoring an n-bit number. The Toffoli depth has been approximated for low parallelization similar to the dlog algorithm. "With work by Ekerå and Håstad [EH17] one may hope to reduce N_b , N_T , and $N_{T,s}$ by a small factor (4).




Figure 9.1: Physical volume required for performing dlog (blue) or factoring (orange) algorithms of different problem sizes (numbers on the right) as a function of the physical error rate. We use a threshold of $p_{th} = 10^{-2}$ and assume a symmetric error model (all kinds of errors equally likely) with equal rates p for initialization, gates, waiting and readout, which makes this plot hardware-agnostic and approximate. A surface code cycle consists of multiple gates, measurement, initialization, and classical processing.

For a general outlook on the required overhead for different physical error rates, we use the steps described in chapter 7 for the surface code to calculate the total volume of an algorithm, i.e., number of physical qubits times the number of surface code cycles (or alternatively the time required for running them). Figure 9.1 shows this volume as a function of the initial physical error rate for several problems. With this, one can get the number of qubits by fixing the computation time to some chosen value or get the runtime by fixing the

number of qubits. The actual time per surface code cycle is given by the time for initialization, readout, two Hadamard and four CNOT gates. The only restriction in balancing time and qubit overhead is the minimal time required to preserve the temporal order of non-Clifford T gates given by the T depth of a circuit. Since the most time-demanding process in the classical feed-forward required by the T gates is one measurement, multiple subsequent T gates can be performed during one surface code cycle, a typical assumption is $t_M = 0.1t_{SC}$. The minimal number of surface code cycles for this assumption is shown in Table 9.6. Usually, this is also the optimal time to choose, since time-scales are still large with current architectures.

	dlog		factoring				
n	160	224	1024	2048	3072	7680	15360
# SC cycles	$8.2 \cdot 10^9$	$2.3 \cdot 10^{10}$	$1.6 \cdot 10^{11}$	$1.4 \cdot 10^{12}$	$5 \cdot 10^{12}$	$8.4 \cdot 10^{13}$	$7.7 \cdot 10^{14}$

Table 9.6: Minimal number of surface-code cycles for performing dlog or factoring algorithms with problem sizes n with one surface code cycle typically corresponding to a time of $\sim 10t_M$. Thus, during one surface code cycle one can run 10 subsequent T gates.

10 Risks of our evaluation scheme

Our evaluation scheme rests on the status of current research and knowledge. Some of these results are extrapolations over many orders of magnitude in size and performance, specifically error rate. We would like to succinctly describe the known risks that this scheme could be wrong, which can only be assessed as research, mostly experimental research, progresses.

10.1 Risks that make quantum computers more reachable

1. We have assumed that cryptanalysis requires long skinny algorithms hence requiring error correction. Discovery of an algorithm that trades time for memory in a way that can be addressed with a small number of gates would make the target processor much smaller. This risk is low, as the latter requirement seems to be exclusive to the simulation of quantum systems with quantum computers.
2. Discovery of physical qubits with extremely low intrinsic error rate. In principle, this is possible—control of qubits can be done with non-dissipative elements that do not produce errors. One candidate for this are topologically protected quantum bits, which currently are sitting on level A as they have not demonstrated two-qubit gates. Now even optimistic analyses [VF17] point out that these qubits will not be error free or have error rates that are low enough to avoid error correction, they just need smaller overhead. Thus, while this could happen, the correction to our scheme would be quantitative only.
3. Discovery of scalable qubits with long-distance interaction with the ability to implement high-dimensional connectivity: This would lead to very high error thresholds and at least logarithmic savings (cf. Section 7.4.3). It is however unlikely that the required parallelism can be reached.
4. Discovery of accidental error avoidance in cryptographically relevant algorithms. This is related to the fact that error estimates following the diamond norm are usually very conservative. Some type of this error avoidance are reported for the QFT-portion of Shor's algorithm [NB15b], but a thorough analysis of error tolerance of the modular exponentiation respectively the underlying group arithmetic is not available (cf. Section 4.1). Current literature focuses on implementing arithmetic efficiently if the problem of errors on the logical level is taken care of by error correction. If some implementation strategy for the arithmetic portion of Shor's approach would turn out to be inherently error-tolerant, this could in principle reduce the error correction overhead. Assuming that Shor's algorithm can be implemented without any error correction seems highly optimistic, but with the current state of the literature it cannot be ruled out that savings in the arithmetic portion of Shor are possible by simply tolerating some imperfections.
5. Implementation of novel, ultra-fast quantum computing platforms in the timescale of fs or as, the shortest directly accessible timescales in physics. This would speed-up physical gate times by three orders of magnitude, making long algorithms more accessible. This has been tried, unsuccessfully, and it is not clear whether that speed advantage would translate into faster qubits.
6. Algorithmic innovations and optimization of the logical encoding: Finding the optimal encoding (see [Jon13] for example Section 7.2.6) and distillation structure (Section 7.2.4.3) is a task that we cannot do in full detail in our analysis. While we take reasonable assumptions for required distance, distillation rounds or logical gate arrangement, an optimized version of a fault-tolerant algorithm found from simulations can be made much more efficient in terms of required qubits and error rates (see for example recent advances in [OC17], or different approaches to fault-tolerant Toffoli gate implementations [Jon13]). These optimizations will be done for sure when thinking about implementing large circuits, the correction will be a constant factor improvement (of maybe one or two orders of magnitude).
7. Extreme progress in error correction with transversal T -gates: Very unlikely.
8. True scaling advantage of the surface code when using lattice surgery, see [FG18] and our section 7.2.5.

10.2 Risks that make quantum computers less reachable

1. Serious deviations in going from levels B to C: Even the best methods to measure operation fidelities on level B reveal an incomplete picture. As there are very few demonstrations of active error correction so far [WL17, OPH⁺16, SBM⁺11, KBF⁺15], and the results are too inconclusive to say whether the expected error correction performance matches the expectations, this has intermediate risk.
2. Discovery of new correlated error mechanisms: Error correction relies on multi-qubit errors being exponentially (in the number of qubits) less likely than single-qubit errors. We have argued that given the two-body nature of interaction, these are extremely unlikely. Now one can view operating a fault tolerant quantum computer with heavy error correction as a precision test for the error model over an enormous scale and novel error mechanisms could creep up. This would be physics beyond the standard model of elementary particle physics (which famously does not contain gravity). Now given the wide success of the standard model and its enormous predictive power at high energies makes it hard to assign a likelihood. If this happened, thresholds of error correction would need to be questioned.
3. Discovery of persistent non-Markovianity: Similar to spatial correlations also temporal error correlations are difficult to catch. This is unlikely, as measurements usually destroy temporal correlations.
4. Insurmountable engineering problems: Assembling large processors cannot guarantee the same quality as the components. The same would hold for temporal stability when scaling operation time, e.g., spurious heating and drifts. Albeit analyzing these operational challenges is done based on what is known for the level C platforms, there can be challenges that only appear while it is attempted.
5. Dominance of coherent errors: Albeit coherent errors have the same error correction threshold as corresponding incoherent errors, the surface code scales less favorably below threshold, which may increase the overhead. Intermediate risk.

11 Global operational criteria for quantum computers

While still elusive, quantum computing research is far enough advanced to project and speculate about operational criteria and requirements for a scaled-up machine. This can be driven by experience gained from classical (super-) computers as well as from the bottom-up operational challenges that were collected in the previous chapter. These criteria are separate from the mid-level requirement for operating quantum error correction. We would like to introduce three classes of operational issues:

- **Extensive parameters:** Numbers characterizing a quantum computer that grow roughly linear with the size of the machine.
- **Critical parameters:** Challenges that need to be overcome in scaling. Operationally, these challenges become critical faster than linearly when scaled.

Further descriptors: Here we look at descriptors that are not growing dramatically and are not critical, but that characterize the operation of a quantum computer and its suitable environment.

11.1 Extensive parameters

An attractive way to project operation is to ask how much of a given resource is required per qubit. This is not easy to answer and the answer should leave space for technological progress. In particular the last two layers are expected to develop dramatically with the entry of industry and engineering in the field—current setups from research laboratories are optimized for flexibility of experimentation, not integration.

11.1.1 Scales of extensive parameters

As a preliminary consideration, it seems necessary to measure the effort per qubit on four different scales.

The bare qubit When choosing a platform, a degree of freedom to encode a qubit, there is a certain scale that even most imaginative engineering cannot overcome, posing a fundamental (and often unreachable) bound for the quantum computer. This is the average diameter of a single ion or atom, the nuclear radius in NMR, the size of a superconducting qubit, the size of a quantum dot are bare qubit scales.

The unit cell Even when building a simple quantum register, a qubit does not go alone. It needs to be addressed by controls (if controls are local, this requires space) that reduce crosstalk (the effect of controls aiming at a specific qubit also affecting other units), it potentially needs to be held in place by external fields. This needs to be treated differently from the bare qubit as different design choices lead to different unit cells. The 3D transmon is, e.g., very small, but needs to be operated in a large machined cavity, different from the planar transmon which thus can be packed much more densely. Unit-cell limits are described in 14.1.3.4, 14.1.3.1, 17.1.6.1.

The periphery layer The controls and read-out attached to qubit unit cells need to be externally connected to electronic and optical elements. They are typically operated under less demanding conditions than the qubits hence requiring different resources. Also, these elements are often shared between platforms - for example microwave electronics for spins in semiconductors and superconducting qubits or laser systems for neutral atoms and ions. Examples are described in 13.1.3.4, 13.1.3.5, 17.1.6.2.

Infrastructure Given the strong need for protection of quantum computing systems, they sit in some type of infrastructure providing suitable operating conditions. For semi- and superconductors these involve cryogenics, for atomic systems they include high vacuum and vibration-isolated optical tables. Here, a critical point occurs when the size of infrastructure units is not sufficient and multiple units need to be connected. Infrastructure challenges are described in 13.1.3.4, 14.1.3.2, 14.1.3.3, 17.1.6.1.

11.1.2 Size

An obvious extensive parameter is the size of a qubit. While usually this is a volume, unit cells can also be quasi 1-dimensional (as in the linear Paul trap), 2-dimensional (as in chip-based superconducting and semiconducting circuits) or also 3-dimensional. In most cases, size is dominated by the periphery layer. In solid-state qubits, the long wavelength of microwaves makes microwave elements centimeter-size, in atomic systems, optics miniaturization is a major challenge.

11.1.3 Power consumption

While quantum computing is reversible, a lot of operations around it are not, so the power consumption per qubit is an issue. Currently, the high-power elements are cryogenics and lasers, which belong to the infrastructure layer and probably do not increase with the number of qubits until reaching the threshold of the need for multiple infrastructure units.

11.1.4 Power dissipation and temperature stability

Next to the power being consumed, it is a separate question where this power is dissipated and how that influences the thermal management of the system (and how much cryogenics are required). Unwanted heating, e.g., plagues semiconductors, where the effective electron temperature is often ten times the temperature of the cryostat. One needs to discriminate power consumption by the means of cooling:

6. room temperature: A/C system
7. 77 K: Liquid Nitrogen
8. 4 K: Liquid He
9. 1 K: Liquid ^3He
10. 10 mK: ^3He - ^4He mixture

and calibrate the amount of coolant needed. Note that some of these coolants (specifically nitrogen) are typically consumed during cooling whereas Helium can be preserved in a closed cycle.

11.1.5 Cycle time

Again, one would like to know the clock speed of the quantum computer as given by its slowest ingredient. This critically depends on the technology being used - it often is believed to be the two-qubit gate but in practice it often is qubit reset, that even takes a full relaxation time or extra overhead for classical reset.

11.1.6 Classical data flow

This issue is most pressing in connection to error correction: As the code and control layer of a quantum processor are classical, one is faced with the need to process data fast and close to the device in a way that grows with computer size. In particular on low-level, this is done with cryogenic electronics, which impacts periphery space and power dissipation.

11.1.7 Reliance on rare materials

Some qubit systems are based on rare materials on some layer. For example, isotopically purified Si without nuclear spins is generated in a laborious process from natural Si. A critical ingredient is ^3He that is needed to reach low temperatures. With reports on the shortage of natural ^4He probably exaggerated, the non-natural

^3He is in short supply already. It has been generated as a by-product of nuclear warfare, specifically hydrogen bombs containing tritium. The little that is generated from US nuclear stockpiles is used for radiation detectors by the US, making ^3He unavailable and prohibitively expensive. A vast quantum computer based on ^3He would likely require a designated nuclear source for ^3He .

11.1.8 Vacuum

Some qubits need to be operated under ultrahigh vacuum. Trapped ions, e.g., use their motional degree of freedom for quantum gates which is at odds with collisions with gas molecules. Given outgassing of materials, one needs to ask to what point vacuum infrastructure can be enlarged.

11.1.9 Production speed

Computers can be scaled based on mass-production. With extreme technology as quantum computers currently made under research conditions, this needs to be addressed. Notable challenges include the enormously long production time for cavities in neutral-atom cavity QED, the difficulty in designated single-dopant implementation in dopant spins in semiconductors. At some scale, also the different speed in mask-production in (parallel) optical lithography as used by the current classical semiconductor industry versus (serial) electron-beam lithography used to make nanostructures needs to be taken into consideration.

11.2 Critical parameters

There are a number of parameters that are mere inconveniences in small laboratory scale systems but that can become prohibitive when scaling up.

11.2.1 Stability

How long can a quantum computer be operated before it needs to be reset/recalibrated? This can be based on effects like the loss of qubits - weakly bound neutral atoms in optical lattices tend to disappear after some time. This can also be due to slow drifts in parameters that occur in imperfectly thermalized systems - it is known that some parameters of Josephson junctions drift on the scale of the day. A crucial example is described in 16.1.2.3.

Unless accommodated in error correcting codes suitable for these problems, these issues can be lethal: Losing a qubit with probability p per unit time means losing a qubit with probability $1-(1-p)^N \simeq Np$ in a large quantum computer per unit time, effectively limiting algorithm run-time to $(Np)^{-1}$ time units.

11.2.2 Yield and scatter

On a level lower than instability, one needs to make sure that the production of a quantum computer is reliable: Are all devices close to their design parameters, are all of them performing on a sufficient level? Can good/bad devices be selected before systems integration? How does this limit the size of a module? For example when qubits cannot be locally controlled but need to be addressed by frequency selection, it is necessary to produce them at the right frequency—is that reliable? This plays a role, e.g., in superconducting flux qubits, where some operation parameters depends doubly exponentially on hard to control fabrication parameters, see also 13.1.2.2.

11.3 Further descriptors

Albeit the previous criteria seem to control most of the deciding criteria to operate a quantum computer, it makes sense to reflect on the overall device. How big is it? Can it be operated overground or in a tall building (i.e., does it need to be insulated from vibrations)? Can it be operated by general data center staff?

11.4 Addressing these operational challenges

These operational challenges are embraced by the most mature quantum computing platforms where basic elements are working well enough so device-level scaling can be addressed. These go in most cases beyond the capabilities of a single university-level research group but need different actors. Typical approaches include:

- 1 Large commercial organizations can address these engineering and scaling challenges, e.g. IBM, Google, D-Wave Systems, MIT Lincoln Laboratories, Rigetti (13.1.5.1)
- 2 Specialized companies can address focused engineering challenges. These can be
 - a dedicated spin-offs from research groups, such as ColdQuanta 16.1.5
 - b research equipment manufacturers that get funded to adapt their products, such as Zurich Instruments (measurement technology, see 13.1.5.2), Oxford Instruments (Cryogenics), or Toptica (Lasers)
 - c specialized engineering firms addressing a few specific challenges, such as BBN (pulse generators, 13.1.5.2), Northrup-Grumman (control electronics, 13.1.5.2), Intel (semiconductor manufacturing)
- 3 Research clusters involving engineering scholars or public engineering institutions enhance the capabilities of quantum researchers, e.g. at QuTech in Delft (15.1.4), the semiconductor growers (15.1.4.1), Sandia national labs (ion traps, 14.1.4.1)

The engineering leaders are mentioned along with the platforms.

11.5 Articulated architectural extrapolations

With the increasing maturity of platforms, extrapolating full quantum computer architectures has become a reasonable proposition - what in fact this study addresses in other chapters. Some platforms have been very early in this - the Kane quantum computer [Kan98] - and needed to correct their assumptions about experimental capabilities. Others are rather hesitant - such as the Josephson qubit community which did not want to repeat the experiences of the classical supercomputing community of the 1980s and 1990s. In this period, rapid single flux quantum (RSFQ) was promised to revolutionize classical supercomputing by allowing for fast clock speeds - which never materialized due to challenges in fabrication, large element size requiring many clock cycles for communication, and supercomputing demands requiring parallelism rather than clock speed. Also, we expect that confidential studies exist, in the form of detailed research proposals and in the form of company strategies. Very recently, two of these studies came out of the IARPA-LogiQ program, see 20.1.1, which has corresponding milestones—and more of these could come out of the remaining LogiQ teams (see again 20.1.1 with two superconducting qubit teams, IBM and TU Delft, and two ion trap teams, Duke and Innsbruck).

A blueprint for a largely microwave based ion-trap quantum computer has been published recently [LWF⁺17]. It argues that this is possible given that there is no prohibitive challenge of laser adjustment. It comes to the conclusion that performing a 2048-bit number Shor factorization will take on the order of 110 days and require a system size of $2 \cdot 10^9$ trapped ions. Shor factoring of a 1024-bit number will take on the order of 14 days. They will require almost the same amount of physical qubits because the required pace of the ancilla qubit generation is the same for a 2048-bit and a 1024-bit factorization. Trapping $2 \cdot 10^9$ ions will

require 23×23 vacuum chambers occupying an area of ca. $103.5 \times 103.5 \text{ m}^2$. Its most surprising result is the power consumption for the surface traps that are made out normal conducting metal, leading to a power consumption of about 1000 W per module, of which this processor requires about 5000, leading to a 5 Megawatt power consumption, which is less than a present-day supercomputer. The paper points on routes with better gates to bring down these numbers.

A competing analysis of the more standard optical ion trap architecture was posted in May [BXN⁺17].

Other than the first one mentioned, this paper is much less concrete in its conclusions. It is a more detailed version of much of our ion trap chapter. Most notable is its analysis of the color code (rather than the surface code) for trapped ions. The color code has a somewhat lower threshold than the surface code and requires longer parity measurements (which are well adapted to the Molmer-Sorensen gate), but it is less complicated to compute on it. The authors also state their expectations for next-generation gate errors, which are all expected to improve by more of an order of magnitude and land somewhat above 10^{-4} .

The book chapter [DSMN16] thoroughly analyzes a photonics quantum computing architecture that uses atoms as nonlinearity and proposes concrete modules for achieving that within topological error correction for cluster states, which is its centerpiece. It estimates an error threshold of around 0.6% but does not extrapolate overhead from the performance of this platform as it is very far above that error rate.

What is the significance of these extrapolations? First of all, the fact that they can be made and have a finite result allows to gauge the distance to a viable architecture and to identify the most mission critical developments. As far as technology is concerned, they are optimistic that some quantitative progress can be made and that scaling up has no unpleasant surprises (e.g. that device performance is not affected by integration in a large machine), but also pessimistic as they do not anticipate breakthroughs that will change the field. So they are probably limited predictors of the *science* of building a quantum computer, but good guides to *engineering* it.

Part II: Platforms

12 Quantum technology and computing platforms

This part summarizes physical platforms for quantum computing as well as algorithms that are relevant for cryptography. It identifies main actors and criteria for the successful operation of a quantum computer.

The field of physical realizations is evolving at quick pace and is producing an ocean of literature as well as large conferences as a result. Also, given the attention quantum technologies in general and quantum computing in particular receive currently, there is a lot of noise created including some research fields trying to relabel themselves. So doing a survey is an impossible task if it is not guided by some clear principles. Here is what we have applied:

1. A serious quantum computing candidate needs to have at least one experimental activity associated with it that identifies the potential existence of a qubit, i.e., check at least the first DiVincenzo criterion . Pure theoretical considerations without any experimental activity would not make the cut.
2. On the same token, an experiment that is not linked to any theoretical proposal how to meet the DiVincenzo criteria in their simplest, non-quantitative form at least in principle, would also be discarded.
3. Quantum technologies are classified in four broad categories originally formulated for the EU [QE16] but now more broadly accepted: sensing and metrology, communication, simulation, and computing. Clearly, this work singles out computing. The dividing line to simulation is not razor-sharp, specifically in the area of digital quantum simulation one essentially runs a specialized quantum algorithm. Also, the EU classifies some approaches, notably quantum annealing, as simulation, even though they can have some cryptographic relevance, so annealing is included here. When describing platforms, we focus on their quantum *computing* aspects, not on the others - for example we do not describe the application of NV-centers in diamond in sensing of magnetic fields and forces, neither do we describe the latest atomic clock technology based on trapped ions.
4. A large part of quantum computing is driven by (in the words of Andrew Steane [Ste03, BKCD02]) climbing *Mount Scalable*—being governed by quantum error correction in needs to improve operation and include more and more qubits. We have tacitly assumed that qubits will ultimately based on an error correcting code . This field is dominated by the topology-inspired surface code [FFSG09, FMCM12], as it provides high error thresholds and only requires nearest-neighbor qubit-qubit coupling. For the sake of this survey, error correction could also be done by other popular codes such as color codes or concatenated *Calderbank-Shor-Steane* (CSS) codes. It needs to be noted that these codes are assembled from physical qubits that are separate functional units - they are agnostic to the type of qubit used as long as their requirements are met. A few exceptions are mentioned explicitly: i) It is believed that topological qubits—qubits that intrinsically, in their microscopic physics realize topological protection—need much less error correction. ii) In quantum annealing, the role of error correction is under debate—some of its proponents highlight that an intermediate amount of classical noise may actually be beneficial. iii) Schemes like cluster state quantum computing are currently in their infancy and connecting them to error correction is certainly a long-shot.
5. The plurality of platforms has some resemblance to the early classical computer age, where implementations went from mechanics to electromechanics, vacuum tubes, solid-state transistors all the way to integrated circuits. This is clearly the stage of the field right now. Even more: Some platforms are quite pluralistic internally, for example semiconductors, whereas others that are more mature are at the same time more focused, such as ion traps. So chapter length is a poor indicator for the quality of a platform.

The description of platforms is driven by sorting and categorizing at least as much as by finding them all. We have ordered them by what we believe (in no strong contradiction with the rest of the community) an order of decreasing potential. This latter ranking has to be taken as preliminary and is based on the European quantum technology roadmap as well as funding priorities. In the end of each chapter we rank the platform in terms of the evaluation system from Part I. We also needed to make a few deliberate choices that could have been done otherwise: NV centers in diamond are part of the semiconductor chapter even

though they are sometimes called nature's trapped ions and borrow a lot of ideas from ion traps. Also, topological qubits did not get their own chapter but were summarized under semiconductor qubits—certainly their currently most advanced candidate is crucially based on a semiconductor. This leads to the seemingly paradox situation of two superconductor-only platforms appearing in the semiconductor chapter, but they form an internally logical group.

We have listed known world-records in gate times, fidelities, and coherence. This is a fast moving target and got updated during the study. Ultimately, the most mature platforms are better characterized by operation fidelities and compatibility with error correction.

Quantum computing continues to have a strong impact on algorithm design, and cryptanalysis remains one of the most prominent potential application domains of quantum computing. It is interesting to observe that more than twenty years after Grover's algorithm and Shor's algorithms have been published, the research community is still working on analyzing their quantitative impact on the cryptanalytic landscape. There are commonly two different aspects to consider when looking at quantum algorithmic innovations in cryptanalysis. First, there is a "true quantum" portion of the algorithmic innovation, which can allow a (sometimes dramatic) speed-up over classical solutions, *assuming appropriate access* to the classical problem description. Second, the classical problem description—which can include things like the arithmetic in a particular cyclic group or the details of a hash function for which a collision is sought—influences the exact operations that need to be mapped to the quantum hardware. The latter aspect can to a wide extent (though not entirely, as elementary operations and cost measures will usually differ) be discussed within a framework of classical reversible computing. Notwithstanding this, the complexity of this "classical portion" is critical for, and can in fact dominate, the overall running time of a quantum attack. This part focuses on a qualitative discussion of pertinent quantum cryptanalytic algorithms, and it should become clear that the current implications of quantum computing for symmetric and asymmetric cryptography are very different. A more quantitative discussion is to be given in Work Package 4.

Operational criteria and "scale up science" are much harder to come by, given the tendency of humans to report successes rather than challenges. They were mostly identified bottom-up but then summarized in the end at a workable executive level. Also, there is no shortage of actors and research groups so our main task was to sort them.

13 Josephson qubits

13.1 Description

Superconducting integrated circuits based on Josephson Junctions are a solid-state based platform of quantum bits. They are viewed as one of the leading realization candidates by the US government and the EU quantum flagship [NIS16, QE16].

This platform is being pursued as a platform for both adiabatic quantum computing / quantum annealing and quantum circuit-based quantum computing. Its basic unit, the qubit, is currently based either on the transmon or the flux qubit design. Coupling and control is mediated through microwave transmission lines, which can also serve as interfaces to other platforms.

This technology is widely pursued by academic, government, and commercial actors.

13.1.1 Basic notions and terminology

Superconductivity, the property of certain metals to conduct electricity without resistance and completely expel magnetic fields, is a macroscopic quantum phenomenon that occurs at low temperatures. In current research on superconducting quantum bits, the materials used are conventional superconductors—elementary metals and alloys—in which superconductivity is well understood. Unconventional superconducting materials including high-temperature materials are currently not pursued for quantum computing in this platform but play a role in topological quantum computing, see respective chapter. The most common materials are Al and Nb which superconducting below $T = 1.2\text{K}$ and $T = 9.3\text{K}$ respectively. Hard superconducting alloys such as NbN and InAs-Al play a minor role. This platform naturally operates at low temperatures and electrical charge there is carried by pairs of elementary charges, Cooper pairs. Superconductivity allows to transmit elementary units of information without losses and is thus important for maintaining quantum coherence.

The key element in these circuits is the Josephson junction. This is a weak link between superconductors made out of non-superconducting material, primarily realized as Josephson tunnel junctions from electrical insulators. Consistent with the need for superconductivity, this is a reactive (non-energy-dissipating) element which, unlike the commonly known capacitor and inductor, is classically nonlinear. This nonlinearity leads to a non-equidistant energy spectrum that is crucial in order to selectively address quantum states as computational (qubit) states.

Contact with these circuits—control and readout as well as inter-qubit coupling—is made using electromagnetic fields in the microwave frequency range. This connects to the more established field of classical superconducting electronics that has both high-speed computing and precision sensing applications.

The combination of low operating temperature and energy scales in the microwave ranges allows for state preparation by cooling.

The main sources of error of operation in this platform originate from decoherence from the condensed environment of this system and the necessity to fabricate complex heterogeneous material systems such as Josephson junctions. It competes with leakage to non-computational states.

13.1.2 Basic qubit layouts

13.1.2.1 Charge qubit derived designs

Charge-qubit derived circuits have been developed from ultra-small circuits showing strong charging effects even with elementary charges, such as single-electron transistors. These devices were originally investigated for metrological applications. As they are known to be very sensitive to charge noise, and as charge noise in manufactured nanostructures is a known problem that so far could not be solved, design evolution in these systems are driven by the need for immunity from slow charge noise. The different evolutionary steps preserve their circuit topology. They are different in the ratio E_J/E_C where E_J/E_C describes the Josephson coupling energy for charge exchange and E_C the charging energy of a single Cooper pair. If this ratio is small, the computational states of the system can be well separated by electrical charge, which also means that they are most sensitive to low-frequency charge noise. This number also gives an indication of geometric size: The Josephson coupling is proportional to the junction area whereas the charging energy is proportional to the inverse capacitance (which in turn is proportional to junction area), hence this ratio is approximately proportional to area squared.

Cooper-Pair Box The Cooper pair box has been the pioneering superconducting qubits design [BDGD05, KYG⁺07, MSS00a, MSS00b]. Its E_J/E_C is much smaller than unity (i.e., $E_J/E_C = 0.08 \pm 0.015$ in ref. [BVJ⁺98]). Its resulting sensitivity to charge noise has limited coherence times to tens of nanoseconds. It is not currently pursued in quantum computing architectures.

Quantronium The quantronium has evolved out of the Cooper pair box with a number of innovations to reduce its sensitivity to charge noise: It has $E_J/E_C \approx 1$ (i.e., $E_J/E_C \approx 1.27$ in ref. [VAC⁺02]). This makes its basis states less charge-like, which requires measurement not by charge sensing but by current sensing, and control not by direct charge manipulation but by resonant microwave control. Furthermore, it is operated at an optimal working point [VAC⁺02, CW08] at which slow charge noise has maximum impact. Charge noise is inherently slow, i. e., it fluctuates at low characteristic frequency f -depending on the sample crossover to white noise can be around 1kHz up to 1 MHz. It thus cannot change the energy $E = hf$ of the qubit hence it only induces phase errors which result from slow energy fluctuations. The optimum working point decouples energy from charge in first order hence greatly reducing the ability of charge noise to change the system energies. Introduction of quantronium has led to a more than 100-fold increase of coherence (from $T_1 \approx 6\text{ns}$ [NPT99] to $T_1 \approx 1.8\mu\text{s}$ [VAC⁺02]) It has been pursued by the Quantronics group at CEA Saclay and its spin-off at Yale University. It was abandoned around 2012 in favor of the transmon qubit given its easier systems integration and connectivity.

Planar transmons The transmon [KYG⁺07] pushes E_J/E_C to even larger values (i. e. $E_J/E_C \approx 50$ in ref. [SHK⁺08]) by introducing an external shunt capacitor. This compresses the variability of the energy due to charge even further and leads to near-immunity of charge noise [HKD⁺09]. As quantronium it is biased at an optimum working point and controlled by microwaves either directly or through a microwave cavity (see below). Measurement is performed through a cavity with photons in the microwave frequency range. Depending on its precise connectivity, it is sometimes called a Xmon (shaped as a cross with four connectors) [BKM⁺14, BLK⁺15, BSL⁺16, KBF⁺15], a starmon [VPK⁺16], or a gmon, a transmon with in-situ tunable interaction with neighboring elements [CNR⁺14]. Its originally perceived drawback of only weakly separated computational states has been overcome by optimal control [MGRW09].

The gatemon [dLvHB⁺15, LPK⁺15] replaces the oxide-based Josephson junction by a junction based on a semiconductor nanowire. This offers the additional control knob that the Josephson coupling can be controlled by an external electrostatic gate (rather than by magnetic flux as in conventional Josephson qubits). This offers the advantage of avoiding flux crosstalk by using control voltage—which could be used to pack qubits more densely. While experiments are encouraging, this technique has not been adopted on a large scale.

The planar transmon is the currently most widely used superconducting qubit.

3D-Transmons The three-dimensional transmon is not so much an alternative qubit as it is a way to connect to a transmon qubit. Other than planar technologies, the transmon qubit is encapsulated in a metallic cavity and only accessed through that cavity. This architecture minimizes the participation of lossy oxides and thus has superior coherence properties, but is also much more difficult to control and to scale due to the physical size of the devices. It is considered to be a valid contender for certain applications [PSB⁺11, RBLD12].

13.1.2.2 Flux qubit derived designs

Structures resembling today's flux qubits were first proposed by Nobel Laureate A.J. Leggett [Leg80] as a candidate for testing quantum physics on a macroscopic scale, even before the conception of quantum computing. They have a loop-type circuit topology interrupted by an odd (effective, see below) number of Josephson junctions and their basis states are described by clockwise and anticlockwise circulating current that produce magnetic fields and fluxes (field integrated over area) of opposite directions that can be used for control and measurement. These properties are very close to other, classical platforms of superconducting electronics, the SQUID (Superconducting QUantum Interference Device)—a sensing platform—and SFQ (Single Flux Quantum) ultra-fast classical digital logic. This gives flux qubits superior connectivity and makes them ideal candidates for quantum annealing but also for reaching ultra-strong coupling to microwaves, a property crucial in quantum simulation. A central challenge for flux qubits is reproducibility, as its non-classical behavior is driven by flux tunneling through a barrier of inductive energy. This term is inversely exponential in $(E_J/E_C)^{1/2}$ where E_J itself is inversely exponential in the thickness of the Josephson junction, a hard-to-control parameter. Still, flux qubits are a serious contender for a range of quantum computing applications including adiabatic computing and annealing. Flux qubits are sensitive to flux noise [ASB⁺13, KSB⁺16].

Single-junction loop The simplest flux qubit consists of a loop interrupted by a single Josephson junction. It makes use of a substantial geometric inductance to form a qubit. As this inductance is roughly proportional to wire length, this fixes qubit sizes at a relatively large value, leading to excellent connectivity but also strong impact of external noise. There are very few experiments [FPC⁺00] on coherent manipulations of flux qubits but the single-junction flux qubits is the workhorse of quantum annealing at D-Wave Systems [JAG⁺11, DJA⁺13, LPS⁺14].

Three (active) junction loop This is the most common flux qubit design for quantum circuit applications. It was conceived at TU Delft and MIT [MOL⁺99, OMT⁺99]. It replaces the geometric inductance by additional Josephson junctions hence leading to a much more compact geometric footprint and superior coherence with a wide variability between samples [BGY⁺11]. Their connectivity and their excellent separation of computational to non-computational states allows strong and fast control [OYL⁺05]. It was also the first qubit in which interactions could be tuned in hardware [HRP⁺06].

In order to maximize connectivity, the qubit and its peripheral elements often share lines. Given that Josephson junctions are typically made in a two-layer overlap geometry, this leads to a change of layer when going around the loop twice, similar to a Möbius strip. As this can lead to uncontrolled offsets, sometimes a large-area (hence very passive) fourth junction is inserted in the loop in order to connect the layers in a controlled way [CNHM03, CBS⁺04].

Capacitively shunted flux qubit In order to rely less on the precision of junction fabrication, the capacitively shunted (C-shunted) flux qubits has been investigated. Parallel capacitance can be controlled much better than Josephson junctions and leads to more reproducible qubit parameters with a small sacrifice in energy separation [SKD⁺10]. Not very well developed yet, it is still discussed as a candidate for coherent quantum annealing [CCG⁺11].

Fluxonium The fluxonium circuit is related to the single-junction loop, replacing the geometric inductance by a linear array of about 100 Josephson junctions acting as a “superinductor”. These qubits have superior coherence but are very hard to operate and integrate. We are not aware of two-fluxonium experiments [PGC⁺14, VPS⁺14]. It has been developed further into a new proposal called flatsonium [SRDR17].

13.1.2.3 Phase qubit derived designs

The phase qubit operates in the regime of large E_J/E_C . Other than the flux qubits, its computational states are not classically macroscopically distinct. Phase qubits are very simple consisting of only a single biased Josephson junction as qubit and readout, a setup already investigated in the 1980s to demonstrate macroscopic quantum tunneling [DMC85]. Initially very successful, this qubit turned out to be plagued by defects in the Josephson junction [SLH⁺04, CMB⁺10, LBM⁺16] and have not reached long coherence times. Only very few groups, notably Ray Simmonds at NIST and Alexey Ustinov at KIT use them.

13.1.2.4 Other designs

A number of other designs have been explored but are of mere historic value:

- the fluxon qubit uses internal degrees of freedom of a long Josephson junction. This work culminated in the demonstration of macroscopic quantum tunneling [TM96, WLL⁺03].
- Junctions from high temperature superconductors in place of conventional superconductors could be appealing to operate at higher temperature. Given their difficult materials science and intrinsic damping, only basic quantum tunneling has been demonstrated [TKL⁺04]. Keeping all components coherent still requires low temperatures.

The specific properties of complex high-temperature junctions were speculated to be useful in the early era of Canadian company D-Wave Systems (and are responsible for their name), but were never experimentally implemented [SZW93, Zag97].

13.1.3 Peripheral elements

With the increasing maturity of this platform, connectivity and peripheral elements play a more and more crucial role in identifying these systems along with their operational challenges.

13.1.3.1 Cavities and waveguides

Superconducting qubits are operated at microwave frequencies between 1 and 20 GHz. This range is dictated on the low end by the ability to cool the system to the ground state in a robust dilution cryostat that can reach about 10 mK base temperature but often cools the electronic load only to about 50 mK, and on the high end by the observation that superconductors lose their superconducting properties at frequencies around their energy gap, which for Al as the weakest superconductor that is widely used is around 80 GHz. In order to manage electromagnetic fields at those frequencies, superconducting coplanar waveguides are used, both as transmission lines and as cavity resonators of finite length. Use of these resonators defines the field of circuit quantum electrodynamics (cQED). These resonators possess better coherence than the qubits (but are not controllable without them) and are used to connect qubits to each other over long distance, as well as in some instances for control [BGW⁺07, Poz12].

13.1.3.2 Direct couplers

Interactions between qubits can be mediated by direct coupling elements based on electrical capacitance or inductance. These couplers are used over short range. In principle, as a key advantage of superconducting qubits over other platforms, these couplers can be made tunable in hardware. This tunability has been demonstrated in [HRP⁺06, CNR⁺14]. While feasible, it currently is mostly applied in niches whereas the tuning of interactions in scalable quantum computer platforms is done using resonance methods [SMCG16, BKM⁺14, KBF⁺15].

13.1.3.3 Amplifiers and detectors

Superconducting qubits in principle offer a variety of read-out options. Specifically, much of the underlying technology has originally been developed for magnetic flux sensing using SQUIDs [CB06] which can go up to high frequencies in a microstrip geometry [KC11]. Also charge sensing using single electron transistors has been pursued. These technologies of direct qubit measurement have largely been replaced by measurement of microwave radiation scattered off the qubits using high electron mobility transistors (HEMT [HMB⁺15]), Josephson Parametric Amplifiers ([BSM⁺10, CBIH⁺08, HVS⁺11]) and their broadband multi-junction version, the traveling wave parametric amplifier (TWPA [MOH⁺15]). An alternative but currently less developed approach uses photon counting [GPX⁺14, GPP⁺15, Il⁺16].

13.1.3.4 Cryogenics

Superconductivity is a low-temperature phenomenon and requirements of coherence require temperatures below 100 mK. This is achieved by dilution cryostats [Pob96] - multistage cooling systems whose coldest stage uses a mixture of the Helium isotopes ³He and ⁴He. These are commercial devices that in some cases are customized to hold a large number of microwave lines. While currently hassle-free workhorses, they pose three challenges: i) Their limited cooling power at low temperatures requires to direct energy dissipation to higher temperature stages ii) the requirement of good shielding and heat management restricts the available sample space to small volumes challenging scaling and iii) worldwide supplies of Helium are low and of the (not naturally abundant) isotope ³He are critically low. Large scale production of this isotope requires nuclear facilities.

13.1.3.5 Microwave components

Requirements of low dissipation at low temperatures as well as isolation and routing of signals require microwave peripherals close to the sample, typically at 1 K. A critical component are non-reciprocal elements, elements that transmit radiation differently depending on their direction, such as gyrators, circulators, isolators, and directional couplers. These all have to be longer than the wavelength of a few centimeters, hence strongly limiting miniaturization and scaling. Several efforts to overcome this limitation are under way [KLC⁺15, CMR⁺16, CR14, BGW⁺15, VD14].

13.1.4 Quantum annealing and its status with superconductors

Adiabatic quantum computing describes the process where the solution of a hard computational problem is encoded in the ground state of a complex Hamiltonian which is hard to reach classically, and using an adiabatic sweep that starts out from an easy Hamiltonian to reach that ground state. A variation of this, quantum annealing, allows faster sweeps as long as the combined action of thermal relaxation and quantum tunneling take the system back to its ground state. These techniques have in common that they require much simpler time-dependent control (ideally only a singly, slowly varied parameter) than implementing a

quantum circuit. This is in particular true for the case of Josephson qubits, where it is a major engineering challenge to apply qubit-specific microwave signals.

In discussing adiabatic quantum computing, one needs to sharply distinguish two classes of applications. We start with the less popular but more cryptographically relevant one.

13.1.4.1 Adiabatic realization of quantum circuits

It has been shown that adiabatic quantum computing is as efficient as circuit-based quantum computing [AvDK⁺07]. The mapping proposed in this paper takes any quantum circuit and maps it onto an annealing architecture and it shows that the energy gap above the ground state (whose inverse sets the time scale for execution of the adiabatic algorithm) shrinks polynomially with the number of gates in the circuit, hence proving that annealing can be as powerful as circuit-based quantum computing—e.g., for Shor’s algorithm. This results assumes that each lattice site in the annealer contains a six-state particle—which can either be directly implemented or simulated by putting more than one physical qubit at each site. It also assumes the presence of three-body couplers, i.e., terms in the Hamiltonian that contain non-trivial operators at three distinct qubits (also referred to as 3-local couplers). Physically, this corresponds to a three-body interaction, which famously does not exist in nature—and in particular in superconducting qubits the capacitive and inductive interaction are all two-body.

As a way out, three- or more-body couplers can be implemented by a technique called perturbative gadgets [KKR06, JF08, BLAG14] that formalize the idea that nonlinear higher energy degrees of freedom can introduce a low-energy interaction whose properties resemble that of a many-body interaction. There is a wealth of current proposals [AvDK⁺07, Bia08, LHZ15, CZW16], none of which has been realized or even seriously attempted. Realizing these is the next frontier in quantum annealing.

13.1.4.2 Adiabatic optimization

Quantum annealing / adiabatic quantum computing naturally lends itself to the solution of hard constrained optimization problems such as 3SAT (an NP complete problem). For these, no efficient quantum circuit is known. There is no proof (or disproof) of quantum speedup for this problem. Current experimental scaling on the D-Wave machine (comments below) indicate that speedup is questionable at best [RWJ⁺14]. It has been shown that for certain extreme cases, there is significant speedup [DBI⁺16] but the result does not allow conclusions for scaling and for generic problems. Currently, it is highly disputed whether this experimental evidence only points at shortcomings of the d-wave machine or hints at the lack of speedup in quantum annealing for 3SAT.

13.1.4.3 Experimental situation

A lot of early aggressive scaling in quantum computing as a whole has been performed in adiabatic quantum computing/quantum annealing. This was largely driven through the Canadian company d-wave systems (see below). The D-wave architecture is optimized for a subclass of hard optimization problems. Even within the paradigm of adiabatic optimization, this machine is not fully general. It does not contain sufficiently general couplers (many-body couplers are not implemented and competing non-commuting interactions are not implemented). Also, the qubits are not very coherent—they would not allow to implement any meaningful quantum circuit. Hence, while there are quantum effects in this machine, speedup is highly questionable [RWJ⁺14] as it has been outperformed by specialized (i.e., sufficiently restricted) classical algorithms.

There is an effort to build a more general annealer at MIT Lincoln Laboratory (description below) that does not have these shortcomings, supported by the IARPA Quantum Enhanced Optimization (QEO) program¹.

While this program declares to also aim for optimization, a QEO annealer would also be able to run Shor's algorithm with acceptable overhead and minimal modification.

All of these annealing hardware platforms are based on flux qubits.

13.1.5 Research groups

13.1.5.1 Current community leaders and their networks

IBM The IBM cooperation at their Watson Research Center in Yorktown Heights, NY is operating a major integrated quantum computing laboratory with a focus on superconducting qubits and a strong theory group. Recently, their laboratory in Rüslikon, Switzerland has started to built their own laboratory, also using the superconducting qubits platform. Their work is in parts funded by government grants. Their theoretical effort is lead by Dr. Jay Gambetta and contains Dr. Sergey Bravyi and Dr. John Smolin as further senior researchers. Their experimental effort is led by Dr. Jerry Chow. Originally, this activity was developed under Dr. Matthias Steffen and Dr. Mark Ketchen. These researchers still appear at conference, publish only minimally but patent actively, it is rumored that they are running a classified quantum computing program. The main activity is in the area of scalable quantum computing with surface code error correction. Their highly integrated effort includes a large number of scientists with interdisciplinary backgrounds. IBM is commercializing quantum computing both as hardware as well as cloud access. Their programming ecosystem QISKIT has reached strong adoption.

Google Quantum Artificial Intelligence Labs (John Martinis, Hartmut Neven) Google initially adopted quantum annealing and built, guided by computer scientist Dr. Hartmut Neven, a theoretical research unit in Venice, CA, which currently employs about ten scientists with a doctoral degree. Initially owing a share of a d-wave machine, they realized its shortcomings and aimed at building both a coherent annealer and a circuit based quantum computer - at which point they included a hardware group by hiring John Martinis.

John Martinis is professor at UC Santa Barbara, CA since 2004. In 2014, he became a research scientist at Google (while maintaining university affiliation) which made him one of the best funded researchers in the field. He runs a dedicated facility in Santa Barbara as well as on-campus laboratories at UC Santa Barbara. Today, he has both a large staff of quantum electronics engineers (a job title applied mostly to very focused researchers trained as experimental physicists) at Google including Dr. Austin Fowler (pioneer of the surface code [FMMC12]) and an academic group at UC Santa Barbara with PhD students. They are today's world leaders in scaling up superconducting qubits, namely Xmons. They have recently published an experiment with 53 qubits and high gate fidelities, claiming quantum supremacy. Martinis' group closely collaborates with the theoretical groups of *Michael R. Geller*, University of Georgia, USA, *Alexander N. Korotkov*, UC Riverside, USA, *Clare C. Yu*, UC Irvine, USA, *Anatoli Polkovnikov*, Boston University, USA, *Enrique Solano*, University of the Basque Country, Spain, and *Frank Wilhelm*, Universität des Saarlandes, Germany, and with the experimental group of *Robert McDermott*, University of Wisconsin Madison, USA.

Robert J. Schoelkopf, Yale University Rob Schoelkopf is Sterling Professor of Applied Physics and Physics and Director of the Yale Quantum Institute. His group's research focuses on the development of superconducting devices for quantum information processing. Together with his collaborators at Yale, Professors Michel Devoret and Steve Girvin, their team created the field of circuit quantum electrodynamics (see above) which allows quantum information to be distributed by microwave signals on wires. His lab has produced many firsts in the field based on these ideas, including the development of a "quantum bus" for information, and the first demonstrations of quantum algorithms and quantum error correction with integrated circuits. The group has been an innovator in high-speed measurement techniques at ultra-low temperatures, and invented numerous devices such as the RF single-electron transistor, the shot noise thermometer, and the transmon qubit. He is expanding in using these paradigm to connect to other platforms. , his publication list is long and thematically broad (20 publications in 2016). His main thrust in quantum computing is currently based on encoding qubits in continuous variables in cavities that are immune to photon loss, so-called *compass states*. While the experiments are impressive, there is no known

way to concatenate these codes so the overall scaling perspective remains elusive. He collaborates predominantly with the other Yale groups of *Michel Devoret*, *Steve Girvin*, *David DeMille* and *Liang Jiang*. During his time as a professor, Rob Schoelkopf supervised a lot of young scientists which became famous later on including *Leonardo DiCarlo*, TU Delft University of Technology, Netherlands, *Andrew Houck*, Princeton University, USA, *Gerhard Kirchmair*, University of Innsbruck, Austria, *Konrad Lehnert*, University of Colorado, Boulder, USA, *Johannes Majer*, Atominstitut TU Vienna, Austria, *David Schuster*, University of Chicago, USA, *Ken Segall*, Colgate University, USA, *Irfan Siddiqi*, University of California, Berkeley, USA, and *Andreas Wallraff*, ETH Zürich, Switzerland.

Michel Devoret, Yale University Michel Devoret co-founded the Quantronics Group at Saclay (see below) and moved to Yale later in his career. Widely regarded as an authority in the field, his research currently focuses on parametric readout techniques beyond the parametric amplifier, including the parametric converter. Some of these are deployed in quantum computing, some in more basic research.

Leonardo DiCarlo, TU Delft After his PhD in Harvard in the group of Charles M. Marcus and his very successful postdoc in Rob Schoelkopf's group, Leonardo DiCarlo became a professor at TU Delft University of Technology, Netherlands, in 2012. Although his group is comparatively young, it became one of the world leading groups in development of superconducting quantum circuits with applications in quantum computing. The group is mainly characterized by the topics of quantum measurements and feedback control for the circuit quantum electrodynamics architecture. In 2013, QuTech, the quantum institute of TU Delft and TNO, was founded. In 2015, American chip manufacturer Intel and QuTech have finalized plans for a ten-year intensive collaboration that has recently materialized in a fully packaged 17-qubit chip. Alongside financial support for QuTech totaling approximately \$50 million, Intel will also contribute expertise, manpower and facilities. DiCarlo is responsible for the superconducting quantum circuits platform during this collaboration time. On top, DiCarlo collaborates with the theoretical groups of *Alexandre Blais*, Université de Sherbrooke, Canada, and *Frank Wilhelm*, Universität des Saarlandes, Germany, and with the experimental group of *Andreas Wallraff*, ETH Zürich, Switzerland.

Andreas Wallraff, ETH Zürich After his postdoc in Rob Schoelkopf's group, Andreas Wallraff joined the Department of Physics at ETH Zürich in January 2006 as a Tenure Track Assistant Professor where he became a Full Professor for Solid State Physics in 2012. His research is focused on the experimental investigation of quantum effects in mesoscopic electronic circuits for performing fundamental quantum optics experiments and also for applications in quantum information processing. His group at ETH Zürich engages in research on micro and nano-electronics, with a particular focus on hybrid quantum systems combining superconducting electronic circuits with semiconductor quantum dots and individual Rydberg atoms, making use of fast and sensitive microwave techniques at ultra-low temperatures. Andreas Wallraff is a member of many scientific networks and collaborations. One can find a long list of collaborators on his webpage including for example *Leonardo DiCarlo*, TU Delft University of Technology, Netherlands, *Alexandre Blais*, Université de Sherbrooke, Canada, and *Enrique Solano*, University of the Basque Country, Spain. He also has a project with IBM Zürich.

D-Wave Systems, Vancouver D-Wave Systems is a spin-off of the University of British Columbia. Originally a pure theory and design team, they identified quantum annealing as a promising (and easy) route to scaling. They initially used the foundry at IPHT Jena and the Laboratory of Dr. Evgeni Il'ichev there to develop experiments and prototype, before starting their own laboratory. Based on significant funding through venture capital and first sales, this group is large (employs more than 60 scientists and engineers). With all criticism of their approach, they have demonstrated a lot of important engineering solutions, including the fact that a cryostat can be run in a data center.

Of the pioneering generation, only Dr. Mohammad Amin is still with the company.

Rigetti Computing, Berkeley Rigetti Computing is a quantum computing startup that was founded in 2013 by former IBM researcher Chad Rigetti and is located in Berkeley and Fremont, California. While exact numbers of employees are not listed on their website, it can be estimated to lie somewhere between 50-100. Even though the number of quantum technology start ups is rising, Rigetti Computing remains the only competitive startup that aims at constructing a full (hardware and software) stack superconducting

quantum computer. Their research spreads over the whole vertical chain of superconducting quantum computing, including quantum integrated circuits, multi-qubit devices and quantum programming (and optimization) tools. After initially being rather secretive, Rigetti Computing has published numerous papers, also in collaboration with other research groups such as the Google quantum AI team. Rigetti Computing has recently opened a quantum cloud service with their own software package *Forest*. After a reorientation period, they have now partnered with Amazon to put quantum computing into their cloud ecosystem.

Wallenberg Center for Quantum Technologies, Chalmers University of Technology Based on a long-standing strong effort at Chalmers, a new, well-funded center is pursuing Josephson junction qubits both within the quantum technologies flagship as well as on their own. Activities are led by Per Delsing and Jonas Bylander on the experimental and Göran Johansson and Göran Wendin on the theoretical side.

Alibaba Quantum lab has been able to recruit a strong and relatively young team of researchers in order to develop both hardware and software access to quantum processors. In a collaboration with various universities, they have a system of 20 qubits connected to a central bus online.

13.1.5.2 Further important contributors

CEA Saclay The Groupe Quantronique at the Commissariat de l'Energie Atomique et Alternative in Saclay, France, is a leading group in all aspects of fundamental solid-state electronics in the nanoscale. Senior researchers include Dr. Daniel Esteve, Dr. Denis Vion, and Dr. Patrice Bertet. They have made a number of important pioneering contributions to the development of qubits and coherence, most notably the Quantonium qubits. Recently, given their broad interest and the peculiarity of the French support system that does not enable large integrated efforts that can be competitive, they have moved to more specialized questions such as quantum sensing. They participate in various research networks.

Konrad Lehnert, JILA Konrad Lehnert is an alumnus of Rob Schoelkopf at the University of Colorado in Boulder and the Joint Institute for Laboratory Astrophysics (JILA). His research focuses on amplification and measurement.

IQM, Helsinki is a spinoff of Aalto University, specifically the group of Mikko Möttönen. They have only started in 2019 but seem to move fast and certainly have strong investor backing.

Raymond Simmonds, Joe Aumentado, NIST Boulder NIST investigates quantum computing in the electric standards division. Ray Simmonds started as a postdoc of John Martinis before Martinis moved to Santa Barbara. He still investigates phase qubits. Aumentado focuses on measurement science.

Irfan Siddiqi, UC Berkeley Irfan Siddiqi is a postdoctoral alumnus of Michel Devoret. Similarly focused on technical aspects of readout, he advanced the traveling wave parametric amplifier (TWPA) which combines the high resolution of the parametric amplifier with larger frequency bandwidth, allowing to integrate multiple qubit channels into the amplifier chain. His young alumni, with similar profiles include Michael Hatridge (University of Pittsburg), Kater Murch (Washington University at St. Louis) and R. Vijayaraghavan (Tata Institute of Fundamental Research, India).

Laboratory for Condensed Matter Physics, Moscow The Condensed-matter physics Laboratory was established in June 2016. Under the direction of Dr. Lev Ioffe and Dr. Yuriy Makhlin, the laboratory brings together leading Russian physicists from research centers of the Russian academy (Landau Institute for theoretical physics, Kapitza Institute for physical problems, Institute of solid-state physics, Institute for Spectroscopy, P.N. Lebedev Physical Institute). Their research divides into three directions: condensed-matter theory, low-temperature physics and nanophysics, and hydrodynamics and turbulence. Particularly, they work on the theory of quantum-coherent phenomena in disordered condensed matter and nanostructures at low temperatures including quantum computing devices.

Quantum Information Electronics Division, RIKEN Center for Emergent Matter Science The Quantum Information Electronics Division divides into 12 specialized research groups including senior researchers Dr. Franco Nori, Dr. Jaw-Shen Tsai and Dr. Yasunobu Nakamura (who built the first working superconducting qubit). Nakamura's group (also at the University of Tokyo) studies quantum electronics in

superconducting circuits and hybrid quantum systems. The targets of their research include quantum information processing and ultra-sensitive measurement based on quantum mechanical properties of the circuits. Tsai's group tries to use this platform to build a superconducting quantum simulator.

Robert McDermott, University of Wisconsin-Madison The group of Robert McDermott, professor at University of Wisconsin-Madison and Martinis' group alumnus, is working on technologies for scalable qubit measurement and control which is a basic requirement for realization of improved large-scale superconducting quantum circuits. One of his main focuses is high-fidelity qubit readout by microwave photon counting with Josephson photomultipliers (JPMs). As a recent breakthrough, McDermott identified a long-standing puzzle related to flux noise in superconducting loops and provided a solution to lower it by advanced sample fabrication. McDermott has several U.S. patents for superconducting quantum devices.

Russian Quantum Center, Moscow The Russian Quantum Center (RQC) hires a lot of native Russian researchers who are currently distributed all over the world. Their Superconducting Quantum Circuits group is led by Prof. Dr. Alexey Ustinov who is professor at Experimental Physics, Karlsruhe Institute of Technology, Germany. The affiliated faculty includes Prof. Oleg Astafiev, Royal Holloway University of London, UK, Prof. Evgeni Il'ichev, Leibniz Institut für photonische Technologien, Jena, Germany, and Prof. Valery Ryazanov, Moscow Institute of Physics and Technology, Russia.

William D. Oliver, MIT Lincoln Laboratory William D. Oliver is a Principal Investigator in the Engineering Quantum Systems Group (MIT campus) and the Quantum Information and Integrated Nanosystems Group (MIT Lincoln Laboratory). His research interests include the materials growth, fabrication, design, and measurement of superconducting qubits, as well as the development of cryogenic packaging and control electronics involving cryogenic CMOS and single-flux quantum digital logic. As a government facility, LL is rather secretive but given his productivity also in the area of classical superconducting supercomputers his group must include about sixty researchers. His group is now a main source for TWPAs and at the core of the new IARPA quantum annealing program.

Britton Plourde, Syracuse University Britton Plourde is professor at Syracuse University and operates a foundry for superconducting circuits by leveraging the Cornell University cleanroom facility.

Fred Wellstood, University of Maryland Wellstood has made pivotal contributions to noise in superconducting devices in his PhD with John Clarke. His group at the Center for Nanophysics and Advanced Materials (formerly Center for Superconductivity) has pioneered investigation of phase qubits and couplers but fell behind in applying time-dependent control. He is now the director of the Joint Quantum Institute, a Center between UMD and NIST Gaithersburg.

Adrian Lupascu, University of Waterloo Lupascu primarily focuses on flux qubits and has recently started to develop towards quantum annealing.

Raytheon-BBN, Cambridge BBN is a high-tech engineering spin-off of MIT that was acquired by the Raytheon group, a defense contractor. Their superconducting qubits laboratory, supports the IBM efforts and produces bespoke analogue control electronics. They are looking to start their own effort but have recently suffered a brain drain towards Rigetti.

Northrup-Grumman, Baltimore Defence contractor Northrup-Grumman acquired superconducting electronics specialist TRW several years ago. They are primarily engaged in classical superconducting supercomputers (group led by Dr. Quentin Herr) but recently started to interface this technology as a control layer to superconducting qubits, with limited success (led by Siddiqi Alumnus Dr. Ofer Naaman). They recently started to engage in quantum annealing.

Zurich Instruments Zurich Instruments makes modern innovative measurement equipment. They are partners of an IARPA program for superconducting qubits where they develop tailored microwave equipment.

Lockheed Martin Lockheed Martin describes themselves unofficially as a software company that also builds planes to run their software on. Thus clearly identifying their need for advanced computing, they invested in a d-wave machine that is physically located at the University of Southern California to explore applications of quantum annealing to aerospace problems. There is no hardware effort known but strong interest in application development.

SQDLab, University of Queensland The group started in February 2013 when Dr. Arkady Fedorov joined UQ as a faculty member after his work at ETH Zürich, Switzerland. Their group research continues to focus on superconducting nanodevices, which both use and probe mysterious principles of quantum physics.

Tim Duty, UNSW Timothy Duty is Professor in the field of experimental condensed matter physics, who leads the superconducting device laboratory at the University of New South Wales (UNSW) in Sydney. They are specialized in fabricating and measuring quantum effects in nanoscale superconducting devices based on Josephson junctions.

Yang Yu, Nanjing Yang Yu is Professor at the National Laboratory of Microstructures of Nanjing University, which is one of China's largest research centers for condensed matter physics. He is an alumnus of MIT. Prof. Yu has contributed in a novel work on creating and manipulating tripartite qubit systems.

Haohua Wang, Zhejiang John Martinis alumnus Haohua Wang is Professor at Zhejiang University, China and has started the Superconducting Quantum Circuit Group in 2010. The Group is dedicated to studying quantum phenomena of superconducting devices at ultra-low temperatures and realizing practical quantum information technology. Prof. Wang is still collaborating and publishing together with John Martinis and also with Andrew Cleland (Exotic Qubits).

Siyuan Han Siyuan Han is Professor for experimental condensed matter physics at the University of Kansas. Besides more fundamental interests in quantum mechanics, he works on the realization of quantum computation with superconducting qubits and has co-authored a pioneering flux qubit paper. Prof. Han has a lot of collaboration with Chinese physicists from Nanjing, Zhejiang, Hangzhou and others.

13.1.5.3 Pioneers

While not active any more, the structure of the community can be tracked to a few pioneers. Anthony “Tony” Leggett (University of Illinois) has pioneered the idea to study macroscopic quantum coherence in superconducting circuits. His work in superfluidity was honored with the 2003 Nobel Prize in Physics. The first experiment along Leggett’s original idea was conducted in the group of James Lukens (SUNY Stony Brook) with Jonathan Friedman (now at Amherst College). John Clarke (UC Berkeley) has investigated SQUID detectors throughout his distinguished career and has explored quantum and noise limits. His alumni include John Martinis, Michel Devoret, and Daniel Esteve. Johan E. “Hans” Mooij (TU Delft) has applied a lot of Leggett’s ideas in experiments and co-invented the flux qubit with Terry Orlando (MIT). Gerd Schön (KIT) was the first theorist to ever formulate a mathematically consistent theory of superconducting qubits.

13.2 Evaluation: Planar transmon

The planar Transmon is currently the leading Josephson qubit circuit, being on a sweet spot with high coherence, connectivity, and reproducibility. The largest known processors (17 Qubits at IBM and 22 at Google) are built from these devices. Older announcements of Google to reach 50 qubits this year are now temporarily withdrawn—a chip will be *fabricated* in time but not tested.

13.2.1 DiVincenzo criteria

13.2.1.1 Scalable qubits

Various organizations have reached large multi-qubit devices

- Rigetti: up to 8 qubits on chip, no significance difference in coherence or fidelities compared to two-qubit chips
- Google: 9 qubits published [KBF⁺15, NRK⁺17], superconducting gmon qubits with tunable frequencies and tunable interactions, 22 qubits shown on conferences [Mar17]
- IBM: 5 qubits published in detail, 16(17) tested and available for experiments now. Successful demonstration of 15-qubit repetition code [WL17].
- TU Delft: 17 qubits currently tested with Intel chip

13.2.1.2 Initialization

Best fidelities reached by fault tolerant state preparation at IBM, initialization can be done by measurement [RvLK⁺12].

13.2.1.3 Universal gates

High gate fidelities have been reported in various architectures and gate implementation schemes. Systems with higher numbers of qubits have shown no significant decrease of fidelities. Also, the possibility for parallel gates has been presented.

- Typical gate fidelities in IBM's 5 qubit setup are 0.997 and 0.965 for single- and 2-qubit gates [LMR⁺17]
- IBM QX3 with 16 qubits promises single- and multi-qubit gate errors of $\sim 2 \cdot 10^{-3}$ and $\sim 4 \cdot 10^{-2}$, CNOT gate times are ~ 600 ns (see IBM QX3), implemented via cross resonance [CBSG17], all gates can be done in parallel
- single- and multi-qubit gate times as low as 10ns and 25ns (Transmon-bus) have been reported [BKM⁺13, GGZ⁺13]
- IBM 5-qubit chip [SMCG16, SBM⁺16]: $F_1 = 0.9995$, $F_2 = 0.991$ RB, two-qubit gate time of 160ns
- Rigetti 8-qubit chip [ROT⁺17]: $F_1 = 0.989$, $F_2 = 0.916$ RB
- Google 5-qubit (xmon) chip [BKM⁺14]: $F_1 = 0.9992$, $F_2 = 0.994$ RB, two qubit gate times of 40ns

13.2.1.4 Other tomographies

Two-qubit gates process tomography fidelities of 0.93 (iSwap) and 0.92 (CZ) have been reported at Rigetti for a two-qubit chip [CDR⁺17].

13.2.1.5 Coherence

Coherence in 2D Transmons is consistently high, following its design principle [KYG⁺07, HKD⁺09]

- IBM 5-qubit chip: $T_1 = 70\mu\text{s}$, $T_2^* = 92\mu\text{s}$ [RGP⁺12]

- Rigetti 2-qubit chip: tunable (fixed) Transmon: $T_1 = 13(34)\mu\text{s}$, $T_2^* = 4-10(20)\mu\text{s}$ [CDR⁺17]
- Rigetti 8-qubit chip, odd qubits tunable: $T_1 = 19\mu\text{s}$, $T_2^* = 12\mu\text{s}$ [ROT⁺17]
- Google single qubits: $T_1 = 44\mu\text{s}$, $T_2 = 20\mu\text{s}$, $T_2^* = 15\mu\text{s}$ [BKM⁺13]
- IBM QX3 (16 qubits) $T_1 \sim 25-50\mu\text{s}$, $T_2 \sim 30-80\mu\text{s}$ IBM QX3
- Google 9-qubit gmon: $T_1 \approx 10-15\mu\text{s}$, $T_2 \approx 5\mu\text{s}$ [NRK⁺17]

13.2.1.6 Readout

- Wallraff, ETH: dispersive readout of Transmon qubit with 0.9825 readout fidelity in 48ns or 0.992 in 88ns [WKG⁺17]
- Martinis: 0.998 in 140ns was best reached fidelity, a more average value is 0.99 in 200ns [JSM⁺14]
- IBM 5 qubits: Single-qubit readout fidelities typically ~ 0.96 [LMR⁺17]
- IBM QX3: readout error $\sim 5 \cdot 10^{-2}$, measurements can be done simultaneously

13.2.2 Summary

Criteria	met?	Comments
Scalable qubits	✓	
Initialization	✓	
Universal gates	✓	
Coherence	✓	
Readout	✓	

Table 13.1: Summary of DiVincenzo criteria for planar Transmon qubits. ✓: Met routinely, ?: Met sometimes or meeting them is controversial, ×: not met.

t_i	t_1	t_2	t_M	p_i	p_1	p_2	p_M	T_2
140ns	10ns	40ns	140ns	$2 \cdot 10^{-3}$	$8 \cdot 10^{-4}$	$6 \cdot 10^{-3}$	$2 \cdot 10^{-3}$	80 μs

Table 13.2: Optimistic assumptions for different error rates and operation times combining the best reached values for each operation. This does not describe a current available setup, but shows what is in principle possible right now or in near future. Times are initialization, 1- and 2-qubit gate and measurement time. Probabilities are error probabilities for the respective processes. A surface code cycle contains 4 two-qubit gates, 2 one-qubit gates, measurement and initialization as well as classical processing.

13.2.3 Fault-tolerant extrapolation

With current experimental advances, running successful error correction is imaginable, although high physical error rates make it quite resource demanding and long measurement times give a limit to the problem size that can be solved in reasonable time. Faster measurement is possible, but at the cost of lower fidelities. With a combination of best reached values for measurement, initialization, gates, and coherence as given in Table 13.2, an estimation of physical qubit costs for running dlog and factoring algorithms was made using the Autotune tool Polyestimate [Fow13b]. Figure 13.1 shows possible realizations given target runtimes of 1 to 100 days and the limit of parallelization due to T depth. Factoring algorithms would only be possible for really long runtimes: For the chosen sets, only $n = 1024$ is possible in 10 days, with $\sim 3 \cdot 10^{10}$

qubits (or in 100 days with 10 times less qubits), $n = 2048$ and $n = 3072$ already require 100 days runtime. Larger problem sizes can not be parallelized to runtimes as low as 100 days as that would violate the temporal order of sequential T gates. The chosen dlog problem sizes seem much more reachable: Problem sizes of up to $n = 356$ could theoretically be run in 1 day, with $10^{10} - 10^{11}$ qubits. Even for those problems, distances still over $d = 100$ in the Clifford part and two distillation rounds for T gates (with distances ~ 130 and ~ 60) were required.

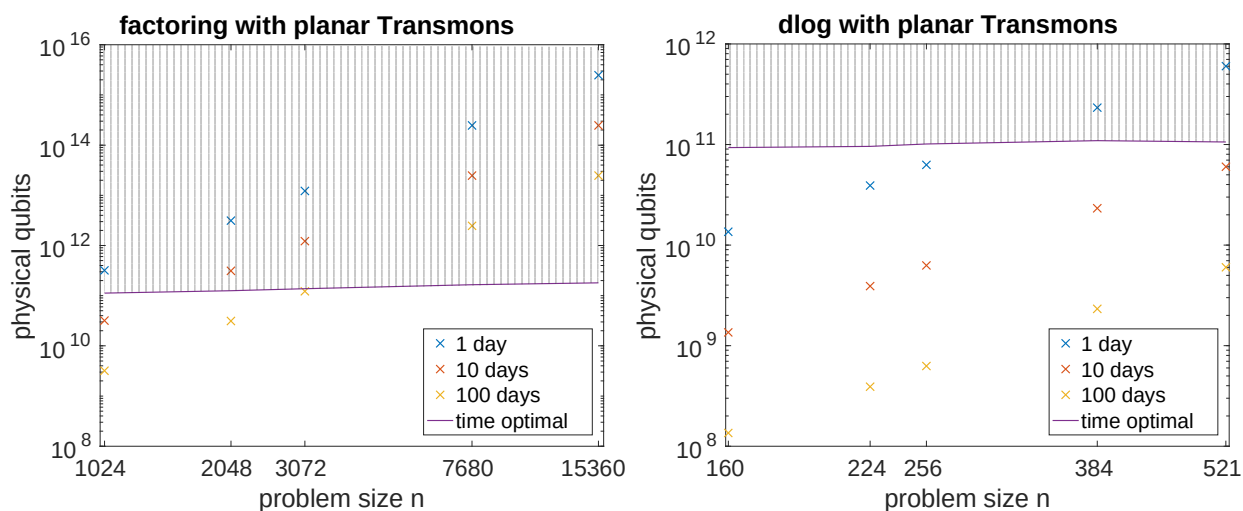


Figure 13.1: Double-logarithmic plot showing possible system sizes when implementing factoring (left) or dlog (right) algorithms for different runtimes, with experimental values from Table 13.2. (rather than the uniform error probabilities of the other chapters, see chapter 7.4.1 for background) The violet line gives the number of qubits for the time-optimized limit, i.e., the minimal required runtime due to the T depth when having maximal parallelization. Implementations above the line would violate the temporal order of T gates and are thus not possible.

13.2.4 Analysis and outlook

Albeit postponed, the crossover to a 50-qubits “quantum advantage” device has been met in 2019 []

2D Transmons have demonstrated all basic ingredients for the demonstration of error correction. For the steps ahead, challenges in engineering, operation, and scale-up science are clearly visible:

- breaking the plane—chains and double-chains of qubits can be approached by control and read-out lines within the plane. Truly scalable architectures need a separate control layer or wire crossings. This is currently a very active area and promising solutions are shown at conferences—this problem is likely being solved soon.
- consistency—qubit fabrication, in particular fabrication of Josephson junction, currently has limited yield. Making a chip where every junction work needs to improve this. As in other areas this has been achieved, it is likely that professional process control will enable this. Companies like Rigetti pride themselves of their fabrication consistency but do not publish any verifiable details.
- cryogenics—although low temperature techniques are well established and robust, there are problems arising when going to even lower temperatures. The lower the temperature, the lower the cooling power. Furthermore, temperatures of 1K require expensive ^4He , lower temperatures need ^3He , which is even harder to get. ^3He is produced as a byproduct of nuclear stockpile stewardship or specialized nuclear reactors and heavily used for neutron detection. The current industrial demand is around 70000 liters per year, with a per-liter price of \$2000. Detailed production figures could not be find and are probably classified. Modern dilution-free cryostats are in principle closed and do not waste any Helium even though its quality degrades on the scale of years. A large modern cryostat like the Oxford Triton XL

[Oxf17], which has been developed to suit the need of quantum information research, requires a filling of 70 liters.

- size—the microwave periphery still consumes majority of space: Per Transmon, which has a size of around 0.1mm^2 on the chip, two circulators of about 10^{-5}m^3 are currently required for readout. Challenges are towards smaller circulators [VD14, MCP⁺17], multiplexed readout—requiring a lower number of circulators, or digital readout [MVP⁺17] without any circulators at all.
- room temperature electronics—rack-mounted classical control electronics currently is mostly laboratory equipment, hence optimized for flexibility, not for space. Conceivably, making this smaller will not be as hard as other scaling challenges. Also, recent research in signal routing allows to use generators for multiple qubits [ADL⁺16].
- closely packing multiple chips, 100.000 qubits would fit on a cold plate of a Triton XL. Assuming 10 readout channels per circulator and somewhat more compact circulators, also 10.000 controls at 1K would fit in Triton XL. A concerted effort that would require to build a cryostat that is ten times as large appears to be possible. With additional effort in miniaturizing controls and qubit footprints and customizing and integrating pulse generation, even 10s of Millions of qubits would be conceivable.
- connecting cryostats—building larger systems would only be possible with superconducting microwave interconnects. With a speed of approximately $c/2$ it would take around 25ns to transfer a signal between two cryostats. However, current remote entanglement protocols require detection and are estimated at around 750ns, creating global slowdown for the clock of such a processor. Accelerating this is a major challenge.
- connectivity: Even though qubits can be arranged on two-dimensional lattices, enabling both-way coupling between all neighboring qubits only works if all engineering targets are met. The IBM QX3 device for example can only perform CNOT gates in one direction (with one qubit marked as control and one as target), some couplings are not possible at all due to an arrangement of qubit frequencies that is not commensurate with the original design. This is expected to improve with more consistent fabrication.

13.3 Evaluation: 3D transmon

3D Transmons reach superior coherence but less flexibility and thus slower control.

13.3.1 DiVincenzo criteria

3D Transmons have passed the DiVincenzo criteria.

13.3.1.1 Scalable qubits

A four qubit device was realized at IBM [PMS⁺16]. Multiple cavities can be coupled through bridge-qubits, i.e., qubits coupled to multiple cavities at once.

13.3.1.2 Initialization

Is done by measurement and postselection: Fidelity of 0.988 [RvLK⁺12] has been measured.

13.3.1.3 Universal gates

Single-qubit gates are realized with local microwave drive, two-qubit CZ gate via driven common resonator.

- Single qubit gates with 0.999 RB-fidelity in 36.7ns, two-qubit gates with 0.98 RB-fidelity in ~ 400ns [PMS⁺16]

13.3.1.4 Coherence

The 3D Transmon was designed with maximum coherence in mind, hence numbers are superior

- $T_1 = 90\mu\text{s}$, $T_2^* = 48\mu\text{s}$, $T_{2E} = 86\mu\text{s}$ [PMS⁺16], cavity decay rate $\kappa = 7.7\text{kHz}$
- Lincoln Lab: $T_1 = 80\mu\text{s}$, $T_2 = 115\mu\text{s}$, $T_{2E} = 154\mu\text{s}$ [JKS⁺15]

13.3.1.5 Readout

Readout of 3D Transmons is done analogous to 2D Transmons [PMS⁺16]. A readout fidelity of 0.99 has been reached and 0.999 seems theoretically reachable.

- 0.981 with homodyne detection enhanced by JPA [RvLK⁺12]
- 0.99 [RD15]
- 0.999 in 60ns theoretically proposed [DDBA13]

13.3.2 Summary

Criteria	met?	Comments
Scalable qubits	✓	
Initialization	✓	
Universal gates	✓	
Coherence	✓	
Readout	✓	

Table 13.3: Summary of DiVincenzo criteria for 3D Transmon qubits.

t_1	t_1	t_2	t_M	p_1	p_1	p_2	p_M	T_2
60ns	40ns	400ns	60ns	0.99	10^{-3}	$2 \cdot 10^{-2}$	0.99	115 μs

Table 13.4: Optimistic assumptions for different error rates and operation times combining the best reached values for each operation. This does not describe a current available setup, but shows what is in principle possible right now or in near future.

13.3.3 Analysis and outlook

3D-transmons are about to cross the fault tolerance threshold, going from B to C. Still, their size leads to technological challenges to scaling, making it unlikely that they will overtake 2D-transmons.

- benefits vs drawbacks: The higher coherence of 3D-transmons comes with a reduction of flexibility, lower fidelities and much higher volume cost that overcompensate the gains. A single 3D-transmon in its cavity requires 50cm^2 on the cold plate of a dilution cryostat. If one manages to remain modular with multiple qubits per cavity, this number can be brought down to : $\sim 20\text{cm}^2$, which is still a large infrastructure challenge compared to 2D-transmons.
- Frequency crowding in larger networks [PMS⁺16]: Many 2-qubit gates work by tuning the frequencies of two qubits into a specific resonance condition. This involves higher energy levels, which might cause

other energy levels to intervene between the qubits. Using WAHWAH [SDEW13, TMW16] control pulses is a possible solution to this.

13.4 Evaluation: Flux qubit

The flux qubit is dominating quantum annealing due to its high connectivity (see Section 8.1 for comments on fault-tolerance and benchmarking). It also presents a superior interface to other quantum systems such as spin, which is interesting for building quantum repeaters.

Flux qubits have demonstrated all DiVincenzo criteria and due to their large anharmonicity allow for ultra-fast gates.

13.4.1 DiVincenzo-Criteria

13.4.1.1 Scalable qubits

Large arrays of flux qubits, albeit with low coherence, have been demonstrated for quantum annealing by D-Wave Systems [BHJ⁺14, BCI⁺16, KXB⁺16, LKEH17]. This seems feasible with more coherent qubits as well. When going to gate-based quantum computing, precise frequency allocation is important which is difficult as it requires unusually precise fabrication. This can be mitigated by using a capacitive shunt or a two-loop design [SKD⁺10, GBY⁺11, SGJ⁺13, YGK⁺16] which are so far not tested as much as the simple flux qubits. With an eye on these challenges, flux qubits can be viewed as scalable.

13.4.1.2 Initialization

Initialization is achieved via cooling, assisted by the large available energy splittings

13.4.1.3 Universal gates

- IQC: single-qubit gate in $\sim 0.5\text{--}1\text{ns}$ with fidelities of 0.996–0.999 [DOS⁺15]
- Mooij: CNOT demonstrated, but with fidelity of 0.4 [PdGHM07]
- Optimal control: single-qubit gates below 1ns and CNOT in 2ns [HG14], limited by leakage (10^{-6}) and decoherence errors ($\sim 10^{-5}$) theoretically proposed and simulated

13.4.1.4 Coherence

- Lincoln Lab: $T_1 = 40\mu\text{s}$, $T_2 = 85\mu\text{s}$, $T_{2E} = 40\mu\text{s}$ [YGK⁺16]

13.4.1.5 Readout

Single shot readout via inductively coupled dc-SQUID [LmcHM05], with measurement time $T_m \approx 300\text{ns}$ and fidelity > 0.8 .

13.4.2 Summary

Criteria	met?	Comments
Scalable qubits	✓	
Initialization	✓	
Universal gates	?	
Coherence	✓	
Readout	✓	

Table 13.5: Summary of DiVincenzo criteria for flux qubits.

13.4.3 Outlook

The key problem for flux qubits is consistent fabrication - given the exponential dependence of the flux tunnel splitting on the Josephson energy, small tolerances in the Josephson energies lead to large errors of that term making targeted placement in frequency space that is needed for scaling up a major challenge. It is expected that they will show high gate fidelities also for two-qubit gates but there is doubt whether reaching level C will be attempted soon. Flux qubits are a great platform for quantum annealing but trailing behind transmons for gates.

13.5 Evaluation: Fluxonium

13.5.1 DiVincenzo-Criteria

13.5.1.1 Scalable qubits

Scaling is challenging due to the space required. Every single qubit needs to be build of up to hundred Josephson Junctions [VPS⁺14].

13.5.1.2 13.5.1.2 Initialization

Can be done via cooling.

13.5.1.3 Universal gates

Given the shear size of fluxonium, two qubit gates have not been attempted.

13.5.1.4 Coherence

- T_1 increase to values above 1ms [PGC⁺14, VPS⁺14].
- Theoretical proposal of flatsonium [SRDR17], with expected dephasing times of $T_\phi \sim 10$ ms.

13.5.1.5 Readout

Quantum non-demolition projective measurements within a time interval much shorter than T_1 , $5\mu\text{s}$ single-shot projective measurement [VPS⁺14]

13.5.2 Summary

Criteria	met?	Comments
Scalable qubits	?	
Initialization	✓	
Universal gates	×	
Coherence	✓	
Readout	✓	

Table 13.6: Summary of DiVincenzo criteria for fluxonium qubits.

13.5.3 Outlook

Given that lack of two-qubit gates, fluxonium is on level A. It is likely that it will catch up, but its sheer size makes it hard to imagine scaling in the sense of this study.

14 Ion traps

14.1 Description

Trapped ions are among the most promising candidates for the realization of quantum computers and quantum simulation. They are currently leading the field in both number of qubits and gate quality. The hyperfine levels in the ground state of the ions that encode the qubit have high coherence times and can be controlled with lasers.

14.1.1 Terminology

This platform is part of atomic, molecular, and optical (AMO) physics. The qubits are encoded in the quantum states of single electrons in the outer shell of a positively charged ion. Given their positive charge and the relative strength of achievable electrostatic forces, there is a well defined handle to trap ions. It is in fact possible to hold ions trapped very close to a well-defined position and spaced far enough so the ions can be individually addressed externally in order to drive gates with only limited and easy active cooling. This technology has been originally developed for a number of applications including metrology in the form of atomic clocks, which is compatible with the requirements of low error rate posed by quantum computing. It turns out that the trap is intimately related to the coupling of qubits hence it will be described under the qubit heading. Ion trap setups are typically placed on large vibration-isolated optical tables and operate in ultrahigh-vacuum systems and are addressed by lasers.

14.1.2 Qubits

14.1.2.1 Ions

Simple atomic ions with a single valence electron such as alkaline earths and particular transition metals are used for storing qubits. The choice of ions is driven by the desire to only have a single electron in the outer shell of the atom after ionization, hence ions used are earth-alkali metals (second column of the periodic table) such as Be^+ , Mg^+ , Ca^+ , Sr^+ , and Ba^+ [MK13][Bra17] as well as a few transition metals (Zn^+ , Cd^+ , Yb^+ , and Hg^+). The atoms have to be isotopically pure for the trap to work (see there) which is natural as all ion evaporation techniques can be made mass-selective. Choice of ion is driven by the preferred way to drive single qubit rotation (see there) and within that range by convenience of lasers and other devices used to manipulate the electrons. The amounts of material are minimal, scarcity is not a problem even if the preferred isotope is not the most abundant. Quantum simulation with trapped ions [BR12, KCK⁺10, JLH⁺14] is another driver for technology development.

There are limited applications of molecular ions currently in quantum simulation but not in quantum computing besides few very early basic investigations [DM12].

14.1.2.2 Trap technologies

While the electrostatic force is very strong on the atomic scale and allows for excellent access to the motion of the ion, it is not possible to hold an ion in a stable position due to electrostatic force alone, a fact known as Earnshaw's theorem in electrodynamics. Solutions to this problem involve the use of slowly time-dependent field that compensate for any instability by synchronization to the motion of the ion (this synchronization requires the ion mass to be known, i.e., the ion system to be isotopically pure). There are several common technologies to achieve this.

The linear Paul trap Named after Nobel laureate Wolfgang Paul [PS53], and reviewed in [Pau90], this trap holds ions in a linear crystalline (i.e., nearly evenly spaced) string in a quadrupole geometry. It is the current workhorse of ion trap quantum computing given its excellent optical access and the relatively small amount of metal needed. Its linear geometry and the resulting potential instabilities and vibrations of the ion string lead to the expectation that the current world record of 14 entangled ions [MSB⁺11] achieved by the Blatt group exhausts its capabilities - even if more ions can be trapped, it is unclear whether they can be coherently manipulated.

The Penning trap The Penning trap is an ion trap mostly used in high-energy physics applications. Although it allows to trap 2D-arrays of qubits for quantum simulation with ~ 300 ions [BSK⁺12], it lacks addressing and control of individual qubits and is not currently implemented for quantum computing [BKM16].

Surface traps In surface traps, all electrodes are part of a flat, segmented metallic surface providing potentials similar to that of a Paul trap [KMW02]. The segments can be controlled separately allowing to move and transport ions on top of the surface, hence not facing the limitation of 1D ion crystals imposed in Paul traps and allowing for much more flexible scaling and implementation of quantum algorithms [HLBH11]. As a general trend, linear surface traps can perform close to regular Paul traps and at least Sandia has reported high-fidelity two-qubit operations in a segmented trap at the APS March meeting in March 2017 [Mau07], but not published yet. More details are described below. Also, 2D surface traps are an active research field [SRW⁺14].

14.1.2.3 Single-qubit gates, encoding and control

There are multiple approaches to identify what states of the ion are used as qubit states. These are driven by the trade-off of long and stable quantum coherence and accessibility for control.

The optical qubit Qubit logic states are encoded in fine structure levels of the ion separated by frequencies that are in or close to the optical range. In order to maintain quantum coherence, the state that is higher in energy is metastable and can only decay through rare quadrupole transitions, with the flip-side being that rather strong lasers are required to drive single qubit gates. The most common ion in this application is ^{40}Ca and it is used by the Blatt group (Innsbruck University) and its many alumni. Fault tolerant topological encoding in Paul traps with 7 ions [NMM⁺14] and repetitive quantum error correction [SBM⁺11] have been demonstrated. Early basic studies attempt to use the Rydberg levels of an ion [FBS⁺15] (Rydberg physics is described below for neutral atoms).

The hyperfine qubit Qubit logical states are encoded in hyperfine levels of an ion. These are states of the electron that are differentiated only by their interaction with the atomic nucleus but otherwise are both part of the ground state manifold of the system. This makes them immune to decay and leads to enormous coherence time $> 1000\text{s}$ [BKM16], 50s without magnetic shielding [HAB⁺14]. Their transition frequency is given by the properties of the ion and is in the microwave frequency range (12.6 GHz for $^{171}\text{Yb}^+$). Gaps to hyperfine splitted excited electronic states lie in the optical regime, and allow for laser access for Doppler cooling to confine the ion near the bottom of the trap. In order to drive these microwave transitions with lasers, a scheme called Raman drive is used, which has the transition frequency as a difference of two (generally much higher) laser frequencies. The single qubit gates can be made ultra-fast $\sim 50\text{ps}$ [CMQ⁺10] and robust [RWL⁺18]. Recently, microwave-based quantum gates [MW01][LWF⁺17] were proposed and implemented in conjunction with static magnetic field gradients, making lasers unnecessary for single qubit gate operations (see under scale-up and technology).

14.1.2.4 Entanglement and two-qubit gates

Local operations The electrostatic Coulomb interaction between ions is used for entangling gates. Ion strings have, very much like a string of pearls connected by springs, collective modes of vibrations that can

be used to couple ions, even over long distances. The original Cirac-Zoller gate [CZ95] is used as well as the Mølmer-Sørensen gate [SM00] which has the advantage of being independent of the background vibration of the ion string. Qubit superposition is transformed to superposition of the ion's position mostly done with laser fields to perform two-qubit gates [BW08a]. The gate speed in a string of N ions is proportional to $N^{-1/2}$ implying slower gates in longer chains, with typical gate times $\sim 10\text{--}500\mu\text{s}$ [BHL⁺16]. These were improved using frequency modulation [LLF⁺18]. In surface traps, these gates would be implemented between short strings of qubits that are assembled using qubit transport for the gate of interest. Individual optical addressing and pulse shaping techniques enable error-correcting encoded qubits [MK13], with up to 100 qubits. Given these limitations, further scaling of that type of limitation requires functional surface traps. Alternatives rely on microwave control fields [OWC⁺11, TBJ⁺11] using microwave dressed states which are robust against magnetic field noise.

Very recently, high-speed (0.2 ns) gates that do not rely on the Coulomb interaction have been demonstrated with promising yet not leading fidelities [WMJM17].

Long-range operation by communication As an alternative to surface traps, long-range communication via photonic links is used for two-qubit gates in a distributed network of manageable-size Paul traps [MK13]. As most photonic methods (see there) these protocols involve post-selection, i.e., some operations are executed probabilistically. The overhead for this mode of operation is not problematic for scaling statements, albeit practically quite substantial. The conversion from stationary to flying qubits, i.e., quantum state transfer from an ion to a photon, has been demonstrated [SCB⁺13], as well as quantum teleportation between two ions via photons [OMM⁺09]. A detailed resource analysis of this approach will be done in AP 3 for general distributed quantum computing and then in AP 4 for this concrete platform.

14.1.2.5 Initialization and read-out

The electronic fine-structure ground state is easily initialized by simple cooling given its large energy separation. Initializing hyperfine states requires a laser-driven technique called optical pumping [Kas50, Kas67]. Readout is performed by attempting to drive a transition between one of the computational states and an auxiliary state and collecting the resulting fluorescence using photodetection (electron shelving technique) [MSW⁺08]. All these techniques are well established and reliable but require for additional lasers.

14.1.3 Scalability and technology

The first proposal of a quantum charge-coupled device (QCCD) [KMW02] for scaling beyond ~ 100 qubits assumed a large number of interconnected ion traps. It introduces interaction regions for logic operations and memory region for storage. Shuttling of ions between these smaller trapping zones requires exquisite control of shuttling ion positions. Also different types of ions are used for gates and transport [HHJ⁺09, BKM16]. A different approach is based on microwaves ion trap X-junction arrays [LWF⁺17] with different zones, microwave-based gate zones, readout zones and loading zones, using global laser fields for state readout and fast ion shuttling. Short distance ion transport was discussed in [BOV⁺09].

14.1.3.1 Trap heating

For scaling up and operation it needs to be noted that anomalous heating is a problem not completely understood yet [BKRB15]—ion traps in vacuum heat up in time hence perturbing everything that is trapped in them. Known to be an effect of the metal surface, this can be addressed using surface science, for example with in-situ ion bombardment [HCW⁺12], but also by minimizing the amount of metal used in the trapping system. Reduction of noise can also be achieved by using cryogenic traps [LGA⁺08]. Inevitable Johnson noise [Joh28, Nyq28] can be almost completely canceled in high-temperature superconducting surface ion traps below the critical temperature.

14.1.3.2 Lasers and temperature

The precision of the ion trap setups requires high stability of the ambient temperature, below 0.1 K. This is in parts because of the important role of the trap that cannot tolerate any thermal expansion. The sophisticated techniques behind the different steps of these setups, including gate drive, optical pumping and readout require a multitude of lasers at different frequencies. Also, the motion of the ions needs to be constantly cooled. Other than in solid state setups, cooling here is not done by heat exchange with a coolant such as Helium, rather, by further laser-based techniques [Phi98]. This leads to very complex setups containing many lasers all of which consume power (not a big problem per se as this only grows slowly with the number of qubits) but most importantly all of which dissipate heat. This requires, in most cases, to operate lasers in a separate room, contributing to the complexity of the system. A separate problem is anomalous heating of the trap, which is only partially understood.

14.1.3.3 Vacuum systems

Ion traps need to be operated in ultrahigh vacuum, as the motional degrees of freedom used for two qubit gates cannot tolerate collision with residual gas atoms. While small ultrahigh vacuum units as part of the research infrastructure are reliable routine equipment, new challenges arise when these grow to large volume.

14.1.3.4 Crosstalk

When driving operations, qubits need to be clearly addressed, i.e., it is important to make sure that controls are qubit specific. Absence of this capability is called cross-talk. Achieving this with lasers requires to space ions by more than a focus area hence preventing dense packing of the ions. Solutions include avoiding crosstalk via different types of atomic species, such as $^{171}\text{Yb}^+$ and $^{138}\text{Ba}^+$ [BKM16]. They also include selecting hyperfine transitions by frequency by putting the ions into a strong static magnetic field.

14.1.4 Research groups

There are many ion trap groups given the clean quantum properties of the system. This description lists the groups that are active in trapped ions for quantum computing applications only, not sensing or atomic clocks. Given the high complexity of these setups, these are grouped in relatively few networks.

14.1.4.1 Wineland-Monroe network

David Wineland, NIST Boulder David Wineland (Nobel prize 2012) is the leader of the ion storage group at the US metrology institute, NIST. He has been a leader and pioneer of ion trap-based quantum computing mostly emphasizing hyperfine qubits, and a true authority in the field. He has first proposed surface ion traps. His lab is pursuing a quantum computing activity, co-led with Dietrich Leibfried but not strongly externally networked, in part because they are part of the US government.

Christopher Monroe, JQI Maryland Chris Monroe is a postdoctoral alumnus of David Wineland and works at the Joint Quantum Institute, a research unit between NIST Gaithersburg and the University of Maryland. Monroe is the key scientist (but not the formal leader) of a large concerted effort to build a networked ion trap quantum computer. Monroe was the first to perform quantum gates in an ion trap [MMK⁺95]. He is also lead scientist of IonQ, a company that tries to commercialize ion trap quantum computing. The CEO of IonQ is David Moehring, also a Wineland Alumnus and previous project manager at IARPA.

Jungsang Kim, Duke University Kim is an optical engineer and pioneer of integrated optical switches, a key component in making an ion trap quantum computer compact. He is managing the Monroe - centric research network.

Kenneth Brown, Georgia Tech Ken Brown was trained as a theorist in a chemistry department (PhD with K. Birgitta Whaley at UC Berkeley) but trained as an ion trap experimental physics as a postdoc with Isaac Chuang. Brown is interfacing theory and experiment for the Monroe team and is lead author on important architecture papers.

Dmitry Matsukevich, CQT Singapore Matsukevich is an alumnus of the Monroe group and develops molecular ions for quantum computing.

Jonathan Home, ETH Zürich Jonathan Home was also trained as a postdoc with Wineland but currently mostly collaborates with Rainer Blatt's network. He specializes on various surface trap developments as well as advanced control techniques for efficient gates.

Winfried Hensinger, Sussex Trained in various areas of AMO physics, Hensinger came to ion traps through a postdoc with Chris Monroe. He is pursuing the radical approach of eliminating as many lasers as possible if not all from ion trap setups, replacing them by microwave sources. His recent work on proposing an architecture serves as a pre-studied blueprint for the scaling of ion traps.

Peter Maunz, Sandia National Laboratories Initiated to ion traps through Chris Monroe, Maunz is a specialist in the fabrication of clean surface traps that are the first of their kind to allow high-quality gate operations. He is also involved in process verification research through his collaboration with Robin Blume-Kohout.

14.1.4.2 Blatt network

Rainer Blatt, University of Innsbruck Rainer Blatt is a world leading experimental physicist at the University of Innsbruck and the Austrian Academy of Sciences. He leads a large effort in building a quantum computer based on going from Paul traps to surface traps. This school mostly uses optical qubits. In his group, senior experimentalists working on ion traps include Christian Roos. The strong application push towards scalable quantum computing has been taken over by **Thomas Monz** who is leading a flagship project as well as a company, AQT.

Isaac Chuang, MIT Originally an NMR experimentalist and theory (notable author of the standard textbook in the field) has developed into an ion trap experimentalist with a strong theoretical vision through collaboration with Rainer Blatt. He interfaces theory and experiments in a unique way.

Ferdinand Schmidt-Kaler, Mainz The Schmidt Kaler group has a major effort in trap development, specifically surface traps for the Rainer Blatt operation. Ulrich Poschinger is a senior scientist in that group.

Hartmut Häffner, UC Berkeley Hartmut Häffner's group focuses on high-quality surface traps and traps in ring geometry, applied to quantum simulation and quantum computing.

Others The Blatt school has spun off many other ion trap experimentalists who are not currently pursuing ion traps for computing but for other applications - but would be able to return to computing. Among them are Matthias F. Brandl (Innsbruck), Christian Ospelkaus (Hannover and PTB), Tobias Schaetz (Freiburg), Michael Drewsen (Aarhus), and Jürgen Eschner (Saarbrücken).

14.1.4.3 Others

Christoph Wunderlich, Siegen Trained through high-fidelity trapped-ion quantum logic using near-field microwaves [HSA⁺16] ion trap pioneer Peter Toschek in Hamburg (who has worked on fundamental science but not on quantum computing), Wunderlich is a pioneer of using microwaves to manipulate ions.

Andrew Steane, David Lucas, Oxford University Andrew Steane is an experimentalist with strong theoretical interest, famous also for his contributions to error correction. He runs his laboratory jointly with David Lucas. They hold the world record for lowest error in ion trap two-qubit gates [BHL⁺16] and got very close to that number in microwaves.

14.2 Evaluation

14.2.1 DiVincenzo-Criteria

14.2.1.1 Scalable qubits

Largest sample reported contains 53 trapped ions hyperfine $^{171}\text{Yb}^+$ in a linear rf Pauli trap [ZPH⁺17]. Larger samples need multiple traps and shuttling operations between them [KMW02]. The required traps are being developed but do not allow for competitive gate quality (yet). Ions can escape the trap, therefore continuous replacing needed. Alternative: Distributed quantum computing with temporal communications overhead.

14.2.1.2 Initialization

- Default initialization method is Doppler cooling and additional sideband cooling $^{43}\text{Ca}^+$ [SBT⁺17].
- Doppler and sideband cooled and optically (re)pumped $^9\text{Be}^+$ ion(s) [BWC⁺11, GTL⁺16], sideband cooling [BHL⁺16].
- $^{40}\text{Ca}^+$ qubit reset in 50(10) μs with error 5×10^{-3} , with expected values in parenthesis [BXN⁺17].

14.2.1.3 Universal gates

- Reasonable gate times in near time future based on microwave control lie around $t_{g,1} = 2.5\mu\text{s}$ for single-qubit gates and $t_{g,2} = 10\mu\text{s}$ for two-qubit gates [LWF⁺17] with $^{131}\text{Yb}^+$.
- Entangling gates with $t_g = 1.6\mu\text{s}$ and 0.998 fidelity using tailored laser pulses in a linear Paul trap and hyperfine qubits of $^{43}\text{Ca}^+$ ions [SBT⁺17]. Higher fidelities are possible for gate durations of 30 μs and 100 μs . Previous results in $^{43}\text{Ca}^+$ show single-qubit gates in 7.5 μs and two-qubit gates in 100 μs with fidelities 0.999934 and 0.999 (Bell-state) [BHL⁺16] in room-temperature ion trap. MS-gate duration $t_g = 3.25\text{ms}$ with 0.997 Bell-state fidelity [HSA⁺16]. Limited by off-resonant photon scattering [KRS⁺17] in $^{40}\text{Ca}^+$.
- Gate duration of $t_g = 30\mu\text{s}$ for the MS-gate with Bell-state fidelity 0.9992 [GTL⁺16] and $^9\text{Be}^+$ ions.
- $^{171}\text{Yb}^+$ in linear Paul trap allow fast single-qubit gates $t_g = 50\text{ps}$ with 0.99 fidelity [CMQ⁺10]. In the 5-qubit comparison typical gate times are 20 μs for single-qubit and 250 μs for 2-qubit gates [LMR⁺17] The fidelities for single- and two-qubit gates are typically 0.991 and 0.97, respectively.
- Single-qubit gate times approximately at 5(1) μs with error 5×10^{-5} and two-qubit entanglers at 40(15) μs with 1×10^{-2} (2×10^{-4}) for single species and 80 μs with errors 3×10^{-2} (4×10^{-4}) for dual species ($^{40}\text{Ca}^+$ and $^9\text{Be}^+$), with expected values in parenthesis [BXN⁺17].
- single-qubit gates RB result of 0.99998 [BWC⁺11] with $^9\text{Be}^+$ in a surface trap (hyperfine qubit), cooled down to 4.2K and 0.999962 [GTL⁺16].
- Single-qubit π -rotation in 10 μs and RB of 5.1×10^{-5} [KRS⁺17].

14.2.1.4 Coherence

- $T_2 = 0.38\text{s}$ [BWC⁺11] with $^9\text{Be}^+$. $T_R = 1.5\text{s}$ [GTL⁺16].
- $T_2^* = 50\text{s}$ [HSA⁺16] for $^{43}\text{Ca}^+$.

- Coherence times in segmented Paul trap and $^{40}\text{Ca}^+$ ions enhanced by dynamical decoupling to 1.1s [KRS⁺17].
- Coherence time of 1000s for $^{171}\text{Yb}^+$ [FSLC97]. $T_2 \approx 0.5\text{s}$ magnetic field noise [LMR⁺17] in hyperfine ground-level qubits. Suppressing magnetic-field noise for improvisations.

14.2.1.5 Readout

- Individual qubit measurement of $^{171}\text{Yb}^+$ with nearly 99% efficiency [ZPH⁺17].
- Measuring the fluorescence signal [GTL⁺16] with $^9\text{Be}^+$ single-qubit readout fidelity is 99.7(1)% for state $|0\rangle$ and 99.1(1)% for state $|1\rangle$
- Average readout fidelity for an entire 5-qubit state is 95.7(1)% ,limited by because of crosstalk [LMR⁺17].
- Measurement of $^{40}\text{Ca}^+$ in 400(30) μs with error 10^{-3} [BXN⁺17].

14.2.2 Summary

Criteria	met?	Comments
Scalable qubits	✓	
Initialization	✓	
Universal gates	✓	
Coherence	✓	
Readout	✓	Focus on faster readout

Table 14.1: Summary DiVincenzo criteria for trapped ions.

t_i	t_1	t_2	t_M	p_1	p_1	p_2	p_M	T_2
50 μs	2 μs	10 μs	30 μs	$5 \cdot 10^{-3}$	$5 \cdot 10^{-5}$	10^{-3}	10^{-3}	1s

Table 14.2: Optimistic assumptions for different error rates and operation times combining the best reached values for each operation. This does not describe a current available setup, but shows what is in principle possible right now or in near future. Times are initialization, 1- and 2-qubit gate and measurement time. Probabilities are error probabilities for the respective processes. A surface code cycle contains 4 two-qubit gates, 2 one-qubit gates, measurement and initialization as well as classical processing.

14.2.3 Fault-tolerance extrapolation

Early tests of error correction are already performed, like

- assessment of trapped ions for implementation of the 7-qubit color code [BXN⁺17], using two ion species, $^{40}\text{Ca}^+$ (qubit host) and $^{88}\text{Sr}^+$ (sympathetic cooling), table of recent and anticipated fidelities and gate times included
- repetitive phase flip correction with $^{40}\text{Ca}^+$ ions [SBM⁺11]

but no actual surface code implementation has been performed. Current gate times make it hard to imagine implementing longer algorithms fault-tolerant in reasonable times. Although the lower error rates (compared to Transmon qubits) allow for much lower distances of around $d = 30$ in the Clifford part for the considered algorithms, the product of required time and qubits ends up in the same range as for Transmons. Additionally, computation times need to be much higher, since the allowed amount of parallelization is reached earlier. The smallest dlog algorithms, as shown in Figure 14.1 are already only possible in at least 100 days, larger key sizes would require even longer runtimes. Factoring is not possible at all in the desired

runtimes of up to 100 days and therefore no plot is shown. A proposal for fast gates in ion traps has been put forward in [SJ09].




Figure 14.1: Double-logarithmic plot showing possible system sizes when implementing dlog algorithms for different runtimes, with experimental values from table 14.2. (rather than the uniform error probabilities used in the other chapters, see chapter 7.4.1 for clarification). The purple line gives the number of qubits for the time-optimized limit, i.e., the minimal required runtime due to the T depth when having maximal parallelization. Implementations above the line would violate the temporal order of T gates and are thus not possible.

14.2.4 Analysis and outlook

Trapped ions are extremely clean and flexible controlled quantum systems which in the context of atomic clocks have reached metrological precision, hence providing an excellent platform for high-quality quantum operations. They are currently the most consistent platform in reaching the error correction threshold and thus most definitely in category C. This platform advances continuously. It is currently working on overcoming a steep scaling obstacle: Changing trap technology from Paul to surface traps while maintaining high operational quality.

14.2.4.1 Resources: Space and time

Both measurement and gate times in trapped ions are quite long. While due to their excellent coherence this does not impede high-fidelity operation (level B), it does affect overall algorithmic performance and, given the effectiveness of error correction, eats up the lower overhead. Thus, the needed processor sizes for cryptographic tasks are comparable if not larger than for Josephson qubits. A further challenge is the size of the surrounding apparatus: While ions are small, ion traps are not, hence again, the required machine is comparable to Josephson qubits.

14.2.4.2 Technical feasibility

An excellent analysis that we agree with has been described in [LWF⁺17] 14.1.3. Brute-force scaling up of an ion trap quantum computer to the required sizes for attacks on cryptography would lead to a machine the size of a soccer field and include the challenge of engineering a large ultrahigh vacuum system. Its power consumption (dominated by trap currents) would be comparable to that of a supercomputing facility. Again, this would be a focused and visible research project comparable to the Apollo or Manhattan programs. Rare materials are not needed. As industrial uptake of ion traps is slower than for solid-state platforms (mostly for reasons related to leveraging existing computer technology), crypto-related ion trap research will likely profit less from industrial activities than solid-state programs.

14.2.5 Trapped Rydberg Ions

While brand-new and currently incomplete, it is worth pointing out the alternative route of using Rydberg ions. On a $^{88}\text{Sr}^+$ ion in a linear Paul trap the π -phase gate has been determined by quantum process tomography with a fidelity of 0.78 [HPZ⁺17], using double STIRAP (sequence Stimulated Raman Adiabatic passage) and a lifetime of the Rydberg state of $\tau_{42S} = 2.3\mu\text{s}$.

15 Semiconductor platforms

15.1 Description

Semiconductors are a natural platform for scalable quantum computing [Hei03]. On the one hand, standard semiconductor devices are in fact simple quantum technologies—transistors, diodes, and semiconductor lasers are based on engineering a quantum material and its energy spectrum. On the other hand, the semiconductor industry provides enormous experience in all aspects of semiconductor fabrication and has shown the potential for large-scale integration of classical processors.

Nevertheless, semiconductor platforms are not yet consolidated to the extent that others are, hence providing a complex community with varied approaches.

15.1.1 Categorization and background

We will categorize semiconductor based approaches along three dimensions

- the degree of freedom is used to encode quantum bits—charge/orbital degrees of freedom, electron spin, nuclear spin, or coupled degrees of freedom
- the method of confinement of the continuous degree of freedom of semiconductor material into discrete logical elements for qubits—quantum dots, single defects, topological mechanisms
- the material system that is used—elementary semiconductors (C, Si, Ge), III-V Materials (GaAs family), or more exotic semiconductors

These dimensions are correlated as will be described.

15.1.1.1 Degrees of freedom

Semiconductors are complex multi-faceted systems, hence different degrees of freedom can be used:

Electron charge / Orbitals A straightforward way to encode quantum information is to use the position degree of freedom of a quantum particle—electrons or holes in semiconductors. These are often referred as *charge qubits* as these particles carry electric charge. They are straightforward to control as they couple well to electromagnetic fields. For the same reason, they also couple well to unwanted degrees of freedom of the semiconductor materials hence strongly limiting coherence times in particular due to charge noise similar to Josephson charge qubits. Pure charge qubits have shown simple single-qubit operations very early [HFC⁺03],) but are hardly investigated any more [CRFG17, THWZ16, WKS⁺16]. Some advanced spin-qubit designs contain use of the charge degree of freedom (through the Pauli principle or the spin-orbit interaction) which again, makes them somewhat sensitive to charge noise. When investigating charge qubits, it is important to evaluate their sensitivity to decoherence by lattice vibrations (phonons) which is very strong in piezoelectric materials.

Electron spin Electrons have spin—an intrinsic angular momentum associated with a magnetic moment. When bound in materials, the orbital angular momentum and spin combine (see below). Spin is a natural two-state system that can be addressed using magnetic fields. Note that spin resonance in molecules is described in a different chapter.

Electron spin states are sensitive to spin environments, most notably nuclear spins. These couple magnetically to the electron spin and move slowly, given the much larger mass of nuclei compared to electron spins, driving protocols, device designs, and choice of material.

Nuclear spin The nuclear spin degree of freedom resembles that of the electron with a few key differences: It is extremely well isolated from its environment thus very quantum coherent. It is usually not mobile and much more difficult to engineer. Nuclear spins play a role as auxiliary degrees of freedom in some semiconductor platforms where they are typically manipulated through interaction with an electron.

Coupled degrees of freedom Some semiconductor implementations use more complex degrees of freedom. Specifically candidates for topological quantum computing encode quantum information non-locally in complex many-body states, including the 5/2 fractional quantum Hall effect state, a collective state of electrons and magnetic field in a thin film in a very strong magnetic field, and Majorana fermions in topological systems that originate from the combination of semiconducting materials with a strong spin-orbit interaction and superconductors.

15.1.1.2 Method of confinement

Electrons in semiconductors are generically free to move in the material. This stands in the way of quantum computing as this makes it impossible to address qubits selectively to drive single-qubit gates. It also leads to an undesired continuous energy spectrum that is a source of decoherence and disallows initialization of qubits by cooling. Defining charge qubits would be impossible in this context, operating the other qubits would be equally impossible.

This means that the electrons need to be confined to a small structure in space, small enough to address them selectively. Depending on the details of the confinement potential, the energy splitting of orbital states of electrons confined to a small region of size L is $\Delta E \propto m_*^{-1} L^{-2}$ where m_* is the effective electron mass, described later. There are three main ways of confinement:

Quantum dots A quantum dot is an artificial nanostructure that confines electron motion. They are either defined by nanofabricated metallic gate structures that repel electrons from regions where they are not wanted (lateral quantum dot) or produced free-standing by a materials-growth or etching procedure. The former are very flexible as the gate voltages can be used to fine-tune dot properties, the latter can be made smaller.

It needs to be noted that the terminology of quantum dot has several meanings in literature and is quite fashionable, not all of them are suitable for quantum computing. In particular *colloidal quantum dots* are molecules in solution for optoelectronics, unrelated to quantum technology. *Self-assembled* or *optical quantum dots* are ultra-small semiconductor structures used in optoelectronics and quantum communication but not in quantum computing.

Donors and defects A key component of all modern semiconductor technologies is doping the material with atoms of other materials to change the effective properties of the material. For quantum computing functionality, it is possible to use the electrons around single, isolated dopants or defects. Electrons there are bound relatively tightly. Again, several realizations of this scheme exist and will be described below. The common challenge to those is that atomically precise placement of donors and scalable control of qubits is a frontier in nanotechnological fabrication.

Topological effects An attractive way to reach confined elements are topological effects. There, interactions between many degrees of freedom create a bound state that, in principle, is stable against any external perturbation: Locally perturbing the state of the system does not change the topological genus of the state hence not compromise quantum information.

15.1.1.3 Materials

Semiconductors are based on half-filled electronic shells, hence there is a restricted choice of semiconductors. This state can be reached either by using elements from the fourth column of the periodic table (C, Si, Ge) or compounds that are in the fourth column on average only (GaAs, InSb). Key properties include

- **Electron effective mass:** Effective mass is a property describing the motion of electrons that are exposed to the crystal lattice, it can substantially deviate from the bare electron mass. Smaller effective masses give larger energy splittings or, given target values for energy splittings, allow to reach those values in larger, more manageable confinement length scales.
- **Valley degree of freedom:** This is a property of the motion in the crystal that identifies that there are multiple electron states that behave similarly [YC10] which makes it difficult to isolate a single qubit (rather than a single qubit per valley). As valley degeneracy is a result of symmetry of the crystal it can be removed by breaking that symmetry through mechanical strain.
- **Content of nuclear spins:** Their presence is a strong perturbation to electron spin qubits through dipolar and contact interaction.
- **Piezoelectricity:** Electric fields produced from mechanical deformation. This is a property of the compound and the crystal structure. Piezoelectricity implies strong interaction of lattice vibrations with the orbital degree of freedom and is serious detriment to quantum coherence.
- **Spin-Orbit interaction:** In principle, spin and orbit are coupled degrees of freedom but strength and relevance of this interaction depend on the material. In general, this interaction is strong in semiconductors made from heavy elements. If the spin degree of freedom is used to encode qubits, spin-orbit interaction is detrimental as it contributes to decoherence via the orbital degree of freedom. On the other hand, it is necessary to have strong spin-orbit interaction for topological qubits.

Gallium Arsenide GaAs is a III-V semiconductor (third and fifth column) with traditional applications in opto-electronics. It has no valley degree of freedom and low effective mass. Every nucleus of every isotope has nuclear spin. GaAs is piezoelectric and has weak spin-orbit interaction at least in lateral quantum dots. Initial confinement is reached by sandwiching the material with AlGaAs.

Silicon Si is a semiconducting element that is the backbone of traditional computers. It has a valley degree of freedom and large effective mass. It can be isotopically purified to be nuclear spin free. It is not piezoelectric and has only very weak spin-orbit interaction. Initial confinement is reached by sandwiching with SiGe which also applies the required strain to remove its intrinsic valley degeneracy.

Carbon Carbon comes in a number of allotropes. While Graphene is a great nanomaterial, there is no viable experiment in Graphene for quantum computing. Carbon Nanotubes and Fullerenes can be used to manage molecules (see chapter on molecular quantum computing). The most interesting allotrope for quantum technologies is currently diamond, wide-bandgap semiconductor with the same crystal structure as Si. Defects in diamond, specifically the Nitrogen-Vacancy (NV) center are a leading system for quantum technologies, mainly in quantum sensing. An approach to quantum computing with NV-centers is described below. Diamond is not piezoelectric and can be made nuclear spin free.

Indium Antimonide InSb is a heavy III-V semiconductor that can be easily grown in nanowires. Its strong spin-orbit interactions, alongside InAs and Si/Ge core-shell nanowires, make it the leading platform for topological qubits [SLTDS10, LSDS10, PSV⁺12, vWPB⁺13].

15.1.2 Concrete semiconductor platforms

Based on the background and classification above, this section describes the known and experimentally relevant semiconductor platforms.

15.1.2.1 Lateral quantum dots: Loss-DiVincenzo

In this scheme [LD98], lateral quantum dots containing one electron hold the quantum information encoded in the spin of that electron. Placed in a strong magnetic field produced by a superconducting permanent magnet to produce a Zeeman splitting that is large enough to allow initialization by cooling. Single-qubit gates are performed by electron spin resonance, i.e., by shining in microwave fields resonant

with that Zeeman splitting. Two-qubit gates rely on the exchange interaction between neighboring quantum dots. This interaction can be tuned by an electrostatic gate between the dots that controls the wave function overlap. Readout is based on spin-to-charge conversion and then fast charge readout by a quantum point contact [EHWvB⁺04, TED⁺05].

In materials with nuclear spin such as GaAs, this scheme is very sensitive to noise. Some progress has been made by noise programming, i.e., by polarizing the nuclear spins by electronic means. A further challenge is related to the addressability of qubits: Microwaves even in near field cannot be focused on only one quantum dot and not its neighbor. The most popular solution [OPLTT12] is to prepare a micromagnet next to the dots that produces a magnetic field gradient, hence changing the resonance frequency of the dots and selecting quantum dots by that frequency.

15.1.2.2 Lateral quantum dots: Singlet/Triplet

In order to combat nuclear spin decoherence, the S-T qubit was invented at Harvard [SDP⁺12]. It physically encodes a single qubit in two spins in two neighboring quantum dots and only uses the unpolarized subspace spanned by $\{|01\rangle, |10\rangle\}$ that are not sensitive to the magnetic field. In this subspace, tuning the exchange energy provides single-qubit gates. Two-qubit gates rely on electrostatic interaction between double-quantum-dots, all other techniques resemble the Loss-DiVincenzo technique. Demonstrations of two-qubit gates have not been reported so far. One of the challenges here is scaling given this very complicated architecture. Also, the Pauli principle makes the orbital part of the quantum states between $(|01\rangle \pm |10\rangle)/\sqrt{2}$ distinguishable, making this qubit sensitive to charge noise, albeit weaker than a pure charge qubit.

15.1.2.3 Lateral quantum dots: Encoded universality

Taking the idea of noise protection further, [BKLW00] and [DBK⁺00] shows that even simple interactions that are not universal for quantum computing can be universal on a decoherence-free subspace. This requires to encode a logical qubit in three physical qubits. This is attempted in lateral quantum dots [LTD⁺10]. Given the move of the community to Silicon, which can be made nuclear spin free, this has become less relevant. During the writing process, an example of symmetric triple dot has been operated [MMN⁺17b].

15.1.2.4 Lateral quantum dots: Charge qubits

Originally, pure charge qubits where only the position of an electron in a double quantum dot is used were also pursued [HFC⁺03]. These are sensitive to phonons in piezoelectric GaAs and to slow charge noise in Si and have been abandoned.

15.1.2.5 Quantum dots: Status of Si in these categories

Recently, there has been a dramatic change in overcoming the nanofabrication challenges in Si quantum dots. Currently, both Loss-DiVincenzo and S-T qubit are being pursued. Si profits from the vast experience of the semiconductor industry, which is engaged (Intel) at TU Delft and at Grenoble through LETI, a pre-industrial foundry operated by CEA. Many groups have already demonstrated single-qubit operations. Two-qubit gates are the current research frontier.

15.1.2.6 Si-Donors: Kane proposal, historic work

Implantation of single donors to provide discrete states is primarily done in Silicon. The Kane Proposal [Kan98, Lan95] provided an early blueprint for this architecture, using P donors as long-lived quantum memory, hyperfine interaction to interact with electron spins used as processing elements, and exchange interaction between electron spins to make two-qubit gates. Control is applied by metallic gate fingers similar to those used in quantum dots to tune interactions and drive single-qubit gates. The Kane proposal states that single-qubit gates are performed through magnetic resonance, and two-qubit gates are implemented via the exchange interaction.

The Kane proposal has been a major effort in Australia in the 2000s. It required to newly event solutions for almost all its components and despite good progress came to a slowdown. It is still pursued at the Centre for Quantum Computation and Communication Technology (CQC2T) centered in Melbourne and Sydney (see <http://www.cqc2t.org/>)

This attempt has caught up since 2010 largely due to adapting techniques first developed in quantum dots to this system, so by making more complex gate structures than those envisaged by Kane. This is largely driven by Andrea Morello at UNSW. On chip microwave transmission lines up to 50GHz are used to drive the system. A SiMOS (Silicon Metal On Semiconductor) Single Electron Transistor can be used to read out charge into which the electron spin (and by one more conversion) the nuclear spin is converted. The coherence times for electrons are ~ 1 s and for the nuclei ~ 30 s largely limited by complex spin-spin interaction [TMS⁺15]. The resulting operation of the nuclear degree of freedom as quantum memory was shown in [FSL⁺17]. Single-qubit gate infidelity errors around 10^{-4} were achieved for electron and nucleus [MLS⁺15] and similar values have been verified by gate set tomography (GST) [DMBK⁺16], which also verifies Bell's inequality. Initialization by relaxation is very slow given those relaxation times but can be improved by Bayesian methods. The combined fidelity of initialization, gate and readout ("Triple-Triple") is 99.9%.

The main frontier in this field is to achieve two-qubit gates. Qubit addressing and coupling was first envisaged in [KHDS01] and robust two-qubit gate proposal can be found in [KLHM14], with strong coupling of polaron pairs [DMT⁺14] and long distance coupling is achieved with resonators [TMS⁺15]).

15.1.2.7 NV-Centers: Distributed quantum computing

(Artificial) diamond is an unusually pure and stable wide-gap semiconductor. Its color centers provide a natural trap for electrons with discrete states usable as qubits. The most studied is the NV (Nitrogen-Vacancy) center, i.e., a defect created by replacing a carbon atom with a nitrogen atom and creating a vacancy in the lattice. The primary quantum technology in which color centers are used is quantum sensing, followed by quantum communication. The simplicity and stability in NV centers also makes them quantum computing candidates. A single NV-center contains an electron and a nuclear spin hence presenting a natural two-qubit register, that has been successfully operated [NBS⁺10, DCJ⁺07, RCB⁺11].

This poses a difficult challenge to scale beyond two qubits. Two proposals are pursued: Distributed non-deterministic quantum computing connecting remote NV centers by photonic links [NTD⁺14, TCT⁺10], and implantation of NV centers at suitable positions, similar to spins in silicon [YJG⁺12, JGBW⁺16]. This is more difficult than in silicon given that it cannot rely on experience of the semiconductor industry. An other approach is the coupling via a superconducting transmission line resonator, as recently demonstrated in Vienna [ANP⁺17]. Although coherent coupling of two NV centers was observed, these first experiments could not say anything about the degree of entanglement that can be reached between the spins yet. The company Element 6 is a foundry for diamond that leads efforts towards better fabrication.

15.1.2.8 Topological qubits: Majorana fermions in nanowires

Realizing topological bound states that promise resistance to noise and decoherence requires well fine-tuned models. The currently most promising route is based on a one-dimensional system with strong spin-orbit interaction that create a topologically nontrivial band structure with superconductivity. This is realized in InSb or InAs Nanowires [SLTDS10, LSDS10, PSV⁺12, vWPB⁺13]. These are not naturally superconducting but superconductivity can be induced by covering with a superconducting layer, a phenomenon known as the proximity effect. There is large number of theoretical descriptions of this system and various strategies to build a quantum computer.

On the experimental side, two groups have claimed to have observed Majorana fermions [PSV⁺12, vWPB⁺13, AHM⁺16b]. The evidence for this observation comes from observations of charge transport characteristics through such a nanowire and it is still controversial whether the signatures can be unequivocally assigned to Majorana states. Single- and two-qubit gates have not been observed but are currently attempted at both places.

Similar evidence for Majorana Fermions has now appeared in planar junctions instead of Nanowires [FWS⁺18].

Nevertheless, this field receives strong attention. Much of the theory in this field has been developed by Microsoft Station Q [DSFN06, NSS⁺08, Wil09] at UC Santa Barbara and the current results have motivated a large investment in experimental research in Delft and Copenhagen.

This field is at a decision point. Showing single qubit rotations in Majorana fermions and investigating their quality will help decide whether the coherence of this system is comparable to a single *physical* qubit in other platforms or already is comparable to that of an error-corrected *logical qubit*.

15.1.2.9 Topological qubits: Other candidates

While this is the most promising platform for topological qubits at this point in time, others have been attempted. Some of these are not based on semiconductors:

- The $5/2$ fractional quantum Hall state occurs in clean two-dimensional electron systems in very clean GaAs heterostructures [Wil13] where $\nu = 5/2$ is the filling factor, the ratio of the number of electrons to magnetic flux quanta. It is believed to carry quasi-particles with fractional non-Abelian statistics that could act as topologically protected qubits. While this state has been robustly observed, evidence for this type quasi-particles remains inconclusive.
- Half-integer vortices in certain unconventional (p-wave) superconductors, specifically strontium ruthenate (Sr₂RuO₄) are believed to contain non-Abelian anyons as bound states [SZH⁺16]. While these have been observed [JFV⁺11] and scanning probe measurements seem to indicate that they may contain these quasi-particles, again, manipulation towards single qubit gates and coherence measurements have not been achieved [DSNT06].

Josephson junction arrays can be used as topologically protected qubits [Kit06, IFI⁺02, DFII05]. Recent work shows that even a small unit cell can in fact be protected [BPIG14] and reaches coherence time compatible with, but not better than the best transmons. When further scaling up, this method will lead to circuits of complexity similar to a transmon-based surface code, which in itself contains ideas of topology.

15.1.3 Operational issues

15.1.3.1 Materials, fabrication, charge noise

Some of the materials involved are difficult and unsystematic to fabricate, i.e., in mass production there is a question of yield. GaAs films are grown by molecular beam epitaxy (MBE) by specialized growers [PWW⁺05, PTR⁺17, BHP⁺14, YSP14, RFB⁺10]. Silicon fabrication leverages techniques from commercial fabrication and, based on limited experience, is more reliable. For spin qubits, isotopic purification is required which is a laborious and expensive but unproblematic process. Nanowires for topological quantum computation are grown in process that, again, only a small number of groups master. The process of contacting them is currently not scalable but is not the primary concern of these systems.

15.1.3.2 Cold electronics, complexity

Issues related to cryogenics as well as to microwave resemble those in Josephson qubits. Cold CMOS logic is pursued as a cold control layer.

15.1.4 Research groups

Jason Petta, Princeton Jason Petta works on nanofabrication of several semiconductor quantum information platforms, including double quantum dots, nanowires and ion implantation for precise positioning of NV centers in diamond. He is considered to be one of the world leaders in Si-based quantum dots.

Mark Eriksson, Wisconsin Besides research on superconducting qubits and neutral atoms, the Wisconsin Institute for Quantum Information is an important player in the field of semiconductor qubits, with Mark Eriksson as primary investigator working on silicon-based nanostructure fabrication and measurement of quantum dots. Other important researchers of the institute working on semiconductor qubits include Robert Joynt, Susan Coppersmith and Mark Friesen in theory.

Malcolm Carroll, Sandia National Laboratories Malcolm Carroll is principal investigator of the Quantum Information Science and Technology (QIST) Grand Challenge at Sandia, which aims to produce a silicon-based spin qubit, integrated with CMOS technology and appropriate error correction circuits. Sandia is working on a broad range of technology for national security issues for the United States of America, their research also includes other quantum computation platforms like trapped ion technologies (Peter Maunz).

Center for Quantum Devices (QDev), Copenhagen Led by Charles Marcus (previously of Stanford and Harvard Universities), Qdev is part of the condensed matter physics group at the Niels Bohr Institute and does both theoretical and experimental studies on solid state qubits, nanowires and topological quantum systems. Marcus focuses on topological qubits and is currently moving to Microsoft. Other key researchers in quantum computing there include Ferdinand Kuemmeth. They have been one of the leaders in GaAs and are currently moving towards silicon.

Centre for Quantum Computation and Communication Technology, Australian Research Council Centre of Excellence The Centre for Quantum Computation and Communication Technology is a team of Australian researchers from nine different universities developing a scalable spin-based quantum computing architecture in silicon.

Michelle Simmons is Professor at the University of New South Wales and Director of the Centre for Quantum Computation and Communication Technology. She is doing pioneering work on atomically precise devices for quantum computing in silicon and germanium.

Another important member is Andrea Morello, Professor of Quantum Nanosystems at the School of Electrical Engineering and Telecommunications and the head of the quantum spin control group. His

research aims at building a quantum computer based on single spins in silicon, and is at the forefront of quantum technologies. This includes the first demonstration of single-shot spin readout in silicon [MPZ⁺10], the first spin quantum bits based on the electron [PTD⁺12] and the nucleus [PTD⁺13] of a single phosphorous atom in silicon. His group is also developing advanced techniques to observe and control the interaction between two qubits, including quantum logic gates. His group currently holds the record quantum coherence time for solid-state qubits [MDL⁺14]. Their collaborators include M. Carroll (Sandia National Laboratories) and Gerhard Klimeck (Purdue University).

In Melbourne, the leading scientists are Lloyd Hollenberg, who works on quantum sensing, not only for quantum computing but also for nanobiological applications, and David Jamieson, Chief Investigator of the Victorian node of the Center who uses ion beam techniques for the fabrication of (doped) nanodevices for the control and readout single electron and nuclear spins. The main experimentalist in quantum dots in their team is Andrew Dzurak.

David Reilly, Sidney David Reilly is Professor for experimental physics with a broad research spectrum including quantum dot spin qubits, quantum hardware (with a focus on technology for the control of condensed matter systems) and topological electronics. He was hired by Microsoft to lead the Sidney Group of Station Q with the goal to build a quantum computer based on Majorana fermions.

Andrew Sachrajda, Sherbrooke Andrew Sachrajda is a Canadian Director for the Canadian European Research Initiative on Nanostructures (CERION) and the leader of the quantum physics group at the Institute of Microstructural Sciences of the National Research Council in Ottawa. He focuses on quantum transport techniques and their applications to semiconductor nanostructures. He has been the first to isolate a single electron in a quantum dot. Recently, restructuring of the NRC has made it impossible for him to remain competitive in fundamental research.

Michel Pioro-Ladriere, Sherbrooke Michel Pioro-Ladriere is an Assistant Professor at the Universite de Sherbrooke working on spin coherence in quantum dots, quantum algorithms for spin qubits and coupling spins to microwave photons. He is an alumnus of Seigo Tarucha and specializes in the use of nanomagnets in quantum dot systems.

Jonathan Baugh, Waterloo Jonathan Baugh is an Associate Professor for Chemistry at the Institute for Quantum Computing (IQC) in Waterloo. He is interested in experimental implementations of quantum information processing and storing, with electron spins in quantum dot, but also with electronic and nuclear spins in molecular crystals, using nuclear magnetic resonance techniques. His focus is on finding new pathways to scalable quantum devices.

William Coish, McGill Bill Coish, an alumnus of Daniel Loss is an Associate Professor for theoretical physics at McGill University in Montreal. He specializes in the application of nanoscale condensed matter systems for quantum information science, specifically semiconductor quantum dot platforms.

Stephanie Simmons, Simon Fraser University Stephanie Simmons is an Assistant Professor at the Department of Physics at Simon Fraser University in British Columbia. She leads the Silicon Quantum Technology Lab where they explore the capacity of linking silicon-based spin qubits to photonic qubits. In the field of silicon spin qubit, she also works together with researchers of the Centre for Quantum Computation and Communication Technology in Australia.

Andrew Briggs, Oxford Andrew Briggs is Professor of Nanomaterials at the Department of Materials, University of Oxford and director of the QuEEN program (Quantum Effects in Electronic Nanodevices). Within this program, he studies various materials and implementations for quantum technologies, like nanotubes, single molecules and graphene nanogaps. The Quantum electronic devices group at Oxford, lead by Prof. Briggs, works on carbon-based spin qubits, coupling of nanomechanical systems to electron charge or spins, atomic-scale and single molecule electronics. Collaborators include John Morton and Stephanie Simmons.

John Morton, UCL John Morton is Professor in Nanoelectronics and Nanophotonics at the University College London (UCL). His research focuses on the control of solid state electron and nuclear spins with application in quantum technologies.

QuTech, TU Delft Delft University of Technology has been a leader of quantum nanotechnology for a long time. The research unit QuTech brings them together with other organizations including the Dutch Technology Organization DTO, to aggressively pursue quantum computers. They recently have entered a long-term partnership with Intel and another one with Microsoft. Their founding director, Leo Kouwenhoven is the leader of the topological quantum computing roadmap of QuTech, currently focusing on topological effects in solid state devices like Majorana fermions, and topological qubits. He has made various scientific breakthroughs in the fields of quantum dots and nanowires. In his lab, he recently demonstrated full control of single spins in nanowire double quantum dots (in future, also entanglement between two qubits is planned), furthermore his team works on the observation of Majorana fermions in nanowire-superconductor systems, and spin qubits based on carbon nanotubes. Kouwenhoven was recently hired by Microsoft, who are building a new lab in Delft. Next to Leo DiCarlo, described under Josephson junctions, the most visible quantum computing researcher there is Lieven Vandersypen. Educated in NMR, Vandersypen has made key contributions to spin qubits in GaAs including the leading readout scheme based on spin-charge conversion and is currently moving to Si-based qubits. QuTech has a strong theory component with Stephanie Wehner and soon Barbara Terhal.

Institute Nanosciences et Cyrogenie, Grenoble The Institute Nanosciences et Cyrogenie (Inac) is a joint research institute of the Commissariat à l'énergie atomique et aux énergies alternatives (CEA) and the University of Grenoble. Important researchers are Maude Vinet, Silvano de Franchesci and Xavier Jehl at the Lab for Statistical Physics, Magnetism and Superconductivity. Their work has a focus on hybrid superconducting- semiconducting systems, but also includes semiconductor transport effects in nanodevices, and quantum dots. De Franchesci has made early achievements on the superconducting proximity effect in semiconductor nanowires, and the Kondo effect in quantum dots, furthermore he has built the first quantum dots based on SiGe self-assembled nanocrystals. They collaborate with LETI, a pre-commercial semiconductor foundry.

Jörg Wrachtrup, Stuttgart Jörg Wrachtrup works on scalable systems with spins in diamond, recently developing an array of silicon carbide nanopillars together with Jelena Vuckovic from Stanford University. He is part of the DFG research group on Diamond Materials and Quantum Applications.

Hendrik Bluhm, Aachen Hendrik Bluhm is a Professor at RWTH Aachen and leader of the Quantum Technology Group, performing research on GaAs and SiGe spin qubits, where the principal investigator for SiGe is Lars Schreiber.

David DiVincenzo, Aachen, Jülich David DiVincenzo, who has worked long time for IBM, is a Professor at RWTH Aachen since 2011 and director of the Institute of Theoretical Nanoelectronics at the Peter Grünberg Institute in Jülich. He is well known for his theoretical study on the requirements necessary for a quantum computer, the DiVincenzo Criteria. In 1997 he proposed a quantum computer based on electron spins in quantum dots [LD98] together with Daniel Loss. Besides his research on superconducting qubits, he works on spin qubits in semiconductors in strong collaboration with experimental groups in Aachen, Delft and at IBM and also in topological quantum computing with Majorana qubits. Together with his wife, Barbara Terhal, he also does investigations in the field of error correction.

Daniel Loss, Basel Daniel Loss is Professor for theoretical physics at the university of Basel and director of the Center for Quantum Computing and Quantum Coherence (QC2). Before, he has worked at IBM and as a postdoc with Nobel Laureate Anthony Leggett. His research interests include semiconducting nanostructures and molecular magnets. His theoretical work on spin qubits has stimulated many experimental investigations worldwide.

Seigo Tarucha, Tokyo Seigo Tarucha is Professor for physics at the University of Tokyo who leads the Quantum Functional System Research Group as a part of RIKEN. The group aims at developing semiconductor and superconductor nanostructure devices for quantum computing. In special, they study electronic properties of low-dimensional systems and spin-based quantum computing platforms like self-assembled quantum dots, nanowires and spin-photon interfaces.

Toshimasa Fujisawa, Tokyo Toshimasa Fujisawa is leading a Lab at the Tokyo Institute of Technology in collaboration with NTT Basic Research Laboratories. He studies quantum transport and single electron dynamics in nanostructures with a focus on single electron measurement and detection techniques.

15.1.4.1 Sample growers

Arthur Gossard, Santa Barbara Arthur Gossard, formerly working at Bell Laboratories, is a Professor of Materials and Electrical & Computer Engineering at UC Santa Barbara. He is known for growing the first modulation doped quantum wells, and being involved in the discovery of the quantum confined Stark effect and the fractional quantum Hall Effect. He is specialized in molecular beam epitaxy, growth of quantum wells, semiconductors and metal-semiconductor nanocomposites.

Chris Palmstrom, Santa Barbara Chris Palmstroms work focuses on the understanding of heteroepitaxial growth, combining materials with different properties with the aim to grow novel materials and structures for (opto)electronic, magnetic and micromechanical devices.

Werner Wegscheider, Zürich Werner Wegscheider is Professor at the Department of Physics at ETH Zürich. He is working on instrumental improvement for fabrication of As- and Sb-based semiconductor heterostructures like quantum dots or quantum wires, but also interested in studying their optical and electronic properties.

Michael Pepper, London Michael Pepper is a Professor of Nanoelectronics at the University College London where he is head of the Electronic Materials and Devices Group. When he was still at the Cavendish Laboratory in Cambridge, he was co-authoring the paper announcing the discovery of the quantum hall effect that brought the Nobel prize to Klaus von Klitzing. He was also pioneering the study of low dimensional electron gases. Current research is on new growth mechanisms for carbon-based structures like nanotubes and heterostructure devices, and also on single-atom rare-earth doped silicon nanoclusters for quantum computers.

Element-6 This is a subsidiary of diamond distributor deBeers that specializes in artificial diamonds. They are involved in making tailored diamonds for NV-Center research.

15.2 Evaluation: Spins in quantum dots

Spins in semiconductor quantum dots based on the Si/SiGe Material family have made qualitative progress in the summer of 2017. In particular, two-qubit gates have been realized, completing the gate family. These were demonstrated for the Loss-DiVincenzo qubits with one spin qubit per dot, hence the double dot is a two-qubit system [WPK⁺17, ZSR⁺17]. This makes this the most relevant platform in this family compared to singlet-triplet and triple-dot qubits (see 15.1.2) and the only one on level B. Challenges from residual nuclear spin dephasing and charge noise [ZSR⁺17] persist.

15.2.1 DiVincenzo criteria

15.2.1.1 Scalable qubits

In order to maintain spin coherence, the silicon needs to be isotopically purified. Only double-dots were demonstrated so far, but proposals for scaling exist and have been shown (with inferior coherence) in related semiconductor structures. Similarity to current microelectronics makes it very likely that this can be achieved. The largest realized sample consists of a linear array of nine quantum dots plus three quantum dots for single-shot charge readout [ZHM⁺16].

As an important architectural primitive, 2018 has seen the realization of coupling of single spins to microwave resonators similar to circuit QED in superconductors [SSK+18,SZK+18].

15.2.1.2 Initialization

Si/SiGe double quantum dot initialization fidelity > 0.99 [WPK⁺17] was demonstrated. Initialization and readout of one dot done with spin-selective tunneling to a reservoir [EHWvB⁺04] (Pauli Spin Blockade), initialization of the other dot at a spin relaxation hot-spot [SNS⁺13] and readout via a controlled rotation and dot 2. It is assumed that waiting $7T_1$ leads to 100% initialization of dot 2.

15.2.1.3 Universal gates

Arbitrary single qubits gates are achieved with magnetic and electrical controls, see [TvdWOT06, PLOT⁺08, TKO⁺16]. In 2018, two-qubit gates with fidelities above 90% up to 98% were achieved [HYC+18]. Quantum state tomography of Bell state 0.85–0.89 [WPK⁺17]. Two-qubit gate (CZ) in 100ns and CNOT in 480ns [VYH⁺15], with $T_2^* = 120\mu\text{s}(61\mu\text{s})$ and $T_2 = 28\text{ms}$ with CPMG.

With the complete gate set, full benchmarking of two-qubit gates has been done, we have already quoted key data above. Fast single qubit gates are already at threshold with single-qubit RB results for the left (right) qubit 0.993 (0.997) [ZSR⁺17] and 0.988 (0.98) [WPK⁺17]. Gate set tomography leading to improved calibration of single-dot [DMBK⁺16], with average gate fidelity 0.99942 of single-qubit gates. They are compatible with RB data, pointing at no major artifacts in RB.

15.2.1.4 Coherence

Ramsey and Hahn echo measurements for the left (right) qubit $T_2^* = 1.2(1.4)\mu\text{s}$ $T_{2,echo} = 22(80)\mu\text{s}$ [ZSR⁺17]. Spin relaxation limited [ZSR⁺17]. Spin relaxation time $T_1 > 50\text{ms}$ ($= 3.7\text{ms}$), $T_2^* = 1.0\mu\text{s}$ ($0.6\mu\text{s}$), $T_{2,echo} = 19\mu\text{s}(7\mu\text{s})$ [WPK⁺17]. Together with the gate time this leads to a coherence limit of gate fidelity based on T_2^* of 84%, so gates are coherence limited. The contrast between this and the spin echo time highlights the potential for improvement by composite pulses.

15.2.1.5 Readout

The readout fidelities of the most recent experiments that allow for two-qubit gates are not reported in [ZSR⁺17] and relatively low with 0.73(0.81) [WPK⁺17]. Much higher fidelities have been reached in older samples, such as single-shot in [PSS⁺12] and singlet-triplet high-fidelity single-shot readout with fidelity 0.98 [FCH⁺17]. It is unclear whether there is a fundamental reason for this discrepancy and we expect that newer samples will catch up with the historic higher fidelities.

15.2.2 Summary

Criteria	met?	Comments
Scalable qubits	✓	Needs isotopically pure silicon
Initialization	✓	
Universal gates	✓	
Coherence	✓	
Readout	✓	

Table 15.1: Summary DiVincenzo criteria for spins in quantum dots.

15.2.3 Analysis and outlook

Si/SiGe qubits are now meeting all DiVincenzo criteria in samples up to two qubits. There do not seem to be any principle obstacles to going bigger - the two-qubit gates are extremely new results and certainly there are immediate improvements ahead.

Single-qubit gate fidelities are reaching threshold values. Two-qubit gate fidelities in this platform are way below the fault tolerance threshold for two-dimensional architectures. As we have just seen the first two-qubit gates, this is not a big surprise. The main challenge of this architecture is overcoming intrinsic charge noise - even though we are fundamentally looking at spin qubits, effects like the exchange interaction or spin-orbit interaction inevitably lead to some orbital / charge structure in the qubit state. Charge noise has been well-known for a long time independent of quantum computing, so it cannot be assumed that this is easy to solve. The neighboring Josephson qubit community has overcome noise problems by reducing the impact of noise rather than eliminating the source which could be an option here. Clearly, given the proximity of this technology to classical CMOS technology, one can expect shortcuts in operation once these obstacles are overcome.

15.3 Evaluation: Other quantum dot platforms

Given the consolidation of the field, most other quantum dot platforms are currently not being pursued as they got stuck on level A. For completeness and in order to provide context if some of these problems are solved, we mention the state of the art. Unless mentioned otherwise, techniques resemble those of SiGe spin qubits.

15.3.1 Double-dot qubits, singlet/triplet

The single-triplet qubit is pursued as a strategy to overcome the impact of nuclear spin noise, hence it is pursued in GaAs. Pioneering work in GaAs shows coherence of $T_1 = 16\text{ns}$ and $T_2^* \geq 400\text{ps}$ [PJM⁺04]. Independent single-shot readout of left and right qubit is possible [NSL⁺11], with average fidelity of 0.86.

- Single-qubit gate fidelities of 0.999 theoretically achievable [CBDB14]. Coherence times exceed $200\mu\text{s}$ [BFN⁺11]. A higher coherence time $T_2 = .87\text{ms}$ is possible by decoupling the qubit from its nuclear environment [MMN⁺17a]. Readout in $8\mu\text{s}$ with 0.8 visibility [MMN⁺17a]. T_1 is in excess of $100\mu\text{s}$ [DSH⁺13], $T_2^{\text{echo}} = 9\mu\text{s}$. Typical relaxation times are $\geq 70\mu\text{s}$ [PJY⁺05].
- Entanglement with controlled phase (CPhase) [SDH⁺12]. Two-qubit state tomography reaches a maximum Bell state fidelity of 0.72 at $t = 140\text{ns}$ and slightly lower fidelity for a π -phase at 160ns . Due to crosstalk the readout is done sequentially.

This technically puts Single-Triplet qubits on level B—but as the challenges to higher fidelity seem insurmountable, they are not pursued for quantum computing any more.

In Si: Process tomography of single-qubit gates [KSS⁺14], with 0.85 fidelity for X rotations and 0.94 for Z rotations. gate times ~100ps. $T_1 \approx 23.5\text{ns}$ $T_2^* \approx 10\text{ns}$. No 2-qubit gates.

Criteria	met?	Comments
Two qubit gates	✓	Low fidelity
Coherence	?	Very few operations

Table 15.2: Short summary double-dots.

15.3.2 GaAs single dots

Very long T_1 , ranges from 7 to 85ms [SKS⁺14]. Very short T_2 due to nuclear spins.

Criteria	met?	Comments
Two-qubit gates	×	
Coherence	×	

Table 15.3: Short summary single-dots.

15.3.3 Triple-Dot GaAs

Typical T_2^* for the three-spin qubit experiments is 8–15ns, being dominated by local uncorrelated nuclear field fluctuations, similar to double quantum dots [GGK⁺11]. Dynamical decoupling to increase coherence times [MMN⁺17b], single-shot readout on microsecond timescale. [LTD⁺10] demonstrates initialization, two-spin coherent manipulation, and readout. Fast readout of charge states takes advantage of multiplexed reflectometry. Measurement times are $t_m = 5\text{--}10\mu\text{s}$, while a π -pulse takes 350ps [PJT⁺05]. Typical π -pulses are shorter than T_2^* , ~100ns [TPLO⁺10]. Allows quantum simulations in a dot array [HFJ⁺17].

The move to SiGe has removed the necessity this complex to engineer platform.

Criteria	met?	Comments
Two-qubit gates	×	
Coherence	?	Very few operations

Table 15.4: Short summary triple-dots.

15.3.4 Summary: Alternative quantum dot platforms

Criteria	met?	Comments
Scalable qubits	✓	High mobility GaAs
Initialization	✓	
Universal gates	×	No or poor two-qubit gates
Coherence	?	Sufficient for very few operations in GaAs
Readout	✓	

Table 15.5: Summary DiVincenzo criteria for other Quantum dot platforms.

15.4 Evaluation: Single dopants in Si/SiGe

The coherence time measured using dynamical decoupling is $T_2 = 400\mu\text{s}$, and the Ramsey decay time $T_2^* = 1\mu\text{s}$, RB of single-qubit gates gives 0.9899, initialization and readout times are around 4ms [KJS⁺16]. Long coherence times of electron and nuclei, see list of records [MDL⁺14]. The latter motivates the theoretical proposal of the flip-flop qubit [TMS⁺15]. RB of 0.9995 and 0.9999 for the electron and the nuclei, respectively [MLS⁺15]. Gate set tomography 0.99942 of the same sample [DMBK⁺16]. These excellent numbers predict strong potential for this platform once two-qubit gates have been realized, currently they are on level A.

Criteria	met?	Comments
Scalable qubits	✓	Needs isotopically pure silicon
Initialization	✓	
Universal gates	×	Two-qubit gates missing
Coherence	✓	
Readout	✓	

Table 15.6: Summary DiVincenzo criteria for Single Dots in Si/SiGe.

15.5 Evaluation: NV-Centers

15.5.1 DiVincenzo criteria

15.5.1.1 Scalable qubits

The key challenge in this platform is to integrate beyond the two qubits (electron and nucleus) at a single center without sacrificing its great properties.

Coupling NV-center crystals through a transmission line resonator [ANP⁺17] has been achieved with ensembles, but not with single NV centers. NV center interconnection with optical photons [NTD⁺14] is very successful [HBD⁺15], but generates significant overhead. Spacing of NV centers around 10 to 100nm [BSA17] is theoretically achievable. So far, attempts at direct implantation of an array of NV centers have not been successful.

15.5.1.2 Initialization

Initialization takes a $3\mu\text{s}$ laser pulse [NMR⁺08] and including waiting times can be done in $\leq 8\mu\text{s}$, with a fidelity of ≥ 0.9 [DCJ⁺07].

15.5.1.3 Universal gates

Dynamical decoupling with logical gates [ZSBS14] leads to process fidelities of 0.985 and a gate duration of $35.5\mu\text{s}$. Photon mediated gates between NV centers in an optical cavity [AB16b] for CZ gates take $10\mu\text{s}$. Further improvement down to 200ns [BSA17] are possible. A Bell-state fidelity up to 0.98 for two nuclear spins has been shown [NMR⁺08]. Electron-nuclear spin two-qubit gates take 100ns [JGP⁺04] with average fidelity ~ 0.92 . Single-qubit gates with fidelity 0.89 [BFBA10]. Entanglement between NV center spin and photons of 0.7 [TCT⁺10].

15.5.1.4 Coherence

The typical spin coherence time is $10\mu\text{s}$ [BSA17]. The phase memory time T_2 is found to be around 0.6ms [NMR⁺08] with two nuclear spins. Very short coherence of $T_2 \sim 6\mu\text{s}$ were reported in [JGP⁺04], but up to $60\mu\text{s}$ possible. Ground state manifold nuclear decoherence is $T_2 = 480\mu\text{s}$ [BFBA10], and Hahn echo of nuclear spin at $T_{2e}^* = 495\mu\text{s}$ [DCJ⁺07], such that coherence times span two orders of magnitude.

15.5.1.5 Readout

Projective optical readout [RCB⁺11] at 8.6K has been shown. Nuclear spins are read out by CNOT and electronic spin readout. Electronic-nuclear flip-flop transitions can reduce optical readout fidelity. In principle this allows scaling for multi-nuclear-qubit readout. The average fidelity is 0.93 . Single-shot readout fidelities are at 0.92 [NBS⁺10] and ~ 0.97 [HBD⁺15].

15.5.2 Summary

Criteria	met?	Comments
Scalable qubits	?	Optical link only so far
Initialization	✓	
Universal gates	✓	Local, not long-range
Coherence	✓	Low coherence for long range interaction
Readout	✓	

Table 15.7: Summary DiVincenzo criteria for NV centers.

15.5.3 Outlook

This may be the most controversial evaluation across the range of platforms. NV centers have demonstrated a lot, but not at the same time. According to our criteria, NV centers need to demonstrate convincingly that they can integrate beyond 2 qubits without sacrificing the marvelous speed and quality of operations shown for single centers, placing them on level A. Partial success has been achieved by using multiple nuclear spins, but it is not clear if this provides a route to arbitrary scaling. It is not clear if and when this will be successful given that the NV center community focuses on sensing and photonics currently. It needs to be noted, that some of the Level C milestones have in fact been met, specifically, quantum error correction for multiple electron spins has been demonstrated [WWZ⁺14]. The situation is thus a bit between nuclear magnetic resonance (where it is consensus that scalable initialization has not been and is not likely to be found) and trapped ions (where also a challenge of scaling while maintaining high quality is currently a main milestone, but it has happened at a larger number of qubits, large enough to clearly master level A).

15.6 Evaluation: Topological qubits—Majorana fermions

Topologically protected qubits [AHM⁺16b, LSDS10, SZH⁺16, Wil09] are prominently developed in semiconductor nanowires hosting Majorana zero modes at their edges but are also / have been pursued in other platforms. As described in 15.1.2.8, Majorana qubits are still a very active research field. Although there is still no fully conclusive experimental proof of the existence of Majorana fermions, theory has proposals for all elements of universal quantum computing that might be realized very fast if material requirements are met. Just recently [GCZ⁺17], Microsoft station Q presented a new technique for growing nanowire structures that fulfill all requirements necessary to perform the first braiding experiments. Topologically

protected qubits fall in category A of our evaluation scheme. We will discuss in the end why they may provide a shortcut to category C.

In [PREF17, KKL⁺17], a completely new approach for designing topologically protected qubits was proposed which solves already some of the current experimental problems.

15.6.1 DiVincenzo criteria

15.6.1.1 Scalable qubits

Structures for storing Majorana qubits can be fabricated. A leading approach is creating Majorana zero modes at the ends of semiconducting nanowires [LSDS10, ORvO10] that are covered with superconducting metal in order to induce superconductivity by the proximity effect, and a strong magnetic field in the parallel to the wire. Experimental evidence so far have been indirect and are largely based on transport measurements, and could in principle also be understood without the use of majorana modes [LJH⁺13]. Furthermore, experimental results [MZF⁺12, DRM⁺12, RLF12, CFGR⁺13] that give evidence to Majoranas have low visibility due to temperature effects, which would suggest very low quality of qubits. This situation is motivating researchers to move on and verify Majorana properties by attempting further steps.

In fabrication, there is still ongoing work in material science to find the right material with optimal properties, which are (induced) superconductivity, large spin-orbit coupling, hard superconducting gap and separation of energy scales between topological gap and Majorana splitting. It is important to keep these properties also when the magnetic field is applied.

15.6.1.2 Initialization

Theoretical proposals exist, but have not been realized. The qubits are automatically initialized when created, which can be done via external electrical control.

15.6.1.3 Universal gates

No coherent gates have been demonstrated. Theory proposes to perform (multi-qubit) Clifford gates via braiding of anyons, similar to moving holes in the surface code. Moving anyons can in principle be realized by changing the chemical potential along the wires (via external electrical gates) and connecting them in a network. Creating these network is not much harder than creating single wires, growing crossed wires has already been demonstrated [PvWC⁺13, GCZ⁺17], so once the first good working qubit is realized, the step to even multi-qubit networks and gates is expected to be short. In engineering multi-qubit circuits, a central challenge is to maintain the magnetic field parallel to the wires - components along the wires would induce orbital effects destroying the Majorana physics.

Universal quantum computing is proposed to be possible with the use of additional state distillation for T gates.

Right now, the fact that the energy scales of the topological gap and the Majorana splitting are still close leads to very bad achievable gate qualities. In order to move anyons adiabatically, a better distinction between those energies is essential.

The approach of [PREF17, KKL⁺17] proposes to use so-called Majorana-box qubits, a set of three parallel coupled nanowires, where logical gates are performed via projective measurements, circumventing the need for braiding and orthogonal crossings of wires (which are problematic since the magnetic field should always be parallel to the wire).

15.6.1.4 Coherence

Measurement of coherence times would require rudimentary gate operations which has not been achieved, so no clear statement is possible.

Theoretical predictions are extremely optimistic. One limiting mechanism is quasiparticle poisoning due to subgap bound states. It gives an upper bound to the coherence time. The current maximum in experiments is around 1min [vWKG15]. This could be further enlarged by finding an implementation with a hard superconducting gap. The proposal of [PREF17, KKL⁺17] also predicts less quasiparticle poisoning.

Furthermore, accidental braiding or spurious anyons due to disorder in the chemical potential could for example lead to leakage effects.

15.6.1.5 Readout

Readout of logical states of topological qubits has not been experimentally achieved.

The leading theoretical proposal is fusion of two Majoranas which form a qubit, resulting either in the existence of an electron or vacuum, depending on the qubit state. This difference in charge can then be read out, for example with quantum dots. The readout fidelity then strongly depends on the readout fidelity of the quantum dot. Readout fidelities of 99% are expected within 1 μ s, alternatively, dispersive readout should be possible within 1ms or less, depending on the amplifier [AHM⁺16a].

15.6.2 Summary

Table 15.8 shows that in experiment, none of the DiVincenzo criteria is satisfactorily met. But although not experimentally realized in sufficient accuracy, theoretical proposals for realization of all criteria exist and meeting them mostly hinges on improvements of the underlying materials.

Criteria	met?	Comments
Scalable qubits	?	only evidence
Initialization	×	proposed
Universal gates	×	proposed
Coherence	×	proposed
Readout	×	proposed

Table 15.8: Summary of the experimental status of DiVincenzo criteria for Majorana Fermions.

15.6.3 Outlook

If all necessary material properties are realized and devices can be grown clean enough, then theory predicts very long stability and coherence time of qubits due to topological protection. In that case no additional error correction would be necessary, which would make this platform very scalable and a direct candidate for category C.

Right now, no such perfect devices or materials exist, and it is not yet known if all error channels can be avoided completely, in the end, Majorana qubits might not be better than any other qubit. Future experiments will show in what direction this platform evolves and most importantly, if the observed states really obey the topological protection and non-Abelian braiding statistics [CSDS17] needed for computation.

15.7 Evaluation: Topological qubits—Other candidates

Several other platforms for topologically protected qubits are or have recently been pursued. The experimental situation is a bit more advanced than in nanowires in that there is mounting evidence that the excitations in such systems indeed are non-Abelian excitations as expected for qubits.

15.7.1 Fractional quantum Hall states

Topological protection, coherence measurements, gates and others have not been done. The $5/2$ fractional quantum hall effect requires ultrapure samples and ultrahigh magnetic field and similar to semiconductor nanowires these are the limiting factors. The largest sponsor, Microsoft, has shifted their attention away from this platform. See Table 15.9 for DiVincenzo criteria.

Criteria	met?	Comments
Scalable qubits	✓	inconclusive
Initialization	×	proposed
Universal gates	×	proposed
Coherence	×	theory
Readout	×	proposed

Table 15.9: Summary of the experimental status of DiVincenzo criteria for fractional quantum Hall states.

15.7.2 Strontium ruthenate

As described in 15.1.2.9, again there is inconclusive evidence for topologically protected quasiparticles and no other DiVincenzo criteria have been met, see Table 15.10. Given the difficulty of making samples that allow to build a quantum computer architecture, this platform has been abandoned in favor of nanowires.

Criteria	met?	Comments
Scalable qubits	✓	inconclusive
Initialization	×	proposed
Universal gates	×	proposed
Coherence	×	theory
Readout	×	proposed

Table 15.10: Summary of the experimental status of DiVincenzo criteria for topologically protected qubits in strontium ruthenate.

15.7.3 Josephson junction arrays

Josephson junction chain experiments have clearly measured qubits, performed at least a subset of single qubit operations and have shown increased coherence and readout. On the one hand, that makes them the leader of the topological qubits field right now, on the other hand, as the same building blocks can be used to build a surface code, which is also based on topology, this approach competes with the more conventional Josephson qubits, in particular because the topological protection was only demonstrated qualitatively, but did not beat the next transmon qubits. This research, which is experimentally only pursued by a single group (Misha Gershenson, Rutgers) will probably help building better Josephson qubits and merge with that stream. A summary of fulfilled DiVincenzo criteria is shown in Table 15.11.

Criteria	met?	Comments
Scalable qubits	✓	
Initialization	✓	
Universal gates	×	proposed
Coherence	?	measurements exist, not above transmons
Readout	✓	

Table 15.11: Summary of the experimental status of DiVincenzo criteria for topologically protected qubits in Josephson junction arrays.

16 Neutral atoms

16.1 Description

16.1.1 Terminology

The field of neutral atom qubits uses electrons in atomic shells to carry quantum computing. These atoms are not ionized but neutral, hence not containing the possibility to couple to them with electrostatic forces to trap and couple them. This challenge can be used to structure the field. The lack of electrostatic interaction also means lack of electrostatic repulsion allowing to bring atoms very close together, which allows fast gates and dense packing of atoms. Over the last years, Rydberg gates have emerged as the only serious candidate for neutral atom quantum computing. We will also describe the approaches of collisional gates and of cavity quantum electrodynamics, which are seriously pursued for quantum technologies other than quantum computing as well as fundamental research. We will describe the obstacles that would need to be overcome to be quantum computing contenders.

16.1.2 Rydberg atoms

It needs to be stated upfront that despite their name, Rydberg atoms are by no means special atoms - they are atoms prepared in so-called Rydberg states.

16.1.2.1 Qubits

The qubit states are encoded in electronic states of the atoms of interest. In order to have only one atom in the outer shell, Alkali metals are used. The particular states that are used are hyperfine states, i.e., ground states split by interaction with the atomic nucleus. This scales with the number of protons Z as Z^4 hence motivating the use for heavy elements like i.e., Rb and Cs.

The atoms need to be held in a specific place to be addressable and the qubits to be well-defined. This is accomplished with laser-induced dipolar forces, an indirect effect, hence very weak. [SWM10]. Given the weakness of trapping forces, active cooling (again by lasers) using a variety of methods is imperative and loss of qubits still a risk. Motivated by neutral atoms research, the impact of qubit loss on error correction is currently being studied [Smi16].

The central device to combine trapping and cooling is called magneto-optical trap (MOT). Holding qubits by light makes it natural to produce regular qubit arrays by trapping with standing waves, however, local addressability by laser led to systems with larger lattice spacing such as bottle beams [IIS⁺13] (light beams with high intensity in a tube, similar to the outside of a bottle) and other geometries.

16.1.2.2 Operations, Rydberg gates

Single qubit operations are achieved with the same techniques as in trapped ions. Two-qubit operations are based on Rydberg blockade [SWM10, Saf16, BBL16, RLB⁺14, PBA14]. An electron is in a Rydberg state if it is prepared in a state of very high principal quantum number, i.e., to energies very close to ionization, typically to values of $n = 50 \dots 100$. The size and dipole moment scales with n^2 , and the interaction strength scales with n^4 for long-range dipole-dipole, and n^{11} for short-range van-der-Waals interaction.

Driving atoms into Rydberg states is achieved by precisely tuned lasers. By choosing which atoms are driven through Rydberg states, gates are made selective and given the long range of the dipole-dipole interaction

which decays only with distance cubed, these can be long-range. Initialization and measurement are performed analogous to that in trapped ions.

16.1.2.3 Status of the field and challenges

Currently, physics problems need to be solved for Rydberg atoms. One is improvement of atom loading into their traps, cooling and long-term storage—qubits still disappear or never appear. Also, gate quality needs to improve. With the strong technological activity in the Saffman group under the IARPA-MQCO program, a lot of technological and scale-up innovation has been demonstrated. The combination of these aspects makes further technological challenges hard to predict. Neither the European roadmap nor the current leading US program, IARPA-LOGIQ, highlight Rydberg atoms.

16.1.3 Collisional gates

As an alternative to Rydberg gates, two-qubit interactions can be implemented using collisional gates [MGW⁺03, NTC11, ALB⁺07]. In this case, the trapping field holding the atom is manipulated such that atoms come in close contact, comparable to the length of a molecular bond, i.e., with overlapping electron clouds. The energetics in this state depends on the internal state of the atom, a phenomenon called Feshbach resonance. This approach is highly successful in cases when a lot of two-atom interactions need to be controlled in parallel, i.e., in quantum simulations in optical lattices, where the whole trapping laser field can be moved as a whole. It has been proposed to extend this to quantum computers by moving atoms with a selective optical tweezer - a beam of light that uses dipole forces to interact with the neutral atom similar to the optical lattice, but is tightly focused [WKMS11]. Typical moving times of rearranging arbitrary 2D arrays of Rydberg atoms are 50 ms [BdLL⁺16] and an interaction strength of neighboring atoms in the MHz range. These tweezers are developed for low speed in quantum simulations, fast optical tweezers needed for gates are a far-fetched projection [WKMS11]. Collisional gates have not been considered a promising route to quantum computing since the establishment of Rydberg gates.

16.1.4 Cavity quantum electrodynamics

Cavity quantum electrodynamics with neutral atoms uses the coherent interaction between atoms and photons, single quanta of light. It is primarily a very clean platform for basic research in quantum physics [HR06, Har13] and has applications in quantum communication [Kim08]. Its basic functional element are neutral atoms in Rydberg states are sent through cavities in order to precisely interact the photonic state of the cavity. As this was an early coherent and controlled quantum system, quantum computing proposals were put forward, either using atoms as qubit and cavities for interaction or vice versa. Given the enormous size of the cavity (which is dictated by the requirement to have modes that are resonant with Rydberg transitions) and the complexity producing these, scaling to any reasonable size processor is not pursued. The tools of cavity quantum electrodynamics have been taken over to other approaches involving cavities, specifically circuit QED, described with superconducting qubits. They can also used for long-distance gate in quantum networks [WHD⁺18].

16.1.5 Research groups

Again, as much of this field is now pursuing quantum simulations, foundations, and sensing, we put a tight focus on those groups pursuing quantum computing.

Mark Saffman, University of Wisconsin Mark Saffman was one of the two first researchers to show Rydberg blockade and implement Rydberg gates. He has led a concerted effort to scale up from 2009--16 within the IARPA-MQCO program and now has different (lower) funding to continue. He is mostly focused on building a quantum computer within this field.

ColdQuanta This is a university spinoff in Boulder, CO that develops compact components for cold atom experiments. They are a long-term partner of Saffman.

Philippe Grangier and Antoine Browaeys, Institut d’Optique, Paris Philippe Grangier at Laboratoire Charles Fabry de l’Institut d’Optique/CNRS has observed Rydberg blockade in the same era as Saffman. He is director for the quantum optics research is mainly working on continuous variables quantum cryptography and recently on single-photon nonlinearities (DELPHI) to use Rydberg atoms as a highly non-linear optical medium.

The institute includes the “Quantum Optics – Atoms” group, led by Antoine Browaeys, using arrays of optical tweezers each to manipulate single Rydberg atoms which scales up to 50 spins. Furthermore they trap up to about 300 Rydberg atoms in the same microscopic dipole trap to study collective effects.

Charles Adams, Kevin Weatherill, Durham (UK) These groups have realized Rydberg blockade and pursue its application in quantum technologies.

Gerhard Birkl, Darmstadt Gerhard Birkl’s group is working on an approach similar to Saffman’s, but with a different optical system for trapping the atoms.

Klaus Mølmer, Aarhus Klaus Mølmer is the leading theorist in Rydberg gates.

Laboratoire Kastler-Brossel (LKB), Paris LKB is part of ENS and associated with CNRS and studies many aspects of quantum optics. Its famous activity in cavity quantum electrodynamics with Rydberg atoms in microwave cavities led by Nobel Laureate Serge Haroche, Jean-Michel Raimond, and Michel Brune. They have been pioneers in many areas but, as mentioned above, their platform is not currently pursued for scaling a quantum computer given its size and difficulty of fabrication.

Jeff Kimble, Caltech Jeff Kimble specializes on cavity QED with atoms at optical frequencies as an optical interface, but is not pursuing quantum computing any more.

Gerhard Rempe, MPQ Gerhard Rempe uses cavity QED with atoms at optical frequencies and continues to pursue quantum gates between matter and light as one of his many activities.

16.2 Evaluation: Rydberg atoms

16.2.1 DiVincenzo criteria

16.2.1.1 Scalable qubits

Recently, quantum simulation with a linear array of 51 ^{87}Rb atoms trapped with optical tweezers has been performed [BSK⁺17], aiming for 2D simulators in the near term future, but this does not provide a route to optimal control.

16.2.1.2 Initialization

Laser cooling and optical pumping leads to state preparation fidelities of 0.95 [JHK⁺16].

16.2.1.3 Universal gates

Single qubit gates with RB average fidelities of 0.998 in 2D [XLM⁺15] and 0.996 in 3D [WKWW16], using ^{133}Cs atoms. Bell-state fidelity to verify two-qubit gates have been shown with fidelity 0.79 [MLX⁺15] and 0.81 [JHK⁺16] in ^{133}Cs , and 0.634 in ^{87}Rb [KLFF⁺15]. Optimal control methods might help to improve two-

qubit gate fidelities [TMWS16]. Two-qubit gates are either performed via Rydberg-blockade [MLX⁺15, JHK⁺16] or local spin-exchange with optical tweezers [KLFF⁺15].

16.2.1.4 Coherence

Coherence time in ¹³³Cs have been confirmed to be $T_2' = 7\text{s}$ in [WKWW16] and $T_2^* = 7\text{ms}$ in [XLM⁺15]. Lifetimes of the Rydberg state are around $40\mu\text{s}$ [JHK⁺16].

16.2.1.5 Readout

Readout through measurement of the fluorescence signal, similar to ion traps.

16.2.2 Summary

Criteria	met?	Comments
Scalable qubits	✓	
Initialization	✓	
Universal gates	✓	Low fidelity two-qubit gates
Coherence	✓	
Readout	✓	

Table 16.1: Summary DiVincenzo criteria for Rydberg atoms.

16.2.3 Outlook

A blueprint for a fault-tolerant quantum computer built of optically trapped Rydberg atoms has been developed [ABB17], aiming for 10^4 qubits. Currently, gate fidelities as well as atom loss rates are however on a level where these extrapolations are largely speculative—level B.

17 Photonic qubits

17.1 Description

Photonic qubits comprise qubits that are encoded in quantum states of the light field. The structure of this field is given by a very unbalanced set of quantum computing resources: While light is very flexible to use and can be very coherent, and while single qubits can be easily manipulated with optical elements (mirrors, phase shifters, beam splitters), photons are not known to interact with each other (classical electrodynamics described by Maxwell's equations in vacuum is a strictly linear theory), leading to no natural two-qubit gate. Approaches to photonic qubits can be classified by the way how they compensate for this challenge. In general, while photons are a cornerstone of many other quantum technologies, they are not in the center of attention for quantum computing. Note that boson sampling [AB16a, AA13, LBH⁺16, CHS⁺15] is known as a road to quantum supremacy, i.e., to demonstrate outperforming classical computers in the application of computing the permanent of a matrix, but is believed to not have any applications beyond that (nor has computing the permanent any cryptographic implication). An operational peculiarity is, that given the immense size of the speed of light, manipulating photonic states in time is very challenging, thus most proposals focus on arranging a quantum algorithm in space.

17.1.1 Qubit encoding

There are multiple methods to encode quantum bits in the light field [OFV09]. One is the use of mode occupation. A mode of the light field is a classical solution of Maxwell's equations typically characterized by its spatial structure and polarization-type properties. In quantum optics (quantum field theory in the limit of non-relativistic matter) [Fox06], the quantum states of these modes are described analogous to harmonic oscillators, with the degree of excitations being interpreted as the number of photons.

In this framework, binary encodings of qubits in single photons are most natural for quantum computing. In polarization encoding, a single qubit is encoded in the polarization state (right circular or left circular) of an otherwise identical spatial mode. In dual-rail encoding [KLM01], polarization is not used, rather, the presence or absence of a photon in a spatial mode, for example in arms of an interferometer, defines qubit states. The main challenge is to produce the qubits, as deterministic on-demand single photon sources are difficult to implement (even in spite of the best efforts of the photonics and quantum cryptography communities).

An alternative encoding is Gaussian quantum information [BvL05, WGP⁺12]. There, quantum information can be encoded in semiclassical coherent states that are widely separated [RGM⁺03]. These states are naturally produced by lasers and can be separated using phase shifters. Their nonorthogonality in the form of a small overlap can be compensated for, furthermore there are protocols that make it easy to correct errors caused by photon loss [CMM99]. An even more exotic approach is using squeezed states for continuous variable quantum computing, where the computational basis consists of all squeezed eigenstates of some quadrature variable [BSBN02]. These states can be produced from coherent states using nonlinear elements.

17.1.2 Enhanced nonlinear optics, integrated optics

One way how photons can interact is by using nonlinear optics [Boy03]. Nonlinear optics describes the interaction of light with matter in a way that the matter mostly introduces nonlinearity into the Maxwell equations. These naturally lead to terms in the quantum optical Hamiltonian of powers larger than quadratic in photon amplitudes. These terms are interpreted as effective photon-photon interaction. An example is the Kerr effect, the dependence of the index of refraction on light intensity. This leads, e.g., to

four wave mixing, the scattering of two incoming photons into two outgoing photons, hence naturally implementing a two-qubit gate. Another nonlinear interaction is two-mode squeezing, which leads to correlations in the quadrature amplitudes of two modes that can be used in continuous variable quantum computing.

This approach is challenged by numbers. Even very effective nonlinear materials such as barium-borate in samples that thin enough to not absorb the photons have conversion efficiencies below 10^{-6} , making two-qubit gates hugely ineffective on the single photon level.

Proposed solutions include confining the light to very small volumes using cavities and integrated optics (i.e., optics on a chip rather than discrete optics) [HBR⁺16, PLP⁺11]. This uses the concept of *mode volume*: The energy of a photon of frequency ν is inevitably $h\nu$. The energy density per volume, on the other hand, is E^2/ϵ_0 , thus the typical electric field of a photon is $E \simeq \sqrt{h\nu\epsilon_0/V}$. Stronger fields can exploit nonlinearities more. Another perspective is that confinement of light into a slightly open cavity makes the photon cross the nonlinear medium many times, giving it more opportunities to interact. Some of these approaches use atoms in cavities as a nonlinear medium. While impressive science, the success probabilities of these direct gates by engineered nonlinearity are still too low to be practical [FFE⁺08, FEF⁺08].

17.1.3 KLM proposal

The Knill-Laflamme-Milburn proposal [KLM01] is a very elegant approach to avoid the use of optical nonlinearity and replace it by the (also nonlinear) resources of single photon generation and detection as well as post-selection. The key element is the nonlinear sign (NLS) gate—a gate that conditions a phase shift on the number of photons—that is simulated using ancilla modes and that is only carried out with 1/4 success probability, but this success is certified by ancilla detection. Two of these NLS gates can be combined into a two qubit gate. The overhead of probabilistic gates and post-selection does not alter the complexity class of algorithms, but increases the hardware effort in practice. To be viable, the data qubits are held in optical memory, the entangling gate is performed on ancilla qubits and once successful, the data are teleported. The KLM gate has been demonstrated [OOHT11] with a success probability of 82% that has been gradually improved [MBB⁺16].

17.1.4 Cluster states and one-way quantum computing

At face value, the teleported KLM protocol creates an entangled ancilla state and teleports data on it. This idea can be taken to its extreme the concept of an ancilla factory in preparing all entanglement nondeterministically first and, if successful, proceed with the computation only through measurement and single-qubit operations [BR05]. This one way quantum computing approach, originally proposed by Raussendorf and Briegel [RB01] is equivalent to regular quantum computing. Cluster states can be generated using parametric downconversion in nonlinear crystals [WRR⁺05], coupled quantum dot emitters [ELR10] or, in the continuous variable case, from frequency combs in nonlinear media [FMP09]. In fact, large cluster states have been produced, yet, functional one-way quantum computing which also requires photonic memory has not been implemented. The so far largest cluster state was realized with over one million modes by continuous variable entanglement [YYK⁺16], in the discrete (conventional) case, the cluster size is still in the order of several photons [WRR⁺05, SCS⁺16].

17.1.5 Continuous variables

The continuous variable encoding has already been described above. It encodes information in squeezed states and uses the squeezing effect for two-qubit gates, a nonlinear mechanism based on the Kerr nonlinearity that entangles the field amplitudes of both modes. The gate is usually performed by a two-mode squeezing process [CMP14], or by combining (one-mode) squeezed states of light at beam splitters

[YUvLF08, SHD⁺13]. This can be more effective than using nonlinearity for single photons. Creation and measurement of the qubits can be performed with current experimental equipment (creation via lasers and nonlinear media, measurement with homodyne detection).

17.1.6 Operational challenges

17.1.6.1 Space

Given that here gates are performed in space, by sending light through an apparatus representing the algorithm, optical quantum computers need more physical space the longer the algorithm becomes. With discrete optical elements this becomes forbiddingly large, also given the overhead of post-selection. Integrated optics creates elements of sizes comparable to other qubits (but still adding a time dimension).

17.1.6.2 Photon sources

All discrete-variables need single photon-inputs [VBR08]. Single-photon sources that are deterministic (i.e., we know when a single photon is coming) and on-demand (i.e., we can trigger injection of a photon) are a field of current research and are most likely reached with self-assembled quantum dots [SSA⁺15, SFV⁺02].

17.1.7 Research groups

Again, we focus on groups pursuing quantum computers, not the vast photonic quantum technologies as a whole, which includes, e.g., photonic quantum communication.

Jeremy O'Brien, Bristol Jeremy O'Brien is integrated in the quantum engineering center at Bristol. His group is operating a large concerted effort in integrated quantum optics on a chip and is leading the field, also by a strong vision for a photonic quantum computer reflected in many reviews cited here. One of his most successful alumni is Alberto Peruzzo, now at RMIT University, Melbourne. They have attracted significant investments through their company, PsiQuantum.

Andrew White, Queensland Andrew White is leading a large research group in Australia on both discrete and integrated quantum photonics, currently focusing on boson sampling yet developing tools for other discrete encodings.

Jelena Vuckovic, Stanford Jelena Vuckovic uses self-assembled semiconductor quantum dots as a tool in quantum photonics. This is a vast community but she stands out as being focused on quantum computing rather than communication. Her results mark much of the largest single-photon nonlinearity in the field.

Olivier Pfister, University of Virginia Olivier Pfister is a leading experimentalist in the field of the production of large continuous variable cluster states.

Theorists Terrence Rudolph at Imperial College has developed photonic cluster state quantum computing but recently moved towards quantum foundations. Groundwork for continuous variable quantum computing was laid by Samuel Braunstein, York, and continued by Nicolas Menicucci, Melbourne.

17.2 Evaluation: Single photons

Many elements of processing single photons are developed in the framework of optical quantum communication and cryptography, but can be useful also for quantum computing. Furthermore, single

photons are good candidates for flying qubits in distributed quantum computing, as they can interact with other, fixed qubits.

17.2.1 DiVincenzo criteria

17.2.1.1 Scalable qubits

Single photon sources are available using single atoms/ions [HSG⁺07, KLH⁺04], color centers in diamond [MMK⁺12, BKH⁺17], quantum dots [LDP⁺17], or optical parametric oscillators [KGPUK16] (e.g. nonlinear crystals) combined with heralding. All optical elements can in principle be integrated on chips [HDM⁺16], which makes scaling more reachable. This is still not technically mature but under development.

A problem arising in all encodings is leakage due to photon loss. Also, a number of these strategies occupy a comparably large amount of space.

17.2.1.2 Initialization

Initialization is usually done directly with creation of the photons.

Direct creation of entangled photons is also possible for example with parametric down-conversion and beam splitters, cluster states of up to 6 photons have been created in several groups, although fidelities are still to be improved [ZHL⁺16, LZG⁺07]. Larger cluster states, and on-chip generation are a matter of ongoing research.

17.2.1.3 Universal gates

A universal gate set is available, however, entangling gates are typically non-deterministic and require additional post-selection. Single-qubit gates with free-space optical elements are typically very fast and accurate.

- direct two-qubit gates $F_2 = 0.87$ [OPW⁺03]
- KLM: CZ gate with 0.68 (0.93) process (Hilbert-Schmidt) fidelity [MBB⁺16]
- on chip: CNOT with 0.94 fidelity [PCR⁺08]
- three-qubit gate: controlled-SWAP gate with fidelity 0.85 [OOT⁺17]

17.2.1.4 Coherence

The most limiting effect is photon loss, leading to leakage. Apart from loss, photons usually have high coherence times and only weak interaction with their environment. It needs to be noted that the use of media and integrated optics reduces photon lifetime [HBR⁺16].

17.2.1.5 Readout

Readout is done with photon detectors. When polarization encoding is used, the polarization information can be translated to dual-rail encoding with polarizing beamsplitters or polarization filters. Photon detectors still need to be developed in terms of photon-number resolution, efficiency, dark count rate and speed. Since single photon detectors are required in a lot of different situations, this is still a field of active research.

17.2.2 Summary

Criteria	met?	Comments
Scalable qubits	✓	
Initialization	✓	
Universal gates	✓	
Coherence	✓	
Readout	✓	

Table 17.1: Summary of DiVincenzo criteria for single photons.

17.2.3 Outlook

Single photons quantum computing fulfills all DiVincenzo criteria, thus is a level B platform, but still in an early stage as the ability for larger algorithms or error correction is not yet reached. Major efforts in scaling, loss-reduction, photon indistinguishability and deterministic gates would be necessary to lift this platform on a higher level. Also for cluster states, due to the probabilistic nature of entangling gates, scaling to larger computations is still the biggest problem.

On the other side, the increasing demand for secure communication, pushing research in single photon quantum key distribution techniques, might also bring up some benefits for universal quantum computation.

17.3 Evaluation: Continuous-variable and Gaussian encodings

Continuous-variable encoding is mainly used in the framework of one-way quantum computing, where in principle, two-qubit gates are not required on-demand, but only as an initial process in the creation of the cluster. However, keep in mind that protocols for merging clusters (as described in section 8.2.2), which will be inevitable in large algorithms, also use additional deterministic (or at least high probability) entanglement operations on demand.

An important measure is the strength of the initial entanglement gates. When using two-mode squeezing processes, this is given by the squeezing strength.

Cat states are currently still at the stage of (bad) quantum memory: They can slightly increase the coherence time, but good protected (multi-qubit) gates that make additional error correction redundant are still far. Furthermore, scaling is challenging since every qubit needs its own cavity.

17.3.1 DiVincenzo criteria

17.3.1.1 Scalable qubits

Creation of huge cluster states is not a problem, but creating them with sufficient fidelity is. States as large as 10^6 qubits in a CV cluster state [YYK⁺16] have been created, with the possibility of creating even larger states, however, most states are still below 10 qubits.

17.3.1.2 Initialization

CV cluster states can be created with optical parametric oscillators using for example entanglement between modes of a frequency comb [MFZP07, MdARC⁺14, FMP09] or time multiplexing [YYK⁺16].

Both methods set a limit on the number of “qubits”, either in time or due to the frequency window available experimentally.

Cat states or other coherent encodings in resonators can be initialized by external coupling, for example to a Transmon qubit.

17.3.1.3 Universal gates

In the one-way quantum computing scheme, gates are already included in the initialization and measurement process. For coherent state encodings, single-qubit gates already work quite well, two-qubit gates are problematic:

- Hadamard gate on coherent states with 0.94 state fidelity [TDL⁺11]
- cat states, controlled via Transmon: universal single-qubit gates with 0.985 fidelity in 1 μ s from RB and 0.9925 process tomography [HRO⁺16]. Two-qubit gates are proposed [MLA⁺14], and have just been realized [RGR⁺17] with a process fidelity of 0.83.

17.3.1.4 Coherence

Unprotected states are typically vulnerable to single photon loss, destroying for example superpositions of coherent states.

Cat code qubits are protected against photon loss, the limiting factor is the coherence time of the Transmon coupled to the resonator. Error corrected cat states reach $T_1 = 2.7$ ms [HRO⁺16].

17.3.1.5 Readout

- homodyne detection, can measure arbitrary quadrature (meaning arbitrary basis for two-mode squeezed cluster states)
- coherent states stored in oscillators can be read out with Transmons

17.3.2 Summary

Criteria	met?	Comments
Scalable qubits	✓	
Initialization	✓	
Universal gates	✓	
Coherence	✓	
Readout	✓	

Table 17.2: Summary DiVincenzo criteria for CV and Gaussian encodings.

17.3.3 Outlook

For continuous variables, two qubit gates have barely been demonstrated. While in principle protected, fidelities are actually below any known error correction threshold. Improving these will be crucial. Level B.

18 Molecular approaches

18.1 Description

Several approaches to quantum computing bind quantum bits in molecules that are tailored to the needs of quantum computing. None of these are considered to be the currently most promising, but they merit review in this report.

18.1.1 Liquid-state nuclear magnetic resonance (NMR)

The nuclei of atoms bound in molecules are well-isolated quantum systems. The qubit states correspond to orientations of a nuclear spin $1/2$ that is common in a range of atoms including hydrogen. They can be split in energy using an external magnetic field through the Zeeman effect, similar to the orientation of a magnetic dipole in a classical field. This system is characterized by the weakness of the nuclear magneton, $\mu_n = e\hbar/2m_p$ where m_p is the mass of the proton. The nuclear magneton describes the magnetic moment of a single proton. Comparing to the analogous formula for the Bohr magneton, the magnetic moment of an electron, it is thus weaker by the inverse mass ratio, $\mu_n/\mu_B = m_e/m_p \approx 1/1800$. This weak magnetic moment provides excellent protection from environmental disturbances, but we will later see that it enforces many compromises that ultimately are in the way of scaling. What are now called single-qubit rotations, manipulation of the nuclear spins by external electromagnetic fields at radio frequencies, has been achieved in the early 1950s. Readout can be done through electromagnetic induction - time-varying magnetism induces electrical current in a pickup coil. See [Abr61, Lev01] for further details.

As a consequence of the weakness of the nuclear magnetic moment, most notably, the signal of a single nucleus cannot be measured, requiring that an ensemble of identical molecules is used and measured simultaneously. Now in order to make sure that the elements of the ensemble, the individual molecules, do not decohere each other, they need to be in a liquid solution, where the fast motion of molecules in the liquid averages out detrimental intra-molecular interactions. Thus, the NMR quantum computer is essentially a test tube containing a solution of special molecules placed in a large magnet.

Now this also averages out many potential intra-molecular interactions between nuclei, leaving the “J-coupling” as the only means to drive two-qubit gates. J-coupling uses the interaction of nuclei with the electrons in the chemical bond between them, restricting it to nearest-neighbors. Also, being very indirect, J-couplings are small, with interaction strengths typically leading to 2-qubit gate durations above 100 ms. This is not a problem per se, as decoherence mechanisms due to residual interaction are weak and coherence times are long, typically in the range of seconds, allowing for a ratio of gate time over coherence time around 0.01 [CLK⁺00, VSB⁺01].

The necessity for a liquid solvent (usually water, even though quantum computing experiments with liquid carbon dioxide have been performed) implies that the temperature of the experiment must not be too low. Given that also magnetic Zeeman fields that can be applied steadily (through superconducting permanent magnets) are on the order of 10 T, the energy splitting between qubit states is much lower than temperature, hence preventing effective initialization by cooling. The largest current NMR spectrometer has a 24 T magnet which leads to a Zeeman splitting for protons of about 1 GHz, corresponding to a temperature of about 50 mK, clearly below the freezing point of suitable solvents. Given that measurement is also not projective, special initialization mechanisms have to be invented. These are:

- Pseudopure states [NC00]: Preparation of an initial state such that its deviation of the thermal state from the totally mixed state corresponds to a pure state. This is possible, but the algorithm to achieve it is exponential in the number of qubits, hence annihilating any quantum speedup.

- Active cooling methods, most notably algorithmic cooling [RMBL08, EGMW11]: This method uses an algorithm to effectively move entropy to a heat bath. It can be efficient but cannot be made fault-tolerant, hence again negating scalability.
- Dynamic nuclear polarization [MvTA⁺07]: In this method, electrons (whose Zeeman splitting is three orders of magnitude larger than that of nuclei, again, due to the ratio of effective masses) are cooled to the ground state and then a pulse sequence is used to transfer their polarization to the nuclear system. This technique is current in the focus of the NMR community at large (i. e., also its vast majority that works on NMR for chemical, biological, and medical applications). Polarizations above 50% are routinely reported for special cases in liquids (one must not confuse this with the much larger DNP in solids), corresponding to till more than 33% initialization error, so while this approach is promising it has a long way to go.

A final challenge lies in addressability: Given the small separation of spins inside a molecule by down to a single bond length (about 10^{-10}m), qubits cannot be addressed spatially, they need to be addressed by selecting their specific resonance frequency. To set these apart, one can use different elements (heteronuclear NMR) [MPB⁺15], where the need to have them bind in suitable molecules and having an abundant isotope with spin 1/2 limits the choice to about five (^1H , ^{13}C , ^{15}N , ^{19}F and ^{31}P). To go beyond, the chemical environment that induces a resonance shift can be used - full addressability then means that each atom of a certain element must have a distinct chemical environment. This makes synthesis of large molecules for NMR quantum computing increasingly difficult, even though there is no principal obstacle known.

Given these challenges, liquid-state NMR is currently not on any roadmap as a promising candidate for quantum computing, it is viewed as test-bed for methods used on other platforms [VC05]. It still was the first platform to factor 15 [VSB⁺01], an impressive number of 12 qubits have been controlled and many ideas were pioneered there first. It is quite clear that a breakthrough in initialization would bring this field back.

18.1.2 Solid-state nuclear magnetic resonance in molecules

Given the limiting role played by the liquid solvent, another candidate is quantum computing in molecular crystals. In order to overcome the inter-molecular interactions that drove the initial choice of liquid state NMR, one needs to simulate the motion that averages it out mechanically. Given the peculiarities of the magnetic dipole-dipole interaction, this is achieved by magic-angle spinning [ABE58]: The crystal is mechanically spun around an axis at the magic angle given through $\cos \theta_m = 1/\sqrt{3}$ (meaning $\theta \simeq 54^\circ$). This approach has shown some success but lose some of the exquisite control and coherence of liquid state NMR, also, given the necessity to rotate the sample and the smallness of the nuclear magneton, experiments are not conducted at the required low temperature. Thus, solid-state NMR reduces the need for advanced initialization techniques but does not eliminate it. The workhorse molecule of this approach is crystalline malonic acid which contains five qubits, gate fidelities above 99% have been reported [PFR⁺15].

Note that this is not identical to the Kane proposal, where quantum information is stored in individual nuclei but manipulated through electrons.

18.1.3 Spin quantum computing using clusters

The ability to manufacture controlled magnetic molecules that put spins at a well-defined distance provides an attractive chemical platform for quantum computing. On the one hand, one can use hollow molecules such as the Buckminster-Fullerene C-60 as cages to trap electron spins, on the other hand one can prepare magnetic molecules with electron spins at well-defined places and thus perform solid-state electron spin resonance. While this idea promises to solve the initialization problem of spin resonance, even basic fabrication and operation are still early in their development [BW08b, TGA⁺05, LL01].

18.1.4 Molecular bonds

The ability to precisely control the chemical bonds in molecules using ultra-fast lasers has promoted the idea of a molecular quantum computer [dVRT02, TdVR02]. This proposal encodes the complete quantum computer in the rich level structure of the rotational and vibrational levels of a chemical bond, i.e., a single quantum degree of freedom. Thus, n levels hold only $\log n$ qubits, making this method not scalable.

18.1.5 Research groups

These are all vast communities in general (in liquid state NMR they include hospitals) so we narrowly focus on researchers dedicated to quantum computing.

Isaac Chuang, Neil Gershenfeld, IBM / MIT The original NMR quantum computing goes back to Neil Gershenfeld of MIT Media Lab and Isaac Chuang, then at IBM Almaden. Chuang has performed the original experiment to factor 15 with NMR. Chuang is also a prolific theorist and author of the pioneering and surprisingly long-lived textbook on quantum computing. Both have moved to other fields—Chuang is described further in the ion trap chapter.

David G. Cory / Raymond Laflamme, Waterloo Raymond Laflamme is the director of the Institute for Quantum Computing at the University Waterloo, David Cory (previously at MIT) holds a highly endowed Canada Excellence Research Chair there. This makes IQC the by far leading center of all aspects of NMR-quantum computing. Laflamme is originally trained as cosmologist and has made many contributions, also theoretical, to quantum computing. David Cory is dividing his attention between quantum sensing with neutrons and spin resonance quantum computing.

Steffen Glaser, Munich Steffen Glaser is a professor of organic chemistry at TU Munich and belongs to the vast NMR school of Nobel Laureate Richard Ernst (ETH Zürich). He is developing a fully heteronuclear NMR quantum computing experiment. Also, he has pioneered the development and application of quantum optimal control techniques, notable the GRAPE algorithm [KRK⁺05, GBC⁺15] and is the director of the Virtual Facility for Quantum Control.

Dieter Suter, Dortmund Dieter Suter is a Professor of experimental physics at TU Dortmund, who also did his PhD in physical chemistry with Richard Ernst. His research interests lie mostly in new spectroscopic methods, but in the last few years he started working on quantum information processing, with a focus on dynamical decoupling schemes and implementation of simple quantum algorithms like error correction and single qubit gates.

Jonathan Jones, Oxford Jonathan Jones is a Professor of Physics at Oxford University. He is continuously pursuing liquid-state NMR for quantum computing focusing on advanced pulse techniques, which he also helps apply to other platforms.

Wolfgang Wernsdorfer, Karlsruhe Wolfgang Wernsdorfer holds a large endowed research chair (Humboldt Professorship) in Karlsruhe. He is a pioneer in quantum dynamics in molecular magnets and leads the field.

Regina de Vivie-Riedle, Munich Regina de Vivie-Riedle is a professor chemistry at LMU Munich and an optimal control theorist. She has pioneered quantum computing with molecular bonds.

18.2 Evaluation: Molecular approaches

18.2.1 DiVincenzo criteria

18.2.1.1 Scalable qubits

Example molecules in liquid-state NMR are ^{13}C -labeled trans-crotonic acid (4 qubits) [XHL⁺17], l-histidine molecule with 14 spin-1/2 nuclei (12 qubits+1 qutrit) [NMR⁺06], unlabeled chloroform (1 qubit), and ^{13}C tris(trimethylsilyl)silane-acetylene (3 qubits) [RLL09]. Examples for solid-state NMR include malonic acid molecules (3 qubits) [BMR⁺06] and (5 qubits) [PFR⁺15], and the latter having much stronger dipole coupling to its neighbors, leading to broader resonance peaks. Main problem is that the number of qubits is limited by molecule size. Also, the connectivity is determined by the molecule.

18.2.1.2 Initialization

Initialization is the thermal equilibrium and further polarization in external magnetic field [NC00], which in the liquid state is close to a fully mixed state. State preparation of pseudopure states take rf pulse sequences of milliseconds [BMR⁺06]. Most notably, the overhead here scales exponentially.

Solid-state NMR can be performed at lower temperatures requiring less or no preparation overhead.

18.2.1.3 Universal gates

Gate times lie in the millisecond regime and are controlled with radio frequency pulses (GRAPE optimization). RB reveals single-qubit errors of 0.0126 and two-qubit CNOT errors of 0.0177 in a 4-qubit liquid-state NMR processor [XHL⁺17]. Gate durations are 1ms for single-qubit gates and 7.5--15ms for CNOT. Better RB values in liquid-state NMR can be found in [RLL09], with 0.9999 for single-qubit gates and 0.995 for two-qubit gates, although benchmarking of single and two-qubit gates was performed with different molecules. Pulse durations below 450 μs with fidelities above 0.9 in solid-state NMR [BMR⁺06]. GRAPE optimized pulses below 1 μs and with average two-qubit fidelity ≥ 0.99 .

18.2.1.4 Coherence

Coherence times range from 100s of ms to several seconds [RLL09]. 4-qubit liquid-state NMR processor has $T_2^* = 400\text{ms}$ [XHL⁺17]. High decoherence in Xe around 1000s [HGL⁺98], but is not scalable to two-qubits. Solid-state decoherence is about $T_1 \approx 300\text{s}$ and $T_2^* \approx 2\text{ms}$ [BMR⁺06]. At room temperature, reported coherence times are $T_1 = 27\mu\text{s}$, $T_2 = 5\mu\text{s}$ and $T_2^* = 28\text{ns}$ [PFR⁺15]. Tailored refocusing and perfect controls would eventually lead to $T_2 = 2T_1$ [BCC⁺07].

18.2.1.5 Readout

Measure magnetic moment by induction.

18.2.2 Summary

Criteria	met?	Comments
Scalable qubits	?	Depends on molecule size
Initialization	?	Pseudopure states not scalable (liquid state)
Universal gates	✓	Designed GRAPE pulses are robust to inhomogeneous line broadening
Coherence	✓	around 0.01, required 10^{-4} , due to slow gates
Readout	✓	

Table 18.1: Summary DiVincenzo criteria for liquid- and solid-state molecular approaches.

18.2.3 Analysis and outlook

NMR more like a testbed for methods like RB and sin, due to its yet unresolved scalability issues, but decades of experience with experimental methods in particular for the liquid state.

Solid-state NMR in molecular crystals has the potential to overcome this initialization obstacle, on the expense of losing also many of the benefits of liquid state NMR such as high coherence and good control. The most promising spin resonance incarnations currently are NV centers and single dopants rather than molecular crystals.

19 Exotic qubit candidates

19.1 Electrons on helium

This approach uses cold liquid (but not superfluid) helium to trap and control electrons [SFD⁺10]. Current efforts focus on trapping, manipulating, and detecting single electrons, which is challenging given the limited experience with this system. Ensemble systems are much more promising [YFK⁺16] but currently there is no scalable architecture known for them.

19.2 Quantum nano- and optomechanics

Nanomechanical oscillators have recently been put into the quantum regime [Cle10, TLA⁺11]. These elements need a lot of overhead meaning that they can be used for nice applications, such as frequency conversion between microwave and optical frequencies [BVAC13] which is necessary for building quantum repeaters, but not for the types of quantum computers relevant to this study.

19.3 Research groups

Mark Dykman, Michigan Mark Dykman is Professor in the Condensed Matter Physics theory group of the Michigan State University with a broad research spectrum, ranging from fluctuation phenomena to correlated electrons and quantum information. He is also interested in a lot of exotic topics: He even wrote a PRL on the statistics of disease extinction, and thus also works on the field of electrons in liquid helium.

David Schuster, Chicago David Schuster is an assistant Professor at the Institute for Molecular Engineering at the University of Chicago. His research is dedicated to finding new ways to think about superconducting quantum computing, to which he makes innovative non-integrated contributions. One approach that he investigates in his lab is a hybrid system of electrons on liquid helium above superconducting cavities. He is currently the only visible experimentalist in this area.

Andrew Cleland, Chicago Andrew Cleland is Professor for Molecular Engineering Innovation and Enterprise at the University of Chicago. In his group, he has demonstrated the coupling of mechanical resonators to superconducting qubits where they control individual phonons and perform ground state operations. More recently, he has moved studies on the conversion between microwave and optical frequencies.

John Teufel and Raymond Simmonds, NIST John Teufel and Raymond Simmonds are project leaders in the Advanced Microwave Photonics Group at the National Institute of Standards and Technology (NIST), a non-regular agency of the United State Department of Commerce. The group works on the coupling of superconducting and mechanical resonators and on cooling micromechanical oscillators with the purpose of storing quantum information. They were among the first (after Cleland, at that time in Santa Barbara, who had lower quality factors in his experiment) to cool a high-quality factor mechanical oscillator to the quantum ground state.

20 Big actors and programs

Most actors were described along their platforms. This is adequate, as most organizations still follow a multitude of approaches, being a collaborative place that is slightly more than the sum of its parts—but has these parts very visible.

20.1 National and international programs

This describes programs organized by major sponsors. We are summarizing publicly available information (we are inside some of these programs).

20.1.1 IARPA

The Intelligence Advanced Research Projects Agency (IARPA) is part of the Office of the Director of National Intelligence (ODNI) which is overseen by the Department of Homeland Security (DHS) of the United States of America. It operates large programs in scope and ambition and funding levels. These programs are not exclusively funding US-based groups.

Logical Qubit Program (LogiQ) The LogiQ program aims at building logical qubits that perform better than the underlying physical qubits and asks to develop modular and scalable designs based on the understanding of the program performers. Its four consortia are led by TU Delft, Duke University, IBM, and University of Innsbruck, hence it concentrates on ion trap and superconducting qubit platforms. LogiQ is widely considered to be the far most advanced program aiming at building a fault-tolerant circuit-based quantum computer. The available funding level has not been announced

Quantum-Enhanced Optimization (QEO) The Quantum-Enhanced Optimization aims at building a quantum annealer / adiabatic quantum processor that does not have the limitations of the d-Wave machine, hence including non-stochastic and many-body coupling as well as coherent qubits. This will be crucial in evaluation the power of quantum annealers and will be the first quantum annealing platform that could run the Shor algorithm.

20.1.2 US Department of Energy

Advanced Quantum-Enabled Simulation (AQuES): The Lawrence Berkeley National Laboratory (Berkeley Laboratory) receives funding of \$ 1.5 Mio from the DOE for leading the DOE's AQuES testbed project. Within five years from October 2017 on, the project aims to examine the behavior of near-term quantum devices for computation. The testbed comprises a quantum processor based on superconducting qubits, supported by modular classical control and logic, and a flexible software stack consisting of low-level programming, validation, verification and benchmarking tools. The project is focussing on: (i) quantum and classical hardware scaling and optimization, including the delineation and exploration of the available co-design tradespace for near-term quantum computing algorithms for simulation, (ii) robust and portable characterization methods for quantum verification and validation, including platform-independent metrics to quantify the computational power of quantum simulators, and (iii) scalable hardware and software interfaces. This project is a collaboration between Berkeley Lab and Lawrence Livermore National Laboratory (LLNL), Berkeley Lab will focus on transmon arrays and LLNL will focus on multimode cavity arrays.

Advanced Quantum Testbed (AQT): The goal of the AQT project is to establish a multi-partner scientific collaboration aiming at building a platform for answering basic outstanding questions about quantum computing. AQT will operate as an open resource for the community, allowing external researchers to evaluate superconducting architectures developed by testbed staff and collaborators for simulations in chemistry, materials, and other areas of computation. Furthermore, AQT will afford industry researchers to

explore the most promising approaches, which can then be transformed into finished commercial products by industry. The project is funded by the DOE with \$ 30 Mio with a duration of five years started in October 2018.

20.1.3 Other US agencies

The Laboratory of Physical Sciences (LPS) is a US-Government laboratory at the University of Maryland that has in-house research but also operates funding programs in collaboration with the US Army Research Office (ARO) that are broader than IARPA's. The Defense Advanced Research Projects Agency (DARPA) has a strong track record in computer science and has been a sponsor of quantum technologies in the past but is not very active in quantum computing any more. The Air Force Office of Scientific Research (AFOSR) is running a smaller quantum computing program, so is the Office of Naval Research (ONR).

20.1.4 United Kingdom Quantum Technology Hubs

The United Kingdom was the first European Country to launch a large-scale quantum technology program. Remarkably, only one of the four hubs (NQIT, Networked Quantum Information Technologies, based in Oxford, Managed by Ian Walmsley) is even looking at quantum computing, but only as one of its areas of engagement. Ion trappers Winfried Hensinger (Sussex) and David Lucas (Oxford) are funded by this instrument.

20.1.5 EU Quantum Technologies Flagship

The European Union has funded quantum computing for a long time through its Future and Emerging Technologies (FET) program, which has not caught up with transition from basic science to engineering and development. The Quantum Technologies Flagship was announced in 2016 and is has started on October 1st 2018, investing 1 Billion Euros over a ten year period. Quantum computing is one of its four application pillars. Among the first selected projects there are two projects aiming at building quantum computers. **An Open Superconducting Quantum Computer (OpenSuperQ)** The goal of the collaborative research project OpenSuperQ is the development of a quantum computing system and its sustainable availability at Forschungszentrum Jülich for external users. The hardware, which will be based on superconducting integrated circuits with up to 100 qubits, will make the computer belong to the leading platforms in the world, and the open and integrative approach is globally unique for the OpenSuperQ. Additional necessary technological infrastructure, including a control system and user-friendly cryogenics, will be supported. The software stack will be integrated from user access all the way to low-level control. Being designed as an all-purpose quantum computer, its particular tasks are quantum simulation in chemistry and material science, optimization and machine learning. The project receives funding of EUR 10.33 Mio from the EU Research Framework Program Horizon 2020. It started on Oct 1st in 2018 with an initial duration of three years. Coordinator: Saarland University.

Advanced quantum computing with trapped ions (AQTION) The AQTION project aims at developing a scalable European quantum computer with registers with up to 50 qubits. The qubits will be represented by single-charged atoms or ions, which can be controlled individually with a high performance. Possible applications will range from chemistry via energy distribution to chip-layout optimization. During its duration from October 2018 til September 2021 the project will receive funding of EUR 9.59 Mio from the EU Research Framework Program Horizon 2020. Coordinator: University of Innsbruck.

20.1.6 BMBF

The German Ministry for Education and Research (BMBF) has operated a flagship-related program named Qutega. It has been developed into a framework program of about 0.65 Geuros. The details of its quantum computing pillar are under development.

20.2 Academic institutions

20.2.1 North America

20.2.1.1 USA

NIST The National Institute of Standards of Technology (NIST) has a primary mission as the metrology and standards organization of the US. This mission is interpreted broadly. The high precision requirements of quantum computing are a natural match for researchers with a metrological background. Two NIST sites are primarily invested in quantum research: NIST Boulder, CO, with the Wineland lab (see ion traps) and NIST Gaithersburg, MD, which forms the Joint Quantum institute with the University of Maryland, see below.

DOE Laboratories The US Department of Energies operates a system of dedicated national laboratories, comparable to the Helmholtz-System in Germany. Some of these are engaged in quantum computing: *Lawrence Livermore* supports superconducting qubits research mostly by modeling materials. Its quantum detection program was set back by a recent scandal of scientific misconduct. The *Sandia National Laboratory* fabricates ion traps for scalable quantum computing and supports IARPA with a verification and validation team whose work informs chapter 6. The *Los Alamos National Laboratory* has made many pioneering contributions in quantum computing but is currently not very active.

Joint Quantum Institute / Joint Center for Quantum Information and Computer Science The Joint Quantum Institute (JQI) brings together the University of Maryland at College Park and NIST at Gaithersburg. It is interdisciplinary in bringing together Physicists and, with its sister institute, the Joint Center for Quantum Information and Computer Science, Computer scientists. Its research focus areas are trapped ions and spins in semiconductors. It is directed by AMO physicist Gretchen Campbell and solid-state physicist Fred Wellstood, its sister institute by computer scientist Andrew Childs and physicist Jacob Taylor. The center collaborates with the center for nanosystems and advanced materials.

MIT Center for Ultracold Atoms / Keck Center MIT's center for ultracold atom focuses on using said platform for quantum computing and simulation. Quantum computing theory is advanced in the Keck Center for Extreme Quantum Information Theory, led by Seth Lloyd and Jeff Shapiro.

Wisconsin Institute for Quantum Information This institute brings together the very strong physics groups in neutral atoms, Si-based semiconductor systems, and Josephson qubits at UW Madison.

Yale Quantum Institute The Yale Quantum Institute brings together the various quantum computing groups at Yale, primarily in superconducting qubits.

Berkeley Quantum Information & Computation Center Founded in January 2017, this center brings together the high-caliber quantum computing researchers at UC Berkeley, including superconducting qubits measurement specialist Irfan Siddiqi, ion trapper Hartmut Häffner, semiconductor grower Thomas Schenkel, and theorists Birgitta Whaley (theoretical chemistry) and Umesh Vazirani (computer science).

20.2.1.2 Canada

Institute for Quantum Computing, University of Waterloo IQC is a large center for quantum computing and other quantum technologies. Most of their 26 faculty were hired specifically for this center and they plan to grow to 40 faculty. They are funded by a large private endowment of at least 100 M\$C and matching funds from both university and government. Next to their outstanding programs in quantum key distribution and theory and mathematics for computer science, they have various activities in quantum computing. The labs of David Cory and Raymond Laflamme are leading in molecular NMR and three relatively young groups (Adrian Lupascu, Matteo Mariantoni, Christopher Wilson) pursue superconducting qubits. While none of them is present in any of the cutting edge programs in quantum computing experiment, they have trained many leaders of the field including ex-postdocs Austin Fowler (superconducting qubits and error correction, now at Google) and Jay Gambetta (superconducting qubits theory, now at IBM).

Équipe de Recherche en Physique de l'Information Quantique, Université de Sherbrooke This research center focuses on solid-state qubits, primarily semiconductors. With theorist Alexandre Blais (Josephson qubits), David Poulin (error correction) and experimentalists Andy Sachrajda (also national research center Ottawa) and Michel Pioro-Ladrière (quantum dots) they have a focused and mostly young team - spectacular hires are announced making this a place to watch. In the province Quebec they join forces with universities in Montréal forming the Institut Transdisciplinaire d'Information Quantique (INTRIQ).

Canadian Institute for Advanced Research (CIFAR) Quantum Information program The Canadian Institute for Advanced Research—a joint public/private initiative that serves the function of an academy, has a nationwide program in quantum computing that runs conferences and sponsors research fellows. It is aimed at bringing theory and experiment together.

20.2.2 Europe

The European Institute of Molecular Magnetism (EIMM) This European network brings together researchers in the field of magnetic clusters and molecules, an exotic molecular platform for quantum computing.

20.2.2.1 United Kingdom

As mentioned above, the United Kingdom operates a large program called quantum technology hubs. Next to these, this activity has spawned a number of local centers.

Sussex Centre for Quantum Technologies, Brighton With a strong activity in cold atomic gases for simulation and in quantum sensing in the center, ion trapper Winfried Hensinger is part of this center.

Oxford Quantum, Oxford UK A broad and high-level research center including materials science. Leading quantum computing experimentalists like David Lucas and Andrew Steane are in there, as are theorists like David Deutsch and Artur Ekert.

London Centre for Nanotechnology (LCN), London This vast center between the University College London and Imperial College has a component that looks at quantum computing, including Paul Warburton (one of the few European experimentalists looking at quantum annealing) and semiconductor spin resonance researcher John Morton.

Joint Quantum Centre, Durham-Newcastle Northern England quantum technology center with notable quantum computing activity in Rydberg atoms.

Centre for Quantum Photonics, Bristol This center pursues photonic quantum computing. Next to Jeremy O'Brien (described above) there is a whole center with more scientists supporting this activity, with John Rarity being the other very prominent researcher.

20.2.2.2 Netherlands

QuTech, Delft QuTech has already been described above under “semiconductors”. It brings together scientists with engineers and TU Delft with other research institutes. QuTech is complementing this experimental activity with a strong theory team (Stephanie Wehner, Barbara Terhal). It is collaborating deeply with Intel, mostly on engineering and systems integration, as well as with Microsoft. Founding director Leo Kouwenhoven has recently moved to Microsoft (who are planning to build a hardware laboratory in or close to Delft), the current interim director is Ronald Hanson.

20.2.2.3 France

As is typical, the French research landscape is only in part carried by Universities and more by other government-funded research institutions, that we still list here under academic centers. We are listing those with large local or national centers (all other significant French researchers are listed at their specific platforms).

Commissariat à l'énergie atomique et aux énergies alternatives (CEA) Saclay and Grenoble This large research system is well funded also for activities that are outside its core mission. Most notably, the Quantronics group at their site in Saclay (south of Paris) has historically been a leader and is still a strong player in Josephson qubits. Its site at Grenoble (the Institute Nanosciences et Cryogenie) has recently emerged with a leading activity in Si-based quantum dots, which is based on collaboration with the Laboratoire d'électronique des technologies de l'information (LETI), a pre-commercial semiconductor foundry.

Centre national de la recherche scientifique (CNRS), France Most of the quantum researchers listed in this document are affiliated with the CNRS. Two centers should be mentioned separately: The Institut d'Optique (Paris-South) is the home of a number of prominent quantum researchers and has its own graduate school that is recognized as a grand école. Leading Rydberg experimentalist Philippe Grangier leads the quantum optics group at their laboratoire Fabry. The Institut Neel in Grenoble is a nanoscience institute that is active in solid-state qubits.

20.2.2.4 Germany

The German research landscape is traditionally rather decentralized. Interestingly, there are large activities in quantum technologies but focuses on applications other than quantum computing. Quantum computing is considered to be basic research where in other jurisdictions it is funded as applied research. With the BMBF QuTega program there is a possibility for this to change. Centers that should be mentioned:

Max Planck-Institute for Quantum Optics (MPQ), Garching This research flagship has leading groups in neutral atoms (Immanuel Bloch, Gerhard Rempe) and outstanding theory group of Ignacio Cirac.

Jülich-Aachen Research Alliance (JARA) - Institut für Quanteninformation, Aachen This is a concerted effort between a large Helmholtz Laboratory and University to pursue quantum computing with solid state systems. Their quantum dot activity is centered around Hendrik Bluhm and their theory effort in both semiconductors and Josephson qubits around Humboldt Professor David DiVincenzo. The center also has built an impressive enabling research and engineering complement.

20.2.2.5 Switzerland

The Quantum Engineering Center, ETH Zürich Being a top technology school, ETH Zürich has a number of strong actors in quantum science and technology (mentioned here: Home in ion traps, Wallraff in superconducting qubits). They are bundled in this center, which is set to branch out into more serious engineering.

Basel Center for Quantum Computing and Quantum Coherence (QC2), Basel This center is focused around semiconductor platforms and theory-driven.

National Center of Competence in Research (NCCR) 'QSIT - Quantum Science and Technology': Zürich, Basel, Genf, Lausanne, IBM This nationwide multi-University center brings together all Swiss quantum groups, with quantum computing being pursued by Loss in Basel and the ETH Groups plus IBM.

20.2.2.6 Austria

Austria is a leading force in quantum computing, primarily driven by Vienna and Innsbruck.

Institut für Quantenoptik und Quanteninformation (IQOQI), Innsbruck + Wien This institute of the Austrian Academy of Sciences brings together the strong institutes at both these geographical centers. It is AMO-centered (with Rainer Blatt, Jörg Schmiedmayer, Anton Zeilinger and Peter Zoller being the most prominent names) but develops other platforms as well.

20.2.2.7 Denmark

Center for Quantum Devices (QDev), Copenhagen This center is at the Niels Bohr Institute at the University of Copenhagen. It is focused on semiconductor physics. With its director Charles Marcus being its most prominent member, it has a strong complement of younger faculty, including Ferdinand Kuemmeth and a number of theorists with a condensed-matter nanophysics background. The NBI also houses a strong group in quantum optics and a commercialization platform named QUBIZ. The center closely collaborates with and is sponsored by Microsoft.

20.2.2.8 Sweden

Linnaeus Centre on Engineered Quantum Systems (Linneqs), Chalmers University of Technology This center, located at a private technical university in Gothenburg, brings together researchers in Josephson qubits and includes an engineering qubits.

20.2.2.9 Russia

Russian science is faced with resurrecting after the brain drain of the 1990s. It does so by new universities and centers that complement the traditional academy research landscape. The Russian Mega Grant program is designed to attract foreign-based (mostly but not exclusively russian-born) scientist.

Russian quantum center, Moscow This center has a particular strong complement in superconducting qubits (with Alexey Ustinov (also Karlsruhe Institute of Technology), Oleg Astafiev (also Royal Holloway University of London), and Evgeni Il'ichev (also IPHT Jena)). While resources seem to be good, it is not clear what proportion of these researchers output is really based in Russia and it is certainly not spectacular. Other activities in quantum technologies are not focused on quantum computing.

20.2.3 Asia

20.2.3.1 China

Some scientific activities in China have reached a spectacular level. This certainly includes their quantum key distribution program included its space-based component. In the area of quantum computing, this is not developed to the same level. A number of western-educated scientists have returned to China.

National Laboratory of Microstructures, Nanjing This laboratory brings together a number of researchers in solid-state qubit and appears to be the one experimental laboratory in China that published on a level that is internationally competitive.

20.2.3.2 Japan

RIKEN: Center for Emergent Matter Science RIKEN is a research-only campus. It has strong effort in superconducting qubits with theorist Franco Nori and experimental pioneer Jaw-Shen Tsai. It also has attracted prominent researchers as a secondary affiliation including leading quantum dot researchers Seigo Tarucha (Tokyo), Daniel Loss (Basel) and Josephson Junction pioneer Yasunobu Nakamura (Tokyo). Nakamura and Tarucha are also lead researchers at Institute for Nano Quantum Information Electronics (NanoQuine), Tokyo.

20.2.3.3 Singapore

Centre for quantum technologies, National University of Singapore This is a large and spectacular center with strong principal investigators, some of them with joint appointments in other places. However, experimental quantum computing is not at the center of attention, with ion trapper Dzimtry Matsukevich being the most visible researcher.

20.2.3.4 India

While for a long time, India's fabulous undergraduates from institutions like IIT or IIS all left, some are staying and returning now. None of them leads in quantum computing, some have potential to take off.

Tata Institute for Fundamental Research, Mumbai This is an independent research institute that has attracted Josephson Junction researcher R Vijayaraghavan (Siddiqi alumnus) who is clearly globally competitive and also has a quantum information group in mathematics.

20.2.4 Australia

Australia has a strong focus in quantum computing, both with single donors in semiconductors and optical quantum computing. There are two large national centers.

ARC Centre of Excellence for Engineered Quantum Systems (eQUS), Australia This center is mostly located in Sydney. While quantum computing is not stated as one of their research areas, their quantum measurement and control team is very visible in this area. Experimentally, ion trapper Michael Bierczuk is very visible and so is David Reilly, quantum dot researcher who pursues strong technology development, on the theory side, Steve Flammia makes important contributions to benchmarking, and Stephen Bartlett, Jason Twamley, and Andrew Doherty are very prolific theorists active in a number of platforms.

Centre for Quantum Computation and Communication Technology, (UNSW, Melbourne, Canberra, Griffith, Queensland, Sydney) This center brings together the semiconductor experimentalists in Australia, specifically leading research in single-impurity spin resonance. Most visible people are Andrea Morello and Michelle Simmons. They stand on the shoulders of people like David Jamieson (who works on fabricating these delicate samples) and theorists like Lloyd Hollenberg.

20.2.5 Latin America and Africa

These two continents are actively working on quantum research but—largely due to funding and difficulty to reliably run facilities—focus on theory. Strong theory groups can be found at the University of KwaZulu Natal (South Africa), Campinas and the Federal University of Rio (Brazil).

20.3 Companies

Relevant companies have all been described along the platforms. The only company that is pursuing quantum computing beyond a single platform is Microsoft. They are operating a quantum algorithms research group in Redmond, Station Q—a group devoted to the study of topological quantum computing—in Santa Barbara, and are currently building a laboratory in Delft. They are collaborating with semiconductor and superconducting qubit platforms.

21 Global conclusions

At this point in time, quantum processors that have been realized are far from those needed to attack cryptography, see Figure 21.1. Given the steep and accelerating development in particular after industry has entered the field, it still makes sense to extrapolate what a concerted activity would be capable of.

As scaling is governed by necessary error correction and as the most demanding step of error correction permits a space-time trade-off, this strongly depends on how much time is allowed for decryption—attacking within one day requires an astronomical number of qubits, but 100 or 1000 days of time allow to reduce their number. The size of the cryptosystem that is being attacked is less important—the slow scaling of Shor and dlog with code size are in continued to the fault-tolerant regime.

So even if the development is fast, a quantum computer that would be capable of attacking cryptography would be a large piece of infrastructure - an optical table of the size of a soccer field or an array of large cryostats containing the rare isotope ^3He .

It is interesting to extrapolate what a concerted research effort for building a quantum computer would be able to reach in a few years—a program where an industrialized nation pours a large part of its research and development activities in a single project comparable to the Apollo program and the Manhattan project. Assuming that things go well on this way—steady progress in error rates, economical solutions for bulky peripherals, large chips, improved cryogenics, one can expect that a Josephson processor with one million qubits and an error rate of 10^{-4} is feasible. This would allow to attack 2048 Bit RSA in a couple of hundreds of days. A faster attack would require up to 1000 of those units to be coupled in order to do this in a few days, which requires technological problems that are so far not solved satisfactorily.

Predicting the time scale at which this can be reached has a lot of uncertainties. While one might be able to extrapolate from the current growth in cloud quantum computing as respect to qubit size, we have not seen the scaling of error rates yet, which would be equally important for such an extrapolation.

21.1 Update 2019

2018 and early 2019 have seen strong activities in quantum computing, specifically in the area of non error-corrected near-term applications and hardware. Notably there have been two results on factoring without error correction, which however do not indicate speedup relative to the classical case. Within the known efficient algorithmic paradigms, various innovations have led to small improvements in prefactors of scaling laws.

While not as fast, this has left its mark also on cryptanalytically relevant fault-tolerant quantum computing, motivated by the perspective of practical implementation. These include lattice surgery, hardware-adapted decoders, and a deeper understanding of the correction of coherent errors. While indicating that fault tolerant quantum computers could be closer than anticipated, there are no conclusive results on their large-scale consequences (yet).

On the side of quantum computer hardware performance, there has been gradual progress in most platforms and strong progress in the field of spin qubits in Si-based quantum dots, putting them from the low to the high end of level B.

21.2 Update 2020

2019 and early 2020 have seen a number of important improvements as the community expands. First and foremost, the demonstration of quantum supremacy by Google is an important demonstration of reaching

the next level of quantum processor size and overcoming a number of engineering challenges. From the view of cryptanalysis, this is only one step in a long sequence and the next step should focus on lowering the error rate and characterizing the errors, specifically, if they are predominantly unitary or stochastic. Another general observation is that on the algorithmic side, there are more heuristics than previously. While it is unlikely that any of them disrupts cryptanalysis, definite statements require access to quantum processors. Finally, novel strategies to implement operations involving measurements give new metrics than the traditional gate counts, yet, they currently do not influence our conclusions.




Figure 21.1: Comparison of algorithmic demands with currently achieved hardware performance. The plot shows required resources as number of qubits times rounds of error correction in the surface code for dlog (blue) and factoring (orange) for common key sizes (s , Chapter 9.5 and Figure 9.1) as a function of the physical error rate p . The squares show current realizations assuming one day run time (solid) or 100 days (empty), the yellow area shows expected near-term progress. Both scales are logarithmic. Note that unlike Chapters 13 and 14, we are assuming uniform error rates for gates, initialization, and readout, rather than those obtained from current experiments, see chapter 7.4.1.

Appendix

22 Details on error models

Realistic quantum computers have a large variety of error channels with a phenomenology given in section 6.5. While in most of the large-scale error correction studies a simplified model of purely depolarizing errors is assumed, we outline in section 7.4.1.3 the current status of literature in correcting more general errors. In this section, we show pieces of calculations that are intended to illustrate the effect of error correction on different models.

In the sections of this chapter, we discuss various aspects of this error classification and their impact on error correction, that are more detailed and technical than those treated in the main text. We compare random Pauli errors and arbitrary unitary errors, both as physical operations on the physical qubits of some encoding. We do this for the three-qubit repetition code that protects against bit flips. In order to keep the notation real-valued, we assume that we would like to protect ourselves against Y errors (we use capital letters for the corresponding Pauli matrices). If we did the same with the seven-qubit Steane code, we would be able to present an analogous argument for all single-qubit Pauli errors (and all single-qubit unitary errors). We assume that we would like to protect an arbitrary superposition $|\psi_0\rangle = a|0\rangle + b|1\rangle$, where a and b are some unknown complex coefficients such that the state is normalized. For the three-qubit repetition code that protects against bit flips, the logically encoded state is $|\psi_{0,L}\rangle = a|000\rangle + b|111\rangle$. We also discuss how unitary errors lead to syndromes analogous to those of random Pauli errors and the role of systematic errors. We conclude by asking what the limitations to recovering errors actually are.

22.1 Treatment of random Pauli errors

We start with Pauli errors, but express this more standard case in a language compatible with the case of unitary errors following in the next section. We apply random bit-flip channels Y_i to the physical qubits

$$\rho \mapsto B_i[\rho] = (1-p_i)\rho + p_i Y_i \rho Y_i \quad ,$$

with the probability p_i that a bit-flip occurs on the i th physical qubit. Thus, our initial state described by a pure density matrix $\rho_{0,L} = |\psi_{0,L}\rangle\langle\psi_{0,L}|$ gets corrupted into

$$\begin{aligned} \rho_{C,L} &= B_1 \otimes B_2 \otimes B_3[\rho_{0,L}] \\ &= (1-p_1)(1-p_2)(1-p_3)\rho_{0,L} + (1-p_2)(1-p_1)p_3 Y_3 \rho_{0,L} Y_3 \\ &\quad + (1-p_1)p_2(1-p_3) Y_2 \rho_{0,L} Y_2 + (1-p_1)p_2 p_3 Y_2 Y_3 \rho_{0,L} Y_2 Y_3 \quad . \\ &\quad + p_1(1-p_2)(1-p_3) Y_1 \rho_{0,L} Y_1 + p_1(1-p_2)p_3 Y_1 Y_3 \rho_{0,L} Y_1 Y_3 \\ &\quad + p_1 p_2(1-p_3) Y_1 Y_2 \rho_{0,L} Y_1 Y_2 + p_1 p_2 p_3 Y_1 Y_2 Y_3 \rho_{0,L} Y_1 Y_2 Y_3 \end{aligned}$$

Now we perform a measurement of the two stabilizers $Z_1 Z_2$ and $Z_2 Z_3$. These commute and can hence be measured simultaneously. The four possible parity measurement outcomes are:

- Even-even parity (eigenvalues $+1, +1$): we do not detect an error and thus do not perform any recovery operation, post-measurement density matrix reads

$$\rho_{00} = \frac{1}{p_{00}} [(1-p_1)(1-p_2)(1-p_3)\rho_{0,L} + p_1 p_2 p_3 \rho_{\perp}] \quad ,$$

with total probability $p_{00} = (1-p_1)(1-p_2)(1-p_3) + p_1 p_2 p_3$ and $\rho_{\perp} = Y_1 Y_2 Y_3 \rho_{0,L} Y_1 Y_2 Y_3$.

- Even-odd parity (eigenvalues $+1, -1$): If the first stabilizer detects even ($+1$) and the second detects odd (-1) parity, the error recovery operation is iY_3 and we obtain the state

$$\rho_{01} = \frac{1}{p_{01}} [(1-p_1)(1-p_2)p_3 \rho_{0,L} + p_1 p_2 (1-p_3) \rho_{\perp}] \quad ,$$

with total probability $p_{01} = (1-p_1)(1-p_2)p_3 + p_1 p_2 (1-p_3)$.

- Odd-even parity (eigenvalues -1,+1): post-measurement density matrix after recovery operation iY_1

$$\rho_{10} = \frac{1}{p_{10}} [p_1(1-p_2)(1-p_3)\rho_{0,L} + (1-p_1)p_2p_3\rho_{\perp}]$$

- with total probability $p_{10} = p_1(1-p_2)(1-p_3) + (1-p_1)p_2p_3$.

- Odd-odd parity (eigenvalues -1,-1): post measurement density matrix after recovery operation iY_2

$$\rho_{11} = \frac{1}{p_{11}} [(1-p_1)p_2(1-p_3)\rho_{0,L} + p_1(1-p_2)p_3\rho_{\perp}]$$

with total probability $p_{11} = (1-p_1)p_2(1-p_3) + p_1(1-p_2)p_3$.

So we nicely see that if two or more of these independent errors occur at the same time, error correction creates a logical error. The resulting map is characterized by

$$B_c[\rho_{0,L}] = \sum_{\sigma, \sigma'=0}^1 p_{\sigma\sigma'} \rho_{\sigma\sigma'} = (1-p_c)\rho + p_c\rho_{\perp}$$

with $p_c = p_1p_2p_3 + p_1p_2(1-p_3) + (1-p_1)p_2p_3 + p_1(1-p_2)p_3 = p_1p_2 + p_1p_3 + p_2p_3 - 2p_1p_2p_3$. So if all errors are of the same order, we have turned the error probability from p on the physical level to $3p^2$ for the logical encoded qubit. In conclusion, we have created a bit-flip channel with an error probability that is lower than the original error probability provided we are below a threshold.

22.2 Treatment of unitary errors

We would like to demonstrate the impact and quantification of unitary errors (over- and under-rotations), i.e., the case of only systematic errors. A unitary Y error—an arbitrary rotation around the Y -axis—can be written in the $\{|0\rangle, |1\rangle\}$ -Basis as

$$U_0 = \exp[-i\theta_0 Y] = \begin{bmatrix} \cos \theta_0 & -\sin \theta_0 \\ \sin \theta_0 & \cos \theta_0 \end{bmatrix} .$$

Hence the corrupted state will have the form

$$U_0|\psi_0\rangle = (a \cos \theta_0 - b \sin \theta_0)|0\rangle + (a \sin \theta_0 + b \cos \theta_0)|1\rangle = \cos \theta_0|\psi_0\rangle + \sin \theta_0|\psi_0^{\perp}\rangle ,$$

where we have defined the orthogonal complement $|\psi_0^{\perp}\rangle = a|1\rangle - b|0\rangle$. Clearly, the probability of an error (after measuring) here is $\sin^2\theta_0$.

Now we would like to illustrate the protection of the Shor code that is offered here. We again start from the now logically encoded state $|\psi_0\rangle_L = a|000\rangle + b|111\rangle$ and apply the unitary error to all three qubits to obtain $|\psi_{C,L}\rangle = U_1 \otimes U_2 \otimes U_3|\psi_0\rangle_L$. Note that this formally corresponds to a rotation error in quantum memory but can equivalently represent uncorrelated unitary single-qubit errors that can have occurred before. We also point out that unitary errors can be more complex in general, hence they need to be handled with more complex codes. However, that would blur the current example. With the short-hand notations $c_i = \cos\theta_i$ and $s_i = \sin\theta_i$, we write out the corrupted state

$$\begin{aligned} |\psi_{C,L}\rangle &= (ac_1c_2c_3 - bs_1s_2s_3)|000\rangle + (ac_1c_2s_3 + bs_1s_2c_3)|001\rangle \\ &+ (ac_1s_2c_3 + bs_1c_2s_3)|010\rangle + (ac_1s_2s_3 - bs_1c_2c_3)|011\rangle \\ &+ (as_1c_2c_3 + bc_1s_2s_3)|100\rangle + (as_1c_2s_3 - bc_1s_2c_3)|101\rangle \\ &+ (as_1s_2c_3 - bc_1c_2s_3)|110\rangle + (as_1s_2s_3 + bc_1c_2c_3)|111\rangle \end{aligned} .$$

This state is no longer in the code space spanned by $|000\rangle$ and $|111\rangle$. Again we perform a measurement of the two stabilizers Z_1Z_2 and Z_2Z_3 , with the four possible outcomes:

- Even-even parity (eigenvalues +1, +1): we obtain the post-measurement state

$$\begin{aligned} |\psi_0\rangle &= \frac{1}{\sqrt{p_{00}}}((ac_1c_2c_3 - bs_1s_2s_3)|000\rangle + (as_1s_2s_3 + bc_1c_2c_3)|111\rangle) \\ &= \frac{1}{\sqrt{p_{00}}}(c_1c_2c_3|\psi_{0,L}\rangle + s_1s_2s_3|\psi_{0,L}^\perp\rangle) = U_{00}|\psi_{0,L}\rangle \end{aligned}$$

with probability $p_{00} = c_1^2c_2^2c_3^2 + s_1^2s_2^2s_3^2$. Here, we use the unitary error with angle θ_{00} such that $\tan\theta_{00} = \tan\theta_1 \tan\theta_2 \tan\theta_3$. We observe that the resulting state still has a unitary error and that it resides in the code space.

- Even-odd parity (eigenvalues +1, -1): after applying the error recovery operation iY_3 we obtain the state

$$\begin{aligned} |\psi_0\rangle &= \frac{1}{\sqrt{p_{01}}}((ac_1c_2s_3 + bs_1s_2c_3)|000\rangle - (as_1s_2c_3 - bc_1c_2s_3)|111\rangle) \\ &= \frac{1}{\sqrt{p_{01}}}(c_1c_2s_3|\psi_{0,L}\rangle - s_1s_2c_3|\psi_{0,L}^\perp\rangle) = U_{01}|\psi_{0,L}\rangle \end{aligned}$$

with $p_{01} = c_1^2c_2^2s_3^2 + s_1^2s_2^2c_3^2$ and $\tan\theta_{01} = -\tan\theta_1 \tan\theta_2 \cot\theta_3$.

- Even-odd parity (eigenvalues -1, +1): we apply the error recovery operation iY_1 and obtain analogous results with $p_{10} = s_1^2c_2^2c_3^2 + c_1^2s_2^2s_3^2$ and $\tan\theta_{10} = -\cot\theta_1 \tan\theta_2 \tan\theta_3$.
- Odd-odd parity (eigenvalues -1, -1): we recover with iY_2 and obtain analogous results with $p_{11} = c_1^2s_2^2c_3^2 + s_1^2c_2^2s_3^2$ and $\tan\theta_{11} = -\tan\theta_1 \cot\theta_2 \tan\theta_3$.

In order to formally describe the result of this procedure, we write it as a quantum channel, i.e., a linear, completely positive and trace-preserving (CPTP) map on density matrices. Without error correction, this map is a single unitary error

$$\rho \rightarrow U_0 \rho U_0^\dagger,$$

and in general mixes states inside and outside the logical code space. In the encoded form, this map goes between logical states. It has four branches depending on the measurement outcome, i.e., it is in general not unitary

$$\rho_L \rightarrow \rho_{L,C} = \sum_{\sigma, \sigma'=0}^1 p_{\sigma, \sigma'} U_{\sigma, \sigma'} \rho_L U_{\sigma, \sigma'}^\dagger,$$

where we assumed that we disregard the measurement information. So we have turned a unitary error into a weighted sum of unitary errors.

Let us remark that we have not performed any approximations at this stage (but a number of assumptions, of course). In particular, we have not assumed anything about a and b other than the state is normalized. In particular, either of these coefficients may be very small. We have also not assumed anything about the error angles $\theta_{1,2,3}$ and we see that there is nothing special happening if they are identical (only expressions get more symmetric). Before we introduce assumptions, let us estimate the measurement error resulting from the last result. It reads

$$\rho \rightarrow (1-p)\rho + pY\rho Y,$$

Now let us assume that the error angles are small and similar in size and expand to lowest non-vanishing order. We note that $\theta_{00} \in O(\theta_i^3)$ and $p_{00} \in O(1)$ hence the contribution of that term is $O(\theta_i^6)$. For the other three terms we have $p \in O(\theta_i^2)$ and $\theta \in O(\theta_i)$ hence we have an error probability of $p_{EC} \approx 3\theta_i^4$. Without encoding we have θ_i^2 , hence we have reduced the error as long as $\theta_i^2 < 1/3$. This very high threshold is an artifact of the assumption that measurement and recovery can be made perfect. In lieu of a full threshold calculation as it is done in literature, we should at least assume that the longer piece of algorithm in error-corrected memory increases $\theta_{1,2,3}$ by a constant factor compared to θ_0 which lowers the threshold but keeps it in existence. If one wants to reduce the error rate even further, one would introduce a similar higher-level encoding, i.e., build qubits out of three of these logical qubits.

As a remark on the error channel notation: a fully random bit-flip channel would read

$$\rho \rightarrow (1-p)\rho + pY\rho Y,$$

whereas our unitary error channel can be expressed as

$$\rho \rightarrow \cos^2 \theta \rho + \sin^2 \theta Y \rho Y - \sin \theta \cos \theta [iY, \rho] .$$

So even though one can argue that the over-rotation channel is applied all the time, what counts as an error probability is only the portion that does not conserve the density matrix, i.e., the terms that do not include a commutator. We also remark that the Pauli-twirl of the unitary channel, with $\sigma_0 = 1, \sigma_1 = X, \sigma_2 = Y$ and $\sigma_3 = Z$,

$$\text{reads } \frac{1}{4} \sum_{j=0}^3 (\cos^2 \theta \rho + \sin^2 \theta \sigma_j Y \sigma_j \rho \sigma_j Y \sigma_j - \sin \theta \cos \theta [\sigma_j iY \sigma_j, \rho]) = \cos^2 \theta \rho + \sin^2 \theta Y \rho Y$$

i.e., mapping the unitary error channel onto the random bit-flip channel with probabilities $p = \sin^2 \theta$.

All in all, we see that in this simple yet generic example, a unitary error is corrected the same way as a classically random error with $\sin^2 \theta_i$ playing the role of error probability p_i . What is quite remarkable is that in the case of a syndrome being detected, the scaling of the error angle does not change—it is a combination of a reduced *classical* error probability and a unitary error of the original order of magnitude that leads to protection.

22.3 Mapping between small ubiquitous errors and error syndrome detection

Simplified representations of quantum error correction often describe the error process verbally, thus obfuscating their connection to physical error models. We present a discussion point out that relationship.

Continuous errors:

States of realistic quantum systems are usually mixed, i.e., next to the intrinsic randomness of quantum mechanics there is also a probability distribution over what state the system is in. These are described by a density matrix, a convex sum over said distribution

$$\hat{\rho} = \sum p_i |\psi\rangle\langle\psi| .$$

These can have non-vanishing entropy and can describe phenomena such as decoherence. Under generic assumptions, their time dynamics are described by a Master equation in Lindblad form (of with the Bloch equation is an example

$$i \dot{\hat{\rho}} = -i [\hat{H}, \hat{\rho}] + \sum_i D_i[\rho]$$

where we are using the shorthand notation for a specific linear operation

$$D_i[\hat{\rho}] = \hat{L}_i \hat{\rho} \bar{\hat{L}}_i - \frac{1}{2} \bar{\hat{L}}_i \hat{L}_i \hat{\rho} - \frac{1}{2} \hat{\rho} \bar{\hat{L}}_i \hat{L}_i$$

where the Lindblad operators describe different decoherence and relaxation mechanism. Pure dephasing, e.g., is described by

$$L_1 = \frac{1}{\sqrt{T_2}} \hat{\sigma}_z$$

with the usual dephasing time T_2 . The first term on the RHS of the Lindblad equation describes the unitary component of quantum dynamics thus being equivalent to the Schrödinger equation. In a general, externally controlled quantum computer, the Hamiltonian and the Lindblad operators, are time dependent. We can now ask the question how such dynamics responds to error correction. To this end, we need to be able to describe what happens between syndrome detection cycles, i.e., after starting the computation (or after the last syndrome detection) at time t_i we need to understand the density matrix at the next syndrome detection t_r .

This can be described by a structural theorem on the formal solution of the Lindblad equation, the operator sum representation (or Kraus decomposition), which follows from the semigroup structure generated by the Lindblad equation. It states

$$\hat{\rho}(t_f) = \sum_{i=1}^N \hat{K}_i \hat{\rho}(t_i) \hat{K}_i^\dagger$$

with

$$\sum \hat{K}_i \hat{K}_i^\dagger = \hat{1}$$

where the Kraus operators K_i encode the dynamics. This is a generalization of purely unitary evolution, where $N=1$ and $K_1=U$. This representation can be read as the system being exposed to a combination of different linear maps.

In order to describe syndrome detection (which for the time being is supposed to be perfect, but is generally treated as having its own error in the quantitative error correction literature), we need to understand how measurement can be described in this formalism. Suppose the measured observable O can be described by eigenvalues o_i and corresponding mutually orthogonal eigenspace projectors Π_i . Then, the probability of measuring o_i is and $P_i = \text{Tr}(\Pi_i \hat{\rho})$. The post-measurement density matrix

$$\rho_i = \frac{\Pi_i \hat{\rho} \Pi_i}{P_i}$$

Note that this is normalized, as this state is conditioned on the measurement outcome.

The connection from the Lindblad equation describing time-continuous errors vs the Kraus decomposition describing their impact at given time allows us to match physical error models to the qubit errors that error correction is geared up to.

Let's discuss this along an example. Suppose we are dealing with quantum memory ($H=0$) and pure dephasing (the L_i specified above is the only Lindblad operator). We can then easily show that the corresponding Kraus operators are

$$K_0 = \sqrt{1-p} \hat{1} \quad K_1 = \sqrt{p} \hat{Z}$$

and with

$$p = e^{-\Delta t/T_2}$$

where $p = (1-q)/2$ is the time for which the state is kept in memory.

Discrete errors:

Let us on the other hand study a quantum channel that performs a phase flip with probability p . For any given pure state of a qubit, we can write this as a map

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle \rightarrow \alpha |0\rangle - \beta |1\rangle$$

Lifted to projectors, this is equivalent to

$$|\psi\rangle\langle\psi| \rightarrow \hat{Z} |\psi\rangle\langle\psi| \hat{Z}$$

If we start with a density matrix as convex sum of projectors, an analogous expression holds. If we thus would like to describe the quantum map that corresponds to phase flip with probability p and no error with probability $1-p$ (a large but rare error), we can construct it as

$$\rho \rightarrow (1-p)\rho + p \hat{Z} \rho \hat{Z}$$

which is the exact same expression that we got above from representing the solution of the Lindblad equation with Kraus operators.

We can thus conclude, that at the time of measurement (syndrome measurement or final readout of the computer), continuous pure dephasing is indistinguishable from phase flip with the probability

$$p = e^{-\Delta t/T_2} .$$

Correcting the latter means correcting the former. We can interpret this expression for the probability as a stroboscopic view on the time evolution. Reducing p can thus, at these stroboscopic times, be interpreted as enhancing T_2 .

22.4 Error correction experiments

An impressive experiment on this topic is in Reference [KBF⁺15]. It addresses correction of single-qubit bit flip errors (the correction is done in software as there is only one round of error correction, so this is technically an error detection experiment). Plot 4a shows on its x-axis such stroboscopic time in the form of error correction cycles, including a more than 8-fold reduction of error rate that can be translated into extending T_1 . Operations improved by error detection with cloud access only has been demonstrated in [HF18].

During the writing and review of the second revision, a remarkable experimental implementation of the surface code has been posted by the Wallraff group at ETH [ARL⁺17]. In this experiment, a small plaquette of the code containing seven qubits was implemented, X and Z errors were simultaneously detected and the logical measurement result was corrected based on the outcome of the syndrome measurements. The results indicated that error correction is effective in reducing the output logical error probability as expected. However, simulations based on a simple model of the chip were not successful in quantitatively predicting the data - the experiment was less effective. At this point, this likely points to an incomplete model.

The hardware used are Transmon qubits in a planar layout, see section 13.2.

22.5 Systematic unitary errors

In this discussion we would like to highlight the case of uniform coherent errors. Linking back to section 6.5 these are systematic errors that are always the same - with the additional assumption of uniformity meaning that we at least permit that all qubits suffer from the same coherent error. This can be due to a global miscalibration. Remarkably, the theory of error correction also allows for these to be corrected, as described in section 7.4.1.3 of the literature review as well as in our demonstration in section 22.1 - how can that be? The key lies in the probabilistic character of quantum physics, specifically in translating matrix elements of unitary matrices, such as misrotation angles, into probabilities.

In fact, a small-angle rotation happening in each round of error correction (for example after initialization, see section 23.1) or at each qubit should also not be misinterpreted as an error rate of 100%. Only the probability of a wrong measurement outcome matters in the end, and this probability does not only depend on the frequency of errors but also on their rotation angle. Thus, even if all rotation angles are the same (and typically small), the results of the parity measurements underlying syndrome correction can show odd parity, in which case the measurement also projects the state onto an odd-parity subspace. Thus, even though error correction can not reduce this *physical* error (because it might appear again in the next round of error correction, or it might not be detected due to similar errors of syndrome qubits) it can still raise the probability for the correct *logical* state - as does error correction for purely stochastic errors. The main task of the actual error correction is to keep the *physical* errors constant in each round, i.e., to prevent errors from adding up during a larger calculation by including the syndrome measurements and resetting the physical qubits repeatedly. With that as a foundation, the logical encoding makes sure that the errors of the

logical qubits are polynomially smaller than the physical errors, i.e., that a measurement of the logical state still gives the right result with arbitrary high probability by adding enough physical qubits to the code.

22.6 Unrecoverable errors

What is indeed well known is that quantum error correction cannot detect multi-qubit correlated errors. Multi-qubit errors that commute with stabilizers cannot be detected by them. For example for the three-qubit repetition code is stabilized by Z_1Z_2 and Z_2Z_3 . If we have correlated errors of the form Y_1Y_2 it will commute with the first stabilizer (but not with the second) so it is partially detected, errors of the form $Y_1Y_2Y_3$ commute with both and cannot be detected—if we flip all three qubits at the same time, we will not notice but rather think that this is a logical operation. The only remedy against this errors is to increase the code distance over the maximum correlation length. In the surface code, errors that connect distinct boundaries of the surface are of this kind and increasing the surface makes them less likely. Clearly, a correlated error across the whole processor will be fatal at all times. Correlated errors will show up in RB, so they are not hidden—they are simply hard to recover from.

Note that having a fully correlated error of this type means that seeing an error on both qubits is not less likely than seeing an error on one qubits. Real error models will contain both.

It is thus important to identify sources of correlated errors in physical platforms, also discerning them from uncorrelated but potentially simultaneous errors. We focus on extreme cases of purely unitary and purely random correlated errors.

22.6.1 Correlated coherent errors

Here and henceforth we call the n qubit correlated error operator $E = S_{i_1} \dots S_{i_n}$ with the S_i being single qubit operators in n distinct qubits $i_1 \dots i_n$. We normalize them such that $S_i^2 = 1$, and $S_i = S_i^\dagger$, i.e., assume they are Pauli-type matrices. A unitary error would be

$$U = \exp(-i\phi E) = \cos(\phi) - i\sin(\phi)E,$$

generated by applying an error Hamiltonian $H_{\text{err}} = gE$ for a time $t = \phi/g$. They hence are based on a time evolution equivalent to an n qubit interaction that needs to occur in an undesired way. Physical interactions are two-body at most, hence direct mechanisms for this scenario are limited to $n = 2$. There is a large effort in the field of perturbative gadgets [KKR06, JF08, BLAG14] that aims at constructing interactions with $n > 3$ mediated by high-energy degrees of freedom. A key result of this work is that the effective interaction constant for such interactions scales as $g_{\text{eff}} \propto (g_2/\Omega)^n$ where g_2 is the underlying physical interaction and Ω is the large energy scale of the mediating degree of freedom. This means that higher-order interactions are weak, as testified by the fact that even elaborate experimentation has not produced any multiqubit gadget today. The same effect also makes n -qubit errors highly unlikely.

It needs to be noted that errors from an external, spatially correlated, random field coupling to many individual qubits does usually *not* lead to this type of error. The coupling of that field goes to individual qubits, so there we have $H_{\text{err}} = g \sum S_{ij} \otimes F$ where F describes the field, inducing uncorrelated single qubit errors.

It needs to be noted that of course local errors can propagate through the processor, but this is taken into account in numerical threshold calculations.

22.6.2 Correlated stochastic errors

Errors due to macroscopic unobserved environments that are weakly coupled (as they are in a well-engineered qubit candidate that passes lower-level functionality tests, level A) can be described by Fermi's golden rule using the system-bath coupling H_{SB} as a perturbation. They describe transitions between states $|i\rangle$ and $|f\rangle$ with a rate proportional to the square matrix element $|\langle f|H_{\text{SB}}|i\rangle|^2$ where $|i\rangle$ and $|f\rangle$ are the states

connected by the random, incoherent process. For an n qubit error, these states need to be different in n qubits, hence H_{SB} needs to contain non-unit operators in n qubits. For the same reason as for unitary errors, this is not known in nature for $n > 2$. Higher order errors can thus occur in higher order perturbation theory in higher order perturbation theory, i.e., if H_{SB} contains $k \leq 2$ qubit operators the rate for n -qubit errors will scale as $(H_{SB}/E_q)^{2nk}$ where $E_q \gg H_{SB}$ is a typical qubit energy.

22.7 Decoherence and entanglement with an environment

Decoherence is a primary error channel in quantum computers. It is built up when the quantum system of interest interacts with unobserved environmental degrees of freedom. Here, we show that this entanglement leads to errors that correspond to the standard formulation of quantum error correction, Kraus channels describing stochastic errors, along a toy example.

Consider a qubit with an environmental degree of freedom that are initially not entangled, starting in the state

$$(a|0\rangle + b|1\rangle) \otimes |i\rangle$$

they undergo state-dependent scattering ending up in

$$(a|0\rangle \otimes |a0\rangle + b|1\rangle \otimes |a1\rangle)$$

Now as outlined, we can only operate on the qubit, i.e., all observable we measure will have a structure of

$$O \equiv O_q \otimes 1_E$$

so to describe any outcomes of measurements of such an observable, we need to compute expectation values using the complete density matrix of the form

$$\langle O \rangle = \text{Tr}_{q+E}(O_q \otimes 1_E \rho) = \text{Tr}_q(O_q \otimes \rho_R)$$

where we have defined the reduced density matrix as the trace of the full density matrix over the environment

$$\rho_R = \text{Tr}_E(\rho)$$

For the state undergoing state-dependent scattering outlined above, we can easily write down the reduced density matrix

$$\rho_R = \begin{pmatrix} |a|^2 & a\bar{b}\bar{q} \\ \bar{a}bq & |b|^2 \end{pmatrix}$$

with

$$q = \langle A_1 | A_0 \rangle$$

which we assume to be real-valued from now on. The reduced density matrix is related to the unperturbed version of the density matrix (without scattering)

$$\rho_P = \begin{pmatrix} |a|^2 & a\bar{b} \\ \bar{a}b & |b|^2 \end{pmatrix}$$

by the pure phase error map

$$\rho_R = p\rho_P + (1-p)Z\rho_P Z$$

with error probability of

$$p = (1-q)/2$$

So this is exactly the type of error channel that quantum error correction is designed to correct: There is a bit flip error with probability p . In order to meet the error correction threshold, an experimenter must keep this low, i.e., she must make sure that scattering events are rare and/or that the states of the environment corresponding to the two qubit states have large overlap q , close to unity, i.e., isolate the qubit states. Also, as can be seen in section 7.4.1.3, this same error would be induced by a time-continuous, pure T_2 -process.

23 Fault-tolerance calculation

During the realization of an error correction cycle, errors can happen at several different points: Due to a faulty initialization, a faulty gate, waiting time (i.e., a faulty identity gate) or measurement. As we will show here, all of these errors can be modeled by one (effective) error operator before measurement. We will describe the model for arbitrary unitary errors occurring at all steps of the cycle. The special case of rotation angle π can describe Pauli operators, a rotation angle of 0 corresponds to an identity operator, i.e., no error happening.

An initialization error occurs whenever a qubit is not initialized in the right quantum state. Instead of the state $|\psi\rangle$, the qubit is now in the state $|\psi_C\rangle = U_I|\psi\rangle$. A faulty gate (including identity and multi-qubit gates) U_C instead of U can always be written as the ideal gate followed by an error operator U_G such that $U_C = U_G U$. A wrong measurement basis can be described by a rotation U_M of the state before measurement. Now, if the structure of the error correction algorithm is known, one can calculate the maximal overall error that can occur during a syndrome measurement cycle (and a fault-tolerant algorithm only consists of such cycles). Let us consider a circuit with several gates U_i and some waiting time before measurement. Waiting times can occur whenever other, longer operations are performed in another part of the code in parallel. The total effective gate sequence between qubit preparation and measurement will be $U_M U_W \prod_i (U_{G_i} U_i) U_I$ where U_W is the error due to waiting and U_{G_i} is the error assigned to the gate U_i . If all gates are Clifford gates (which usually is the case—even T gates are realized with Clifford gates and magic states) then we can permute the error operators to the left of the gate sequence [AL07] and calculate the effective error operator such that the sequence now can be separated in the ideal part and the total error: $U_{\text{eff}} \prod_i U_i$. When interpreting the error operators as rotations in the Bloch sphere, then the maximal rotation angle of the effective error U_{eff} is the sum of all single rotation angles, since permuting through Clifford gates does not change the rotation angle but only the rotation axis. Now that we have our error right before measurement, we can simply interpret any error by the probability for effecting a (Pauli) flip in the measurement basis, i.e., to a wrong measurement outcome and thereby a projection to an orthogonal state.

For larger systems, all of the operators can be understood as (usually uncorrelated) multi-qubit operators so that in the end, the error probability can be calculated for each qubit individually.

23.1 Errors from deviations in the measurement and initialization basis

In this section, we will study inaccuracies of the measurement basis, when measurement is also used for initialization. This is a special case of a faulty implementation that has slightly different (but not worse) implications to the error-correction mechanism.

23.1.1 Problem settings

We assume that we have an N -qubit processor outfitted with detection in a single basis. This detection mechanism is designed to measure in a specific physical basis $B = \{|0\rangle, |1\rangle\}$ which can represent quantities like voltage for electrical systems, some spin component in spin systems etc. We assume that this adjustment is not quite accurate, so in reality, qubit measurements are taken in a basis $B_n = \{|0\rangle_n, |1\rangle_n\}$ that is slightly deviating. We call U_n the unitary operation (basis-to-basis map) between those, i.e., $|\sigma\rangle_n = U_n|\sigma\rangle$. This impacts measurement and initialization which in the long run should be done by measurement.

It is peculiar that this error does not occur in what is usually the limiting factor—two-qubit gate fidelity—but places it into measurement and initialization, which are typically less emphasized in expositions of fault tolerance—however, they are being studied in the finer and more detailed points of the state-of-the-art work on the subject. In this topic, Nielsen/Chuang [NC00] shows its age. Also, the key step from error correction to fault tolerance (the former being a part of the latter), it is explored how all steps involve, including measurement. What is more peculiar is that this error can occur as a physical inhomogeneity—in each qubit

it is static, it occurs the same way every time, very much like T_2^* effects. If we were to average over qubits, it could still be random, it could at worst be systematic. We can look at both limiting cases at least.

23.1.2 Formalization

Here is a formalization of this error and how it plays into threshold calculations. Let P be the superoperator of a quantum measurement (that still can be inefficient in accordance with the threshold theorem) in basis B . Then, the physical measurement reads

$$P_n = P U_n^\dagger,$$

where U_n is the superoperator associated with U_n so we measure in a qubit specific basis. This would be the initially most straightforward and physical way to think about this phenomenon.

An alternative view (which is more easily connected to standard error correction) is to turn the table: error correction largely relies on measurement, and measurement and preparation are the prominent ways in which classical data is converted into quantum states. So we define B_n as the basis we would like to encode information in—even if we cannot write it down precisely, it is the one we read and write in. This change of paradigm means, that any gate acquires an error—a single qubit gate U_{1q} becomes $U_{1q,n} = U_n^\dagger U_{1q} U_n$, a two qubit gates between qubits n and m becomes $U_{2q,nm} = U_n^\dagger \otimes U_m^\dagger U_{2q} U_n \otimes U_m$ etc. Thus, if we subscribe to this redefinition, we obtain inhomogeneity-type gate errors.

23.1.3 A tighter bound

Assigning the same error to each gate seems to be overly pessimistic. Let us for the moment assume that all gates are performed without error and represent everything in this ideal basis. Then, an error due to a faulty measurement only occurs before the gates are applied (i.e., after initialization) and directly before measurement. Having a unitary mapping U between the faulty measurement basis and our chosen ideal basis, and ideal gates U_1, U_2, \dots, U_N applied between initialization and measurement, the total channel of operations acting on the qubit after initialization, expressed in the ideal basis is $U^\dagger U_N \dots U_2 U_1 U$. We can now also permute U to the left and write the gate sequence as $U_{\text{eff}} U_N \dots U_2 U_1$, with U_{eff} being an error of the same order of magnitude as U (maximally having doubled the rotation angle). If U commutes with all gates, it even cancels out to no effective error happening.

If we describe the gates as rotated from the ideal basis defined by the actual measurement basis, i.e., taking the alternative view, we get the same result: each gate can be written in the measurement basis as $U^\dagger U_i U$; performing several such gates in a row results in $(U^\dagger U_N U) \dots (U^\dagger U_2 U) (U^\dagger U_1 U) = U^\dagger U_N \dots U_2 U_1 U$.

This scenario shows that these errors should not be counted per gate but per measurement cycle, in between which there can be an arbitrary number of gates, e.g., for parity extraction, hence assigning an error to every gate is overly pessimistic. Only in measurement cycles involving several multi-qubit gates the effects can get bigger, depending on how many qubits are connected by the gates and how the bases of the qubits vary between each other.

23.1.4 Randomized Benchmarking

Are these errors easily detectable? As we are discussing these on a level of qubits that are functioning on a basic level, it is most appropriate to see how RB responds to them. This is rather clear—the basis changes for the gate cancel out during the RB sequence, so this does not contribute to the error per gate but it does contribute to state preparation and measurement errors. This is appropriate as clearly also in a quantum processor that executes multiple gates, these errors are independent of the number of gates during a measurement cycle.

24 Low-level error mitigation

24.1 Dynamical decoupling

Quantum error correction algorithms are not the only way to make algorithms run more stable. Instead of (or in addition to) creating logical qubits, one can also reduce unitary errors directly on the physical level by eliminating unwanted interactions.

One way to do so is dynamical decoupling, which seeks to minimize the unwanted system-bath interactions in an open quantum system. The system-bath Hamiltonian reads $H_0 = H_S \otimes 1 + 1 \otimes H_B + H_{SB}$, where H_{SB} is the system-bath interaction. Now we add a time-dependent perturbation to the Hamiltonian such that the time evolution of $H_1(t)$ is periodic with T_c , i. e., $U_1(t + T_c) = U_1(t)$ [VKL99]. In a frame rotating with H_1 , the Hamiltonian reads

$$H(t) = H_0 + H_1(t) ,$$

and the total time evolution reads $U_{\text{tot}}(T_N) = \exp(-iH'T_N)$, with $T_N = NT_c$ and the effective Hamiltonian H' is determined through the Magnus expansion [BCOR09] in $U_{\text{tot}}(T_c) = \exp(-iH'T_c)$. By proper choice of $H_1(t)$ this decouples the time evolution of system and bath up to the desired order in the Magnus expansion.

Explicitly, the first order of the Magnus expansion reads

$$H'^{(1)} = \frac{1}{T_c} \int_0^{T_c} H^R(t) dt ,$$

and the effective Hamiltonian reads $H' = \sum_{i=1} H'^{(i)}$. This is a technique that originally has been developed and refined in NMR pre quantum-computing. It removes systematic and unitary errors of physical qubits on the expense of more complicated and thus longer operations. Its success can be verified by Level B benchmarking techniques and provides input to the error rates used in fault-tolerance. It is generally effective for systems with low non-unitary errors and large inhomogenities, for example in electron spins perturbed by nuclear spins as well as in ion traps. Its specific applications are given along with the physical platforms.

Although in realistic environments, dynamical decoupling can never completely avoid all unitary errors, it can still be used to further reduce physical error rates in order to have a surface code run more efficiently on those qubits. An optimal balance between the costs of this physical error reduction and error correction on the logical level is analyzed in [PSL13].

24.2 Decoherence-free subspaces

Decoherence-free subspaces [LW03, BKLW00, KBLW01] exploit symmetries of noise mechanisms in order to render noise ineffective. They are common eigenspaces of noise operators and system dynamics. In order to be viable, they need to allow universal computing within that subspace. DFS have been demonstrated [Vio01, Kie01, KMW02, YHO⁺08, RKB⁺16] in artificial scenarios. One can interpret DFS as error correction without syndrome detection.

The challenge is that an argument like the definition of DFS makes rather precise requirements about the noise operators, i.e., the noise needs to be known precisely and ideally there must be a robust symmetry that guarantees a certain structure of the noise operators. So, realistically, if these properties are approximately satisfied, DFS constructions may be used to lower the impact of decoherence but not remove it completely as an engineered machine typically does not obey strong symmetries. Even in this situation, DFS are very specific to one of the noise mechanisms affecting a quantum computer. Comparing DFS against error correction one can notice that DFS offer protection against collective noise if given the promise of a highly symmetric noise structure whereas error correction exploits the uncorrelated nature of noise mechanisms—the former assumption is rarely satisfied, the latter is rather common.

DFS have contributed to isolated solutions of decoherence problems, e.g., in the triple-quantum-dot spin qubit described in 15.1.2. They can potentially contribute to further solutions of this kind, but on a large scale where symmetries are less likely to occur.

25 Practicalities of fault tolerance

25.1 Fault-tolerant implementation of a small rotation: A practical example

As described in Section 7.2.4, a key challenge in fault-tolerant computation is to implement fault tolerant logic gates. For the surface code, this can be done within the code for all Clifford gates. In section 7.2.4.3, the implementation of necessary non-Clifford gates has significant computational overhead, suggesting the use of a minimal gate set consisting of the Clifford (Section 7.2.4) and T gates (Section 7.2.4.3). The reason for the restriction on the T (and T^{-1}) gate is that there it is known how to implement it with an ancilla factory (Section 7.2.4.3), and that only one type of factory is needed. It is possible based on the Solovay-Kitaev Theorem to approximate any non-Clifford gate with Clifford plus T with only logarithmic overhead in time [Sol00, Kit97b]. We would like to exemplify this along the example of a small rotation based on [RS16a] which says that a rotation

can be approximated by the gate sequence

```
HTSHTSHTHTHTHTSHTHTSHTSHTSHTHTHTSHTSHTH
THTSHTHTSHTHTHTHTHTHTSHTSHTSHTHTSHTHTSHT
HTHTHTSHTHTHTSHTHTSHTHTHTHTSHTSHTSHTHTHTSH
TSHTSHTSHTHTSHTSHTSHTSHTHTSHTHTSHTSHTHTHTH
THTSHTHTHTSHTSHTSHTHTSHTSHTHTHTSHTHTHTHT
HTSHTSHTHTHTHTSHTHTHTHTSHTHTHTHTHTHTHTH $\omega^7$ 
```

with a deviation below 10^{-10} , ω here is a global phase. How do we implement it with the surface code? Before performing the cycle, we need to determine the degree of error protection we want. For being able to correct for at least arbitrary single errors on a single-qubit, we need at least a distance 3 code. For state-of-the-art qubits their error rate requires distance 100 code and beyond to also correct multiple errors. In order to initialize a single qubit of the double-cut variety, we need to expand this to 72 physical qubits.

Every single element of the gate sequence above now has to be operated as a logical protected gate on a logical qubit. The following section seeks to introduce a schematic way of representing the structure and sequence of stabilizer measurements that needs to be performed in order to get those logical operations.

25.2 Stabilizer measurement pattern for logical operations

We will present the operation of logical gates in a distance-independent way, only showing in which regions of the physical qubit array stabilizer measurements are deactivated during each step. Figure 25.1 explains the new diagram in comparison with the previously used representation. Our new diagram also serves as step between the representation used in [FMMC12] (and chapter 7) and the three-dimensional space-time building blocks used for example in [Fow12, FD12, FDJ13], as described in section 7.2.5. One picture of our diagram corresponds to one constant-time layer of the three-dimensional representation.




Figure 25.1: Translation of our schematic diagram (left of the “=”) to stabilizer measurements on physical qubits as represented in [FMMC12]. One picture can correspond to one or more surface code rounds. Gray regions show where stabilizer measurements on the qubit array are turned off during these rounds. The edge color represents the type of active stabilizers next to the deactivated region. This is a distance 3 example, however, the schematic diagram is valid independent of distance, where gray regions can also include more than one deactivated stabilizer, i.e., represent larger holes (proportions unaccounted for). Inside the holes, additional direct action on data qubits may be necessary, which is also not shown.

Every logical operation can be performed by just running a sequence of stabilizer measurements with some deactivated regions. The sequence of required stabilizer measurements for logical initialization, readout, Hadamard and CNOT implementations are shown in Figure 25.2–Figure 25.5. Every sequence might require several runs of each step (to track temporal error chains), depending on the actual distance to be preserved.




Figure 25.2: Sequence of different stabilizer measurement patterns for initialization of a Z-cut qubit in a Z eigenstate (a) or X eigenstate (b). Time runs from left to right.




Figure 25.3: Sequence of different stabilizer measurement patterns for measurement of a Z-cut qubit in the Z basis (a) or X basis (b).



Figure 25.4: Sequence of different stabilizer measurement patterns for a logical Hadamard. Uncolored pictures represent additional physical gates between the stabilizer measurements (number of qubits chosen arbitrary). Here, the physical Hadamards in step 4 can be done during the measurement of the stabilizer cycle before, the physical SWAP gates (step 5–6) can be performed while stabilizer measurements are running in the rest of the array. This sequence basically shows a transformation to a different qubit encoding (step 3), where the logical X and Z operators are just vertical and horizontal lines, on which a Hadamard can be performed easily (steps 4–6) and a transformation back to the original shape of the qubit.




Figure 25.5: Sequence of different stabilizer measurement patterns for a logical CNOT, in analogy to the braiding diagram in Figure 7.6 (d). Logical qubit operators are defined as horizontal operator chains between the holes, i. e., the two top holes build a logical qubit, and the two bottom ones build another qubit. Steps 2–3 show the creation of an X -cut qubit (green) in a Z eigenstate, step 4 is a projective measurement (with initialization following in step 5) of the upper Z -cut qubit in the X -basis, step 6 a measurement (destructive) of the X -cut qubit in the Z basis.

Logical S and T gates are performed indirectly using logical Hadamard and CNOT gates with special ancilla states (see Figure 7.7), that are prepared separately by a logical state distillation algorithm (as described in section 7.2.4.3), and a logical measurement. Due to its non-deterministic character from the measurement, every T gate additionally needs one step of classical information processing and feed-forward to decide whether or not to apply a corrective gate for a successful execution.

So, any gate sequence to be performed fault-tolerant translates to a sequence of stabilizer measurement patterns, consisting of the sequences for logical initialization, measurement, Hadamard and CNOT, and an additional ancilla factory running in parallel, producing states to be consumed for S and T gates.

Error syndromes are collected throughout the whole process in every single cycle to keep track of all physical errors happening and correcting their effect in the very end of the calculation (section 7.2.2). Basically, whenever a syndrome measurement detects an eigenvalue change of a physical qubit, we just redefine the logical qubit state by exchanging the two states of the flipped physical qubit—this means we correct how we interpret the final measurement. An error on the logical state only occurs when more than $d/2$ qubits have flipped after *one* syndrome measurement cycle. So, even though a large number of errors might occur during the whole gate sequence, they can still be detected (and therefore also corrected in the end) as long as they are spread enough over time so that less than $d/2$ error happen in a row at the same time or on the same qubit (see section 7.2.2.3). In the end, both error rates of the algorithmic approximation and of the surface code logical errors combine to one higher effective error rate that gives the deviation from the desired ideal gate.

26 List of records

In the following tables we give the best known values (summer 2017) for the maximal number of qubits controlled in an experiment n , the times for single-qubit gates $t_{g,1}$ and—if existing—two-qubit gates $t_{g,2}$, coherence times T_1 , T_2 , T_2^* , $T_{1\rho}$, T_{2E} (Hahn-Echo), and fidelities F_1 and F_2 (for single- and two-qubit states), see Chapter 6 for further description of stated fidelities, gate types, and other specifications. Note that in many cases we have given multiple records for the same platform—that is in cases when some numbers look impressive but some other crucial milestones (such as two-qubit gate operation) were not met. This is important as these can be design trade-offs—one can produce very clean qubits by removing most of their connectivity.

Type of qubit	Group	n	Gate times	Coherence time	Fidelity	
Transmon (planar)	IBM [SMCG16, RGP ⁺ 12, SBM ⁺ 16,McKWS+17]	5	$t_{g,1} = 13\text{ns}$ $t_{g,2} = 160\text{ns}$	$T_1 = 70\mu\text{s}$ $T_2^* = 92\mu\text{s}$	$F_1 = 0.9998$ $F_2 = 0.991$	R B
	Google [BKM ⁺ 14]	9	$t_{g,2} = 40\text{ns}$		$F_1 = 0.9992$ $F_2 = 0.994$	R B
	Google [BKM ⁺ 13]	1	$t_{g,1} = 10\text{ns}$	$T_1 = 44\mu\text{s}$ $T_2 = 20\mu\text{s}$ $T_2^* = 15\mu\text{s}$		
Transmon (3D)	IBM [PMS ⁺ 16]	4	$t_{g,2} \sim 300\text{--}700\text{ns}$	$T_1 \sim 90\mu\text{s}$ $T_2^* \sim 48\mu\text{s}$ $T_{2E} = 46\text{--}86\mu\text{s}$	$F_1 > 0.999$ $F_2 = 0.93\text{--}0.97$	R B
	Lincoln Labs [JKS ⁺ 15]			$T_1 = 80\mu\text{s}$ $T_2^* = 115\mu\text{s}$ $T_{2E} = 154\mu\text{s}$		
Flux qubit	Lincoln Labs [YGK ⁺ 16]			$T_1 = 40\mu\text{s}$ $T_2 = 85\mu\text{s}$ $T_{2E} = 40\mu\text{s}$		
	IQC [DOS ⁺ 15]		$t_{g,1} \sim 0.5\text{--}1\text{ns}$		$F_1 \sim 0.996\text{--}0.999$	sf
	Mooij [PdGHM07]		CNOT		$F_2 = 0.4$	

Table 26.1: Records for Josephson qubits.

Type of qubit	Group	n	Gate times	Coherence time	Fidelity	
⁹ Be ⁺ microwave controlled surface trap	Wineland [BWC ⁺ 11]			$T_2 = 0.38\text{s}$	$F_1 = 0.99998$	RB
⁹ Be ⁺	Wineland [GTL ⁺ 16]		$t_{g,2} = 30\mu\text{s}$	$T_R = 1.5\text{s}$	$F_1 = 0.999962$ $F_2 = 0.9992$	RB, Bell
⁴³ Ca ⁺ in room-temp. surface trap	Lucas [HSA ⁺ 16]		$t_{g,2} = 3.25\text{ms}$	$T_2^* = 50\text{s}$	$F_2 = 0.997$	MS
⁴³ Ca ⁺ room-temp. trap	Lucas [BHL ⁺ 16]		$t_{g,1} = 7.5\mu\text{s}$ $t_{g,2} = 100\mu\text{s}$		$F_1 = 0.999934$ $F_2 = 0.999$	
Surface trap proposal	Hensinger [LWF ⁺ 17]		$t_{g,1} = 2.5\mu\text{s}$ $t_{g,2} = 10\mu\text{s}$			a
¹⁷¹ Yb ⁺ linear Paul trap	Hensinger [RWL ⁺ 18]		$t_{g,1} = 100\text{ns}$		$F_1 = 0.9999$	RB
¹⁷¹ Yb ⁺ linear Paul trap with 5 qubits	Brown [LLF ⁺ 18]		$t_{g,2} = 90\mu\text{s}$		$F_2 = 0.983$	
¹⁷¹ Yb ⁺ linear Paul trap, ultra-fast	Monroe [WMJM17]		$t_{g,2} = 200\text{ps}$		$F_2 = 0.76$	

laser					
-------	--	--	--	--	--

Table 26.2: Records for trapped ions.

Type of qubit	Group	n	Gate times	Coherence time	Fidelity	
e ⁻ in Si	Morello [MDL ⁺ 14]			$T_R = 160\mu s$ $T_{2E} = 1.1ms$ $T_{CPMG} = 560ms$	$F_1 = 0.99953$	cf
³¹ P ⁺ in Si	Morello [MDL ⁺ 14]			$T_R = 600ms$ $T_{2E} = 1.8s$ $T_{CPMG} = 35.6s$		
Double quantum dot	Morello [VYH ⁺ 15]		$t_{g,2} = 100ns$	$T_2^* = 120\mu s$ $T_{2,CZ}^* = 8.3\mu s$		
Si/SiGe quantum dot	Vandersypen [WPK ⁺ 17, XWH ⁺ 18]		$t_{g,1} = 100ns$ $t_{g,2} = 500ns$	$T_1 > 50ms$ $T_{2,DD} = 19\mu s$	$F_1 = 0.99$ $F_2 = 0.92$	Bell RB
	Dzurak [HYC ⁺ 18]			$T_{2,DD} = 30\mu s$	$F_1 = 0.993$ $F_2 = 0.98$	RB
	Tarucha [YTO ⁺ 18]		$t_{g,1} = 120ns$	$T_{2,DD} = 99\mu s$	$F_1 = 0.999$	
GaAs	Kuemmeth [MMN ⁺ 17a]			$T_{2,DD} = 0.87ms$		
NV center	Hanson [RCB ⁺ 11]				$F_1 = 0.93$	avg
	Wrachtrup [JGP ⁺ 04]		$t_{g,2} \sim 100ns$	$T_2 = 6\mu s$	$F_2 \sim 0.92$	avg
	Awschalom [BFBA10]			$T_2 = 480\mu s$	$F_1 = 0.89$	
	[HSM06]			$T_1 = 380s$		
	Suter [ZSBS14]		$t_{g,1} = 35.5\mu s$		$F_1 = 0.985$	
SiV	Wrachtrup/ Vuckovic [RWN ⁺ 17]	100		$T_2 = 160\mu s$		

Table 26.3: Records for semiconductor platforms.

Type of qubit	Group	n	Gate times	Coherence time	Fidelity	
¹³³ Cs	[WKWW16]		$t_{g,1} = 80\mu s$	$T_2' = 7s$	$F_1 = 0.996$	RB, mw
⁸⁷ Rb	[KLFF ⁺ 15]			$T_2 = 1ms$ a T_2 was measured only for a single qubit	$F_2 = 0.69$	
Cavity	[WHD ⁺ 18]		$T_{g,2} = 2\mu s$		$F_2 = 0.76$	Bell
¹³³ Cs	[JHK ⁺ 16]		$t_{g,2} = 2\mu s$	$T_R = 40\mu s$	$F_2 = 0.81$	Bell

Table 26.4: Records for neutral atoms.

Type of qubit	Group	n	Gate times	Coherence time	Fidelity	
Single photon	Everywhere		$t_{g,1}$: short enough	long enough		
	O'Brien [OPW ⁺ 03]				$F_2 = 0.87$	ps, sf
Single photon (KLM)	White [MBB ⁺ 16]		$t_{g,2} > 300s$		$F_2 = 0.68$ (0.93)	ps, pf (hs)
Single photon (on chip)	O'Brien [HBR ⁺ 16]		$t_{g,2} \sim 20ms$	200ns	$F_2 > 0.82$	ps, sf
Coherent state/cat state	Andersen [TDL ⁺ 11]				$F_1 = 0.78$ (0.94)	ps, avg (sf)
	Schoelkopf		$t_{g,1} = 1\mu s$	$T_1 = 2.7ms$	$F_1 = 0.99$	RB

	[HRO ⁺ 16]					
Single photon cluster state	Several groups	6				ps
CV cluster state	[YYK ⁺ 16]	10 ⁶	$t_{g,2} = 160\text{ns}$			

Table 26.5: Records for photonic qubits.

Type of qubit	Group	n	Gate times	Coherence time	Fidelity	
NMR (liquid-state)	IQC [NMR ⁺ 06]	12		$T_1, T_2 : O(100\text{ms})$		
	IQC [RLL09]	1	$t_{g,1} = 24\mu\text{s}$ (100 μs)	$T_1 = 7\text{s}$ $T_2 = 4.5\text{s}$ $T_2^* = 0.45\text{s}$	$F_1 = 0.9999$ (0.99999)	RB (G)
	IQC [RLL09]	3	$t_{g,1} = 1.2\text{ms}$ $t_{g,2} = 2\text{--}4\text{ms}$	$T_1 = 2\text{--}5\text{s}$ $T_2 \sim 2\text{s}$	$F_2 = 0.995$	RB, G
NMR (solid-state)	IQC [BMR ⁺ 06]	3	$t_{g,3} \sim 450\mu\text{s}$	$T_1 = O(100\text{s})$ $T_2^* = O(1\text{ms})$	$F > 0.9$	
	IQC [PFR ⁺ 15]	5	$t_g = O(100\text{ns})$		$F_2 > 0.99$	avg

Table 26.6: Records for molecular approaches.

26.1 Abbreviations

These abbreviations also concern the different ways to characterize fidelity, which are analyzed in-depth in Chapter 5.

- ps** postselected, heralded
- G** GRAPE pulses used
- mw** microwave pulse
- MS** Mølmer-Sørensen gate
- a** assumed values
- sf** state fidelity
- cf** control fidelity
- pf** process fidelity
- hs** normalized Hilbert-Schmidt fidelity
- avg** average gate fidelity
- RB** average gate fidelity deduced from RB

Reference Documentation

- [AA13] S. Aaronson and A. Arkhipov. The computational complexity of linear optics. *Theory of Computing*, 9(4):143–252, 2013.
- [AAM⁺19] F. Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574, 505–510, 2019.
- [AB16a] S. Aaronson and D. J. Brod. Bosonsampling with lost photons. *Phys. Rev. A*, 93(1):012335, 2016.
- [AB16b] A. Auer and G. Burkard. Long-range photon-mediated gate scheme between nuclear spin qubits in diamond. *Phys. Rev. B*, 93:035402, Jan 2016.
- [ABB17] J. M. Auger, S. Bergamini, and D. E. Browne. A blueprint for fault-tolerant quantum computation with Rydberg atoms, *Phys. Rev. A* 96, 052320, 2017.
- [ABE58] E. R. Andrew, A. Bradbury, and R. G. Eades. Nuclear magnetic resonance spectra from a crystal rotated at high speed. *Nature*, 182(4650):1659–1659, 1958.
- [ABO99] D. Aharonov and M. Ben-Or. Fault-tolerant quantum computation with constant error rate, 1999, arXiv:quant-ph/9906129.
- [Abr61] A. Abragam. *The Principles of nuclear Magnetism*. Oxford University Press, London, 1961.
- [ADL⁺16] S. Asaad, C. Dickel, N. K. Langford, S. Poletto, A. Bruno, M. A. Rol, D. Deurloo, and L. DiCarlo. Independent, extensible control of same-frequency superconducting qubits by selective broadcasting. *Nat. Partn. J. Quantum Inf.*, 2:16029, 2016.
- [AGP08] P. Aliferis, D. Gottesman, and J. Preskill. Accuracy threshold for postselected quantum computation. *Quantum Inf. Comput.*, 8:181–244, 2008.
- [AGL+18] B. Amento-Adelmann, M. Grassl, B. Langenberg, Y.-K. Liu, E. Schoute, and R Steinwandt. *Quantum Cryptanalysis of Block Ciphers: A Case Study*. Poster at Quantum Information Processing QIP 2018, 2018.
- [AHM⁺16a] D. Aasen, M. Hell, R. V. Mishmash, A. Higginbotham, J. Danon, M. Leijnse, T. S. Jespersen, J. A. Folk, C. M. Marcus, K. Flensberg, and J. Alicea. Milestones toward majorana-based quantum computing. *Phys. Rev. X*, 6(3):031016, 2016.
- [AHM⁺16b] S. M. Albrecht, A. P. Higginbotham, M. Madsen, F. Kueemmeth, T. S. Jespersen, J. Nygård, P. Krogstrup, and C. M. Marcus. Exponential protection of zero modes in majorana islands. *Nature*, 531(7593):206–209, 2016.
- [AKR10] B. Altshuler, H. Krovi, and J. Roland. Anderson localization makes adiabatic quantum optimization fail. *Proc. Natl. Acad. Sci. U.S.A.*, 107(28):12446–12450, 2010.
- [AL06] P. Aliferis and D. W. Leung. Simple proof of fault tolerance in the graph-state model. *Phys. Rev. A*, 73(3):032308, 2006.
- [AL07] R. Alicki and K. Lendi. *Quantum dynamical semigroups and applications*. Lectures Notes in Physics. Springer, Berlin, 2007.
- [AL16] T. Albash and D. A. Lidar. *Adiabatic Quantum Computing*, 2016.
- [AL17] T. Albash and D. A. Lidar. Evidence for a limited quantum speedup on a quantum annealer, 2017, arXiv:1705.07452.
- [AL18] T. Albash and D. A. Lidar. Demonstration of a Scaling Advantage for a Quantum Annealer over Simulated Annealing, *Phys. Rev. X* 8, 031016, 2018.

- [ALB⁺07] M. Anderlini, P. J. Lee, B. L. Brown, J. Sebby-Strabley, W. D. Phillips, and J. V. Porto. Controlled exchange interaction between pairs of neutral atoms in an optical lattice. *Nature*, 448(7152):452–456, 2007.
- [Amb10] A. Ambainis. *Variable time amplitude amplification and a faster quantum algorithm for solving systems of linear equations*. arXiv:1010.4458, 2010.
- [AMG+16a] M. Amy, O. D. Matteo, V. Gheorghiu, M. Mosca, A. Parent, and J. Schanck. Estimating the Cost of Generic Quantum Pre-image Attacks on SHA-2 and SHA-3. In *Selected Areas in Cryptography - SAC 2016*, Springer LNCS vol. 10532, pp. 317–337, 2016.
- [Ami15] M. H. Amin. Searching for quantum speedup in quasistatic quantum annealers. *Phys. Rev. A*, 92(5):052323, 2015.
- [AMM14] M. Amy, D. Maslov, and M. Mosca. Polynomial-time T-depth Optimization of Clifford+T Circuits via Matroid Partitioning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 33(10): 1476–1489, 2014.
- [AMMR13] M. Amy, D. Maslov, M. Mosca, and M. Roetteler. A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, 32(6):818–830, 2013.
- [AND⁺17] V. V. Albert, K. Noh, K. Duivenvoorden, R. T. Brierley, P. Reinhold, C. Vuillot, L. Li, C. Shen, S. M. Girvin, B. M. Terhal, and L. Jiang. Performance and structure of bosonic codes, 2017, arXiv:1708.05010.
- [ANP⁺17] T. Astner, S. Nevlacsil, N. Peterschofsky, A. Angerer, S. Rotter, S. Putz, J. Schmiedmayer, and J. Majer. Coherent coupling of remote spin ensembles via a cavity bus. *Phys. Rev. Lett.*, 118(14):140502, 2017.
- [AOGC18a] E. R. Anschuetz, J. P. Olson, A. Aspuru-Guzik, and Y. Cao. Variational Quantum Factoring. In *Quantum Technology and Optimization Problems, First International Workshop, QTOP 2019*, Springer LNCS vol. 11413, pp. 74–85, 2019.
- [ARL⁺17] C.K.Andersen et al, Repeated Quantum Error Detection in a Surface Code, <https://arxiv.org/abs/1912.09410>
- [ARS13] B. Amento, M. Rötteler, and R. Steinwandt. Efficient quantum circuits for binary elliptic curve arithmetic: reducing T -gate complexity. *Quantum Inf. Comput.*, 13:631–644, 2013.
- [ASA⁺18] M. Almazrooie, A. Samsudin, R. Abdullah, and K. N. Mutter. Quantum reversible circuit of AES- 128. *Quantum Information Processing*, 5, 2018.
- [ASB⁺13] S. M. Anton, I. A. B. Sognaes, J. S. Birenbaum, S. R. O’Kelley, C. J. Fourie, and J. Clarke. Mean square flux noise in squids and qubits: numerical calculations. *Supercond. Sci. Technol.*, 26(7):075022, 2013.
- [ATB16] R. N. Alexander, P. S. Turner, and S. D. Bartlett. Randomized benchmarking in measurement-based quantum computing. *Phys. Rev. A*, 94(3):032303, 2016.
- [AvDK⁺07] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev. Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation. *SIAM J. Comput.*, 37(1):166–194, 2007, <http://dx.doi.org/10.1137/S0097539705447323>.
- [Bac13] D. Bacon. *Quantum Error Correction*, chapter Introduction to quantum error correction, pages 46–76. Cambridge University Press, 2013.
- [BB17] G. Banegas and D. J. Bernstein. Low-communication parallel quantum multi-target preimage search. In *Selected Areas in Cryptography – SAC 2017*, 2017.
- [BBC⁺17] L. S. Bishop, S. Bravyi, A. Cross, J. M. Gambetta, and J. Smolin. Quantum volume, 2017, https://dal.objectstorage.open.softlayer.com/v1/AUTH_039c3bf6e6e54d76b8e66152e2f87877/community-documents/quatum-volumehp08co1vbo0cc8fr.pdf.
- [BBHT98] M. Boyer, G. Brassard, P. Høyer, and A. Tapp. Tight bounds on quantum searching. *Fortschritte der Physik*, 46:493–506, 1998.

- [BBL16] A. Browaeys, D. Barredo, and T. Lahaye. Experimental investigations of dipole-dipole interactions between a few rydberg atoms. *J. Phys. B: At., Mol. Opt. Phys.*, 49(15):152001, 2016.
- [BBM17] D. Bernstein, J.-F. Biance, and M. Mosca. A Low-Resource Quantum Factoring Algorithm. Post-Quantum Cryptography – PQCrypto 2017, Springer LNCS vol. 10346, pp. 330–346, 2017.
- [BBV⁺16] V. B. Braginsky, I. A. Bilenko, S. P. Vyatchanin, M. L. Gorodetsky, V. P. Mitrofanov, L. G. Prokhorov, S. E. Strigin, and F. Y. Khalili. Background to the discovery of gravitational waves. *Phys. Usp.*, 59(9):879–885, 2016.
- [BC15] S. Bravyi and A. Cross. Doubled color codes, 2015, arXiv:1509.03239.
- [BCC⁺07] J. Baugh, J. Chamilliar, C. M. Chandrashekar, M. Ditty, A. Hubbard, R. Laflamme, M. Laforest, D. Maslov, O. Moussa, C. Negrevergne, M. Silva, S. Simmons, C. A. Ryan, D. G. Cory, J. S. Hodges, and C. Ramanathan. Quantum information processing using nuclear and electron magnetic resonance: review and prospects. 2007, arXiv:0710.1447.
- [BCI⁺16] Z. Bian, F. Chudak, R. Israel, B. Lackey, W. G. Macready, and A. Roy. Mapping constrained optimization problems to quantum annealing with application to fault diagnosis, 2016, arXiv:1603.03111.
- [BCOR09] S. Blanes, F. Casas, J. Oteo, and J. Ros. The Magnus expansion and some of its applications. *Phys. Rep.*, 470(5-6):151–238, 2009.
- [BDGD05] K. Bladh, T. Duty, D. Gunnarsson, and P. Delsing. The single cooper-pair box as a charge qubit. *New J. Phys.*, 7(1):180, 2005.
- [BdLL⁺16] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys. An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays. *Science*, 354(6315):1021–1023, 2016.
- [BDPA11] G. Bertoni, J. Daemen, M. Peeters, and G. V. Assche. The KECCAK reference. <http://keccak.noekeon.org/>, 2011.
- [Bea03] S. Beauregard. Circuit for Shor’s algorithm using $2n + 3$ qubits. *Quantum Inf. Comput.*, 3(2):175–185, 2003.
- [BEKP18] S. Bravyi, M. Englbrecht, R. König, and N. Peard. Correcting coherent errors with surface codes. *NPJ Quant. Inf.*, 4:55, 2018.
- [BFBA10] B. B. Buckley, G. D. Fuchs, L. C. Bassett, and D. D. Awschalom. Spin-light coherence for single-spin measurement and control in diamond. *Science*, 330(6008):1212–1215, 2010, <http://science.sciencemag.org/content/330/6008/1212.full.pdf>.
- [BFN⁺11] H. Bluhm, S. Foletti, I. Neder, M. Rudner, D. Mahalu, V. Umansky, and A. Yacoby. Dephasing time of gaas electron-spin qubits coupled to a nuclear bath exceeding 200[thinsp][mu]s. *Nat Phys*, 7(2):109–113, February 2011.
- [BGW⁺07] A. Blais, J. Gambetta, A. Wallraff, D. I. Schuster, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Quantum-information processing with circuit quantum electrodynamics. *Phys. Rev. A*, 75(3):032329, 2007.
- [BGW⁺15] S. Barzanjeh, S. Guha, C. Weedbrook, D. Vitali, J. H. Shapiro, and S. Pirandola. Microwave quantum illumination. *Phys. Rev. Lett.*, 114(8):080503, 2015.
- [BGY⁺11] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver. Noise spectroscopy through dynamical decoupling with a superconducting flux qubit. *Nat. Phys.*, 7(7):565–570, 2011.
- [BH12] S. Bravyi and J. Haah. Magic-state distillation with low overhead. *Phys. Rev. A*, 86(5):052329, 2012.
- [BHJ⁺14] P. I. Bunyk, E. M. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, A. J. Przybysz, and J. Whittaker. Architectural considerations in the design of a superconducting quantum annealing processor. *IEEE Transactions on Applied Superconductivity*, 24(4):1–10, Aug 2014.

- [BHL⁺16] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Phys. Rev. Lett.*, 117(6):060504, 2016.
- [BHP⁺14] C. W. Berry, M. R. Hashemi, S. Preu, H. Lu, A. C. Gossard, and M. Jarrahi. High power terahertz generation from eras: Ingaas plasmonic photomixers. In *2014 39th International Conference on Infrared, Millimeter, and Terahertz waves (IRMMW-THz)*, pages 1–2, 2014.
- [BHNP⁺19] X. Bonnetain, A. Hosoyamada, M. Naya-Plasencia, Y. Sasaki, and A. Schrottenloher, *Quantum Attacks Without Superposition Queries: The Offline Simon’s Algorithm*. In *Advances in Cryptology – ASIACRYPT 2019*, vol. 11921 of LNCS, pp. 552 – 583, Springer, 2019.
- [BHT98] G. Brassard, P. Høyer, and A. Tapp. Quantum Algorithm for the Collision Problem. In *Third Latin American Symposium on Theoretical Informatics (LATIN ’98)*, volume 1380 of *Lecture Notes in Computer Science*, pages 163–169. Springer, 1998.
- [Bia08] J. D. Biamonte. Nonperturbative k -body to two-body commuting conversion hamiltonians and embedding problem instances into ising spins. *Phys. Rev. A*, 77(5):052331, 2008.
- [BIS⁺16] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, J. M. Martinis, and H. Neven. Characterizing quantum supremacy in near-term devices, 2016, arXiv:1608.00263.
- [BK98] S. B. Bravyi and A. Y. Kitaev. Quantum codes on a lattice with boundary, 1998, arXiv:quant-ph/9811052.
- [BK05] S. Bravyi and A. Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Phys. Rev. A*, 71(2):022316, 2005.
- [BK13] S. Bravyi and R. König. Classification of topologically protected gates for local stabilizer codes. *Phys. Rev. Lett.*, 110(17):170503, 2013.
- [BKCD02] R. Blume-Kohout, C. Caves, and I. Deutsch. Climbing mount scalable: Physical resource requirements for a scalable quantum computer. *Found. Phys.*, 32(11):1641–1670, 2002.
- [BKGN⁺13] R. Blume-Kohout, J. K. Gamble, E. Nielsen, J. Mizrahi, J. D. Sterk, and P. Maunz. Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit, 2013, arXiv:1310.4492.
- [BKGN⁺17] R. Blume-Kohout, J. K. Gamble, E. Nielsen, K. Rudinger, J. Mizrahi, K. Fortier, and P. Maunz. Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography. *Nat. Commun.*, 8, 2017.
- [BKH⁺17] J. Benedikter, H. Kaupp, T. Hümmer, Y. Liang, A. Bommer, C. Becher, A. Krueger, J. M. Smith, T. W. Hänsch, and D. Hunger. Cavity-enhanced single-photon source based on the silicon-vacancy center in diamond. *Phys. Rev. Applied*, 7:024031, Feb 2017.
- [BKLW00] D. Bacon, J. Kempe, D. A. Lidar, and K. B. Whaley. Universal fault-tolerant quantum computation on decoherence-free subspaces. *Phys. Rev. Lett.*, 85(8):1758–1761, 2000.
- [BKM⁺13] R. Barends, J. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, C. Neill, P. O’Malley, P. Roushan, J. Wenner, T. C. White, A. N. Cleland, and J. M. Martinis. Coherent josephson qubit suitable for scalable quantum integrated circuits. *Phys. Rev. Lett.*, 111(8):080502, 2013.
- [BKM⁺14] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014.
- [BKM16] K. R. Brown, J. Kim, and C. Monroe. Co-designing a scalable quantum computer with trapped atomic ions. *Nat. Partn. J. Quantum Inf.*, 2:16034, 2016.

- [BKM18] A. Botea, A. Kishimoto, and Radu Marinescu. On the Complexity of Quantum Circuit Compilation. 11th Annual Symposium on Combinatorial Search (SoCS 2018), AAAI Publications, 2018, <https://aaai.org/ocs/index.php/SOCS/SOCS18/paper/view/17959>.
- [BKRB15] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt. Ion-trap measurements of electric-field noise near surfaces. *Rev. Mod. Phys.*, 87(4):1419–1482, 2015.
- [BLAG14] R. Babbush, P. J. Love, and A. Aspuru-Guzik. Adiabatic quantum simulation of quantum chemistry. *Sci. Rep.*, 4:6603 EP, 2014.
- [BLK⁺15] R. Barends, L. Lamata, J. Kelly, L. García-Álvarez, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O’Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, E. Solano, and J. M. Martinis. Digital quantum simulation of fermionic models with a superconducting circuit. *Nat. Commun.*, 6:7654, 2015.
- [BLPV16] H.-P. Breuer, E.-M. Laine, J. Piilo, and B. Vacchini. Colloquium. *Rev. Mod. Phys.*, 88(2):021002, 2016.
- [BMD07] H. Bombin and M. A. Martin-Delgado. Topological computation without braiding. *Phys. Rev. Lett.*, 98(16):160502, 2007.
- [BMR⁺06] J. Baugh, O. Moussa, C. A. Ryan, R. Laflamme, C. Ramanathan, T. F. Havel, and D. G. Cory. Solid-state nmr three-qubit homonuclear system for quantum-information processing: Control and characterization. *Phys. Rev. A*, 73(2):022305, 2006.
- [BNB16] B. J. Brown, N. H. Nickerson, and D. E. Browne. Fault-tolerant error correction with the gauge color code. *Nat. Commun.*, 7:12302, 2016.
- [BNP18] X. Bonnetain and M. Naya-Plasencia. Hidden Shift Quantum Cryptanalysis and Implications. *Advances in Cryptology - ASIACRYPT 2018, Lecture Notes in Computer Science* vol. 11272, pp. 560–592, Springer, 2018.
- [BNPS19] X. Bonnetain, M. Naya-Plasencia, and A. Schrottenloher. *Quantum Security Analysis of AES*. IACR Transactions on Symmetric Cryptology, 2019(2), 55–93, 2019.
- [BOV⁺09] R. B. Blakestad, C. Ospelkaus, A. P. VanDevender, J. M. Amini, J. Britton, D. Leibfried, and D. J. Wineland. High-fidelity transport of trapped-ion qubits through an X-junction trap array. *Phys. Rev. Lett.*, 102(15):153002, 2009.
- [Boy03] R. W. Boyd. *Nonlinear Optics (Second Edition)*. Academic Press, 2003.
- [BP02] H. P. Breuer and F. Petruccione. *The theory of open quantum systems*. Oxford University Press, 2002.
- [BPIG14] M. T. Bell, J. Paramanandam, L. B. Ioffe, and M. E. Gershenson. Protected josephson rhombus chains. *Phys. Rev. Lett.*, 112(16):167001, 2014.
- [BR05] D. E. Browne and T. Rudolph. Resource-efficient linear optical quantum computation. *Phys. Rev. Lett.*, 95(1):010501, 2005.
- [BR12] R. Blatt and C. F. Roos. Quantum simulations with trapped ions. *Nat. Phys.*, 8(4):277–284, 2012.
- [Bra98] S. L. Braunstein. Error correction for continuous quantum variables. *Phys. Rev. Lett.*, 80(18):4084–4087, 1998.
- [Bra17] M. F. Brandl. A quantum von Neumann architecture for large-scale quantum computing in systems with long coherence times, such as trapped ions, 2017, arxiv:1702.02583v1, <http://arxiv.org/abs/1702.02583v1>.
- [BS15] P. Budhathoki and R. Steinwandt. Automatic synthesis of quantum circuits for point addition on ordinary binary elliptic curves. *Quantum Inf. Process.*, 14(1):201–216, 2015.
- [BSA17] G. Burkard, V. O. Shkolnikov, and D. D. Awschalom. Designing a cavity-mediated quantum cphase gate between nv spin qubits in diamond. *Phys. Rev. B*, 95:205420, May 2017.

- [BSBN02] S. D. Bartlett, B. C. Sanders, S. L. Braunstein, and K. Nemoto. Efficient classical simulation of continuous variable quantum information processes. *Phys. Rev. Lett.*, 88(9):097904, 2002.
- [BSK⁺12] J. W. Britton, B. C. Sawyer, A. C. Keith, C.-C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger. Engineered two-dimensional ising interactions in a trapped-ion quantum simulator with hundreds of spins. *Nature*, 484(7395):489–492, 2012.
- [BSK⁺17] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin. Probing many-body dynamics on a 51-atom quantum simulator, 2017, arXiv:1707.04344.
- [BSL⁺16] R. Barends, A. Shabani, L. Lamata, J. Kelly, A. Mezzacapo, U. L. Heras, R. Babbush, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, P. J. J. O’Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, E. Solano, H. Neven, and J. M. Martinis. Digitized adiabatic quantum computing with a superconducting circuit. *Nature*, 534(7606):222–226, 2016.
- [BSM⁺10] N. Bergeal, F. Schackert, M. Metcalfe, R. Vijay, V. E. Manucharyan, L. Frunzio, D. E. Prober, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret. Phase-preserving amplification near the quantum limit with a josephson ring modulator. *Nature*, 465(7294):64–68, 2010.
- [BSV01] V. Braginsky, S. Strigin, and S. Vyatchanin. Parametric oscillatory instability in fabry-perot interferometer. *Phys. Lett. A*, 287(5):331 – 338, 2001.
- [BSV14] S. Bravyi, M. Suchara, and A. Vargo. Efficient algorithms for maximum likelihood decoding in the surface code. *Phys. Rev. A*, 90(3):032326, 2014.
- [BVAC13] J. Bochmann, A. Vainsencher, D. D. Awschalom, and A. N. Cleland. Nanomechanical coupling between microwave and optical photons. *Nat. Phys.*, 9(11):712–716, 2013.
- [BVJ⁺98] V. Bouchiat, D. Vion, P. Joyez, D. Esteve, and M. H. Devoret. Quantum coherence with a single cooper pair. *Phys. Scr.*, 1998(T76):165, 1998.
- [BvL05] S. L. Braunstein and P. van Loock. Quantum information with continuous variables. *Rev. Mod. Phys.*, 77(2):513–577, 2005.
- [BW08a] R. Blatt and D. Wineland. Entangled states of trapped atomic ions. *Nature*, 453(7198):1008–1015, 2008.
- [BW08b] L. Bogani and W. Wernsdorfer. Molecular spintronics using single-molecule magnets. *Nat. Mater.*, 7(3):179–186, 2008.
- [BWC⁺11] K. R. Brown, A. C. Wilson, Y. Colombe, C. Ospelkaus, A. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland. Single-qubit-gate error below 10^{-4} in a trapped ion. *Phys. Rev. A*, 84(3):030303, 2011.
- [BWG⁺18] S.J. Beale, J.J. Wallmann, M. Gutierrez, K.R. Brown, and R. Laflamme, Quantum error correction decoheres noise, 2018, arXiv:1805.08802.]
- [BWM⁺16] B. Bauer, D. Wecker, A. J. Millis, M. B. Hastings, and M. Troyer. Hybrid quantum-classical approach to correlated materials. *Phys. Rev. X*, 6(3):031045, 2016.
- [BXN⁺17] A. Bermudez, X. Xu, R. Nigmatullin, J. O’Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller. Assessing the progress of trapped-ion processors towards fault-tolerant quantum computation, 2017, arXiv:1705.02771.
- [CB06] J. Clarke and A. Braginski. *The SQUID Handbook: Fundamentals and Technology of SQUIDS and SQUID Systems*. Wiley, 2006.
- [CBDB14] P. Cerfontaine, T. Botzem, D. P. DiVincenzo, and H. Bluhm. High-fidelity single-qubit gates for two-electron spin qubits in gaas. *Phys. Rev. Lett.*, 113:150501, Oct 2014.

- [CBIH⁺08] M. A. Castellanos-Beltran, K. D. Irwin, G. C. Hilton, L. R. Vale, and K. W. Lehnert. Amplification and squeezing of quantum noise with a tunable Josephson metamaterial. *Nat. Phys.*, 4(12):929–931, 2008.
- [CBS⁺04] I. Chiorescu, P. Bertet, K. Semba, Y. Nakamura, C. J. P. M. Harmans, and J. E. Mooij. Coherent dynamics of a flux qubit coupled to a harmonic oscillator. *Nature*, 431(7005):159–162, 2004.
- [CBSG17] A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta. Open quantum assembly language, 2017, arXiv:1707.03429.
- [CCG⁺11] J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, J. A. Smolin, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen. Simple all-microwave entangling gate for fixed-frequency superconducting qubits. *Phys. Rev. Lett.*, 107(8):080502, 2011.
- [CDG⁺10] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf. Introduction to quantum noise, measurement, and amplification. *Rev. Mod. Phys.*, 82(2):1155–1208, 2010.
- [CDR⁺17] S. Caldwell, N. Didier, C. A. Ryan, E. A. Sete, A. Hudson, P. Karalekas, R. Manenti, M. Reagor, M. P. da Silva, R. Sinclair, E. Acala, N. Alidoust, J. Angeles, A. Bestwick, M. Block, B. Bloom, A. Bradley, C. Bui, L. Capelluto, R. Chilcott, J. Cordova, G. Crossman, M. Curtis, S. Deshpande, T. E. Bouayadi, D. Girshovich, S. Hong, K. Kuang, M. Lenihan, T. Manning, J. Marshall, Y. Mohan, W. O’Brien, C. Osborn, J. Otterbach, A. Papageorge, J. P. Paquette, M. Pelstring, A. Polloreno, G. Prawiroatmodjo, V. Rawat, R. Renzas, N. Rubin, D. Russell, M. Rust, D. Scarabelli, M. Scheer, M. Selvanayagam, R. Smith, A. Staley, M. Suska, N. Tezak, T. W. To, M. Vahidpour, N. Vodrahalli, T. Whyland, K. Yadav, W. Zeng, and C. Rigetti. Parametrically-activated entangling gates using transmon qubits, 2017, arXiv:1706.06562.
- [CFLLS14] E. Crosson, E. Farhi, C. Y. Lin, H. Lin, P. Shor, Different Strategies for Optimization Using the Quantum Adiabatic Algorithm,, <https://arxiv.org/abs/1401.7320>
- [CFGR⁺13] H. O. H. Churchill, V. Fatemi, K. Grove-Rasmussen, M. T. Deng, P. Caroff, H. Q. Xu, and C. M. Marcus. Superconductor-nanowire devices from tunneling to the multichannel regime: Zero-bias oscillations and magnetoconductance crossover. *Phys. Rev. B*, 87(24):241401, 2013.
- [CGFF17] J. Combes, C. Granade, C. Ferrie, and S. T. Flammia. Logical randomized benchmarking, 2017, arXiv:1702.03688.
- [CHS⁺15] J. Carolan, C. Harrold, C. Sparrow, E. Martín-López, N. J. Russell, J. W. Silverstone, P. J. Shadbolt, N. Matsuda, M. Oguma, M. Itoh, G. D. Marshall, M. G. Thompson, J. C. F. Matthews, T. Hashimoto, J. L. O’Brien, and A. Laing. Universal linear optics. *Science*, 349(6249):711–716, 2015.
- [ChGa18] Y.-A. Chen, X.-S. Gao. *Quantum algorithms for Boolean equation solving and quantum algebraic attack on cryptosystems*. arXiv:1712.06239v3, 2018.
- [CJS13] B. D. Clader, B. C. Jacobs, and C. R. Sprouse. Preconditioned Quantum Linear System Algorithm. *Phys. Rev. Lett.*, 110(250504), 2013.
- [CKS17] A. M. Childs, R. Kothari, and R. D. Somma. *Quantum Algorithm for Systems of Linear Equations with Exponentially Improved Dependence on Precision*. SIAM Journal on Computing 46.6, pp. 1920–1950, 2017.
- [Cle10] A. Cleland. Nanoelectromechanical resonators. In K. Sattler, editor, *Handbook of Nanophysics: Functional Nanomaterials*, chapter 37, pages 266–290. CRC Press (Taylor & Francis), 2010.
- [CLK⁺00] D. Cory, R. Laflamme, E. Knill, L. Viola, T. Havel, N. Boulant, G. Boutis, E. Fortunato, S. Lloyd, R. Martinez, C. Negrevergne, M. Pravia, Y. Sharf, G. Teklemariam, Y. Weinstein, and W. Zurek. Nmr based quantum information processing: Achievements and prospects. *Fortschr. Phys.*, 48(9-11):875–907, 2000.
- [CMB⁺10] J. H. Cole, C. Müller, P. Bushev, G. J. Grabovskij, J. Lisenfeld, A. Lukashenko, A. V. Ustinov, and A. Shnirman. Quantitative evaluation of defect-models in superconducting phase qubits. *Appl. Phys. Lett.*, 97(25):252501, 2010.

- [CMB⁺16] A. W. Cross, E. Magesan, L. S. Bishop, J. A. Smolin, and J. M. Gambetta. Scalable randomised benchmarking of non-clifford gates. *Nat. Partn. J. Quantum Inf.*, 2:16012 EP, 2016.
- [CMM99] P. T. Cochrane, G. J. Milburn, and W. J. Munro. Macroscopically distinct quantum-superposition states as a bosonic code for amplitude damping. *Phys. Rev. A*, 59(4):2631–2634, 1999.
- [CMP14] M. Chen, N. C. Menicucci, and O. Pfister. Experimental realization of multipartite entanglement of 60 modes of a quantum optical frequency comb. *Phys. Rev. Lett.*, 112(12):120505, 2014.
- [CMQ⁺10] W. C. Campbell, J. Mizrahi, Q. Quraishi, C. Senko, D. Hayes, D. Hucul, D. N. Matsukevich, P. Maunz, and C. Monroe. ultra-fast gates for single atomic qubits. *Phys. Rev. Lett.*, 105(9):090502, 2010.
- [CMR⁺16] B. J. Chapman, B. A. Moores, E. I. Rosenthal, J. Kerckhoff, and K. W. Lehnert. General purpose multiplexing device for cryogenic microwave systems. *Appl. Phys. Lett.*, 108(22):222602, 2016.
- [CNHM03] I. Chiorescu, Y. Nakamura, C. J. P. M. Harmans, and J. E. Mooij. Coherent quantum dynamics of a superconducting flux qubit. *Science*, 299(5614):1869–1871, 2003.
- [CNPS17] A. Chailloux, M. Naya-Plasencia, and A. Schrottenloher. *An efficient quantum collision search algorithm and implications on symmetric cryptography*. In ASIACRYPT (2), vol. 10625 of LNCS, pp. 211–240. Springer, 2017.
- [CNR⁺14] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O’Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, M. R. Geller, A. N. Cleland, and J. M. Martinis. Qubit architecture with high coherence and fast tunable coupling. *Phys. Rev. Lett.*, 113(22):220502, 2014.
- [Cop94] D. Coppersmith. An Approximate Fourier Transform Useful in Quantum Factoring, 1994.
- [CR14] J. I. Colless and D. J. Reilly. Modular cryogenic interconnects for multi-qubit devices. *Rev. Sci. Instrum.*, 85(11):114706, 2014.
- [CRFG17] D. S. A. Coden, R. H. Romero, A. Ferrón, and S. S. Gomez. Optimal control of a charge qubit in a double quantum dot with a coulomb impurity. *Physica E: Low-dimensional Systems and Nanostructures*, 86:36–43, 2017.
- [CRKW17] T. Chasseur, D. M. Reich, C. P. Koch, and F. K. Wilhelm. Hybrid benchmarking of arbitrary quantum gates. *Phys. Rev. A*, 95(6), 2017.
- [CSDS17] D. J. Clarke, J. D. Sau, and S. Das Sarma. Probability and braiding statistics in majorana nanowires. *Phys. Rev. B*, 95(15):155451, 2017.
- [CTV16] E. T. Campbell, B. M. Terhal, and C. Vuillot. The steep road towards robust and universal quantum computation, 2016, arXiv:1612.07330.
- [CW00] R. Cleve and J. Watrous. Fast parallel circuits for the quantum Fourier transform. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 526–536, 2000.
- [CW08] J. Clarke and F. K. Wilhelm. Superconducting quantum bits. *Nature*, 453(7198):1031–1042, 2008.
- [CW15] T. Chasseur and F. K. Wilhelm. Complete randomized benchmarking protocol accounting for leakage errors. *Phys. Rev. A*, 92(4), 2015.
- [CZ95] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Phys. Rev. Lett.*, 74(20):4091–4094, 1995.
- [CZW16] N. Chancellor, S. Zohren, and P. A. Warburton. Circuit design for multi-body interactions in superconducting quantum annealing system with applications to a scalable architecture, 2016, arXiv:1603.09521.
- [DB15] N. S. Dattani and N. Bryans. Quantum factorization of 56153 with only 4 qubits. arXiv:1411.6758, 2015.

- [DBI⁺16] V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babbush, V. Smelyanskiy, J. Martinis, and H. Neven. What is the computational value of finite-range tunneling? *Phys. Rev. X*, 6(3):031015, 2016.
- [DBK⁺00] D. DiVincenzo, D. Bacon, J. Kempe, G. Burkard, and K. Whaley. Universal quantum computation with the exchange interaction. *Nature*, 408(6810):339–342, 2000.
- [DCJ⁺07] M. V. G. Dutt, L. Childress, L. Jiang, E. Togan, J. Maze, F. Jelezko, A. S. Zibrov, P. R. Hemmer, and M. D. Lukin. Quantum register based on individual electronic and nuclear spin qubits in diamond. *Science*, 316(5829):1312–1316, 2007, <http://science.sciencemag.org/content/316/5829/1312.full.pdf>.
- [DDBA13] I. Diniz, E. Dumur, O. Buisson, and A. Auffèves. ultra-fast quantum nondemolition measurements based on a diamond-shaped artificial atom. *Phys. Rev. A*, 87(3):033837, 2013.
- [DDW16] P.-L. Dallaire-Demers and F. K. Wilhelm. Quantum gates and architecture for the quantum simulation of the fermi-hubbard model. *Phys. Rev. A*, 94(6):062304, 2016.
- [DFII05] B. Douçot, M. V. Feigel'man, L. B. Ioffe, and A. S. Ioselevich. Protected qubits and chern-simons theories in Josephson junction arrays. *Phys. Rev. B*, 71(2):024505, 2005.
- [DH81] P. Dutta and P. M. Horn. Low-frequency fluctuations in solids: 1/f noise. *Rev. Mod. Phys.*, 53(3):497–516, 1981.
- [DHMS⁺18] D. Dervovic, M. Herbster, P. Mountney, S. Severini, N. Usher, and L. Wossnig. *Quantum linear systems algorithms: a primer*. arXiv:1802.08227v1, 2018.
- [DHN06a] C. M. Dawson, H. L. Haselgrove, and M. A. Nielsen. Noise thresholds for optical cluster-state quantum computation. *Phys. Rev. A*, 73(5):052306, 2006.
- [DHN06b] C. M. Dawson, H. L. Haselgrove, and M. A. Nielsen. Noise thresholds for optical quantum computers. *Phys. Rev. Lett.*, 96(2):020501, 2006.
- [DiV00] D. P. DiVincenzo. The physical implementation of quantum computation, 2000, arXiv:quant-ph/0002077v3, <http://arXiv.org/abs/quant-ph/0002077v3>.
- [DJA⁺13] N. G. Dickson, M. W. Johnson, M. H. Amin, R. Harris, F. Altomare, A. J. Berkley, P. Bunyk, J. Cai, E. M. Chapple, P. Chavez, F. Cioata, T. Cirip, P. deBuen, M. Drew-Brook, C. Enderud, S. Gildert, F. Hamze, J. P. Hilton, E. Hoskinson, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Lanting, T. Mahon, R. Neufeld, T. Oh, I. Perminov, C. Petroff, A. Przybysz, C. Rich, P. Spear, A. Tcaciuc, M. C. Thom, E. Tolkacheva, S. Uchaikin, J. Wang, A. B. Wilson, Z. Merali, and G. Rose. Thermally assisted quantum annealing of a 16-qubit problem. *Nat. Commun.*, 4:1903, 2013.
- [DK12] S. L. Danilishin and F. Y. Khalili. Quantum measurement theory in gravitational-wave detectors. *Living Rev. Relativ.*, 15(1):5, 2012.
- [dLvHB⁺15] G. de Lange, B. van Heck, A. Bruno, D. J. van Woerkom, A. Geresdi, S. R. Plissard, E. P. A. M. Bakkers, A. R. Akhmerov, and L. DiCarlo. Realization of microwave quantum circuits using hybrid superconducting-semiconducting nanowire josephson elements. *Phys. Rev. Lett.*, 115(12):127002, 2015.
- [DM12] S. Ding and D. N. Matsukevich. Quantum logic for the control and manipulation of molecular ions using a frequency comb. *New J. Phys.*, 14(2):023028, 2012.
- [DMBK⁺16] J. P. Dehollain, J. T. Muhonen, R. Blume-Kohout, K. M. Rudinger, J. K. Gamble, E. Nielsen, A. Laucht, S. Simmons, R. Kalra, A. S. Dzurak, and A. Morello. Optimization of a solid-state electron spin qubit using gate set tomography. *New J. Phys.*, 18(10):103018, 2016.
- [DMC85] M. H. Devoret, J. M. Martinis, and J. Clarke. Measurements of macroscopic quantum tunneling out of the zero-voltage state of a current-biased josephson junction. *Phys. Rev. Lett.*, 55(18):1908–1911, 1985.
- [DMT⁺14] J. P. Dehollain, J. T. Muhonen, K. Y. Tan, A. Saraiva, D. N. Jamieson, A. S. Dzurak, and A. Morello. Single-shot readout and relaxation of singlet and triplet states in exchange-coupled ³¹P electron spins in silicon. *Phys. Rev. Lett.*, 112(23):236801, 2014.

- [DOS⁺15] C. Deng, J.-L. Orgiazzi, F. Shen, S. Ashhab, and A. Lupascu. Observation of floquet states in a strongly driven artificial atom. *Phys. Rev. Lett.*, 115(13):133601, 2015.
- [DR99] J. Daemen and V. Rijmen. The Rijndael Block Cipher, 1999.
- [DRM⁺12] A. Das, Y. Ronen, Y. Most, Y. Oreg, M. Heiblum, and H. Shtrikman. Zero-bias peaks and splitting in an al-InAs nanowire topological superconductor as a signature of majorana fermions. *Nat. Phys.*, 8(12):887–895, 2012.
- [DSFN06] S. Das Sarma, M. Freedman, and C. Nayak. Topological quantum computation. *Phys. Today*, 2006.
- [DSH⁺13] O. E. Dial, M. D. Shulman, S. P. Harvey, H. Bluhm, V. Umansky, and A. Yacoby. Charge noise spectroscopy using coherent exchange oscillations in a singlet-triplet qubit. *Phys. Rev. Lett.*, 110:146804, Apr 2013.
- [DSMN16] S. J. Devitt, A. M. Stephens, W. J. Munro, and K. Nemoto. *Analysis of an Atom-Optical Architecture for Quantum Computation*, pages 407–437. Springer Japan, Tokyo, 2016.
- [DSNT06] S. Das Sarma, C. Nayak, and S. Tewari. Proposal to stabilize and detect half-quantum vortices in strontium ruthenate thin films: Non-abelian braiding statistics of vortices in a $p_x + ip_y$ superconductor. *Phys. Rev. B*, 73(22):220502, 2006.
- [dVRT02] R. de Vivie-Riedle and C. Tesch. Molecular quantum computing: Implementation of global quantum gates applying optimal control theory. In *The Thirteenth International Conference on ultra-fast Phenomena*, page FB2. Optical Society of America, 2002.
- [EAÅ05] J. Emerson, R. Alicki, and K. Åyzykowski. Scalable noise estimation with random unitary operators. *J. Opt. B: Quantum Semiclassical Opt.*, 7(10):S347, 2005.
- [ECMG14] J. M. Epstein, A. W. Cross, E. Magesan, and J. M. Gambetta. Investigating the limits of randomized benchmarking protocols. *Phys. Rev. A*, 89(6):062321, 2014.
- [Edm65] J. Edmonds. Paths, trees, and flowers. *Canad. J. Math.*, 17(0):449–467, 1965.
- [EGMW11] Y. Elias, H. Gilboa, T. Mor, and Y. Weinstein. Heat-bath cooling of spins in two amino acids. *Chem. Phys. Lett.*, 517(4-6):126 – 131, 2011.
- [EH17] M. Ekerå and J. Håstad. Quantum algorithms for computing short discrete logarithms and factoring rsa integers. Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017, Lecture Notes in Computer Science vol. 10346, pp. 347–363, Springer, 2017.
- [EHWvB⁺04] J. M. Elzerman, R. Hanson, L. H. Willems van Beveren, B. Witkamp, L. M. K. Vandersypen, and L. P. Kouwenhoven. Single-shot read-out of an individual electron spin in a quantum dot. *Nature*, 430(6998):431–435, 2004.
- [EHWR+19] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, E. Kashefi, Quantum certification and benchmarking, <https://arxiv.org/abs/1910.06343>
- [Eke16] M. Ekerå. Modifying Shor's algorithm to compute short discrete logarithms. Cryptology ePrint Archive, Report 2016/1128, 2016.
- [Eke18] M. Ekerå. Quantum algorithms for computing general discrete logarithms and orders with tradeoffs. Cryptology ePrint Archive, Report 2018/797, 2018
- [Eke19] M. Ekerå. Revisiting Shor's quantum algorithm for computing general discrete logarithms. arXiv:1905.09084v1, 2019.
- [ELR10] S. E. Economou, N. Lindner, and T. Rudolph. Optically generated 2-dimensional photonic cluster state from coupled quantum dots. *Phys. Rev. Lett.*, 105(9):093601, 2010.

- [FBS⁺15] T. Feldker, P. Bachor, M. Stappel, D. Kolbe, R. Gerritsma, J. Walz, and F. Schmidt-Kaler. Rydberg excitation of a single trapped ion. *Phys. Rev. Lett.*, 115(17):173001, 2015.
- [FCH⁺17] M. A. Fogarty, K. W. Chan, B. Hensen, W. Huang, T. Tanttu, C. H. Yang, A. Laucht, M. Veldhorst, F. E. Hudson, K. M. Itoh, D. Culcer, A. Morello, and A. S. Dzurak. Integrated silicon qubit platform with single-spin addressability, exchange control and robust single-shot singlet-triplet readout, 2017, arXiv:1708.03445.
- [FD12] A. G. Fowler and S. J. Devitt. A bridge to lower overhead quantum computation, 2012, arXiv:1209.0510.
- [FDJ13] A. G. Fowler, S. J. Devitt, and C. Jones. Surface code implementation of block code state distillation. *Sci. Rep.*, 3(1), 2013.
- [FEF⁺08] I. Fushman, D. Englund, A. Faraon, N. Stoltz, P. Petroff, and J. Vučković. Controlled phase shifts with a single quantum dot. *Science*, 320(5877):769–772, 2008, <http://science.sciencemag.org/content/320/5877/769.full.pdf>.
- [FFE⁺08] A. Faraon, I. Fushman, D. Englund, N. Stoltz, P. Petroff, and J. Vuckovic. Coherent generation of non-classical light on a chip via photon-induced tunnelling and blockade. *Nat. Phys.*, 4(11):859–863, 2008.
- [FFSG09] A. G. Fowler, A. G. Fowler, A. M. Stephens, and P. Groszkowski. High threshold universal quantum computation on the surface code. *Phys. Rev. A*, 80:052312, 2009.
- [FG09] A. G. Fowler and K. Goyal. Topological cluster state quantum computing. *Quantum Inf. Comput.*, 9(9):721–738, 2009.
- [FG18] A.G. Fowler and C. Gidney. Low overhead quantum computation using lattice surgery. ArXiv:1808.06709.
- [FGG14] E. Farhi, J. Goldstone, and S. Gutmann. A Quantum Approximate Optimization Algorithm. ArXiv:1411.4028
- [FHZ14] D. Fangwei, W. Hong, and M. Zhi. Quantum Collision Search Algorithm Against New FORK-256. *J. Electron.*, 31, 2014.
- [FM14] A. G. Fowler and J. M. Martinis. Quantifying the effects of local many-qubit errors and nonlocal two-qubit errors on the surface code. *Phys. Rev. A*, 89(3):032316, 2014.
- [FML⁺17] C. Figgatt, D. Maslov, K. A. Landsman, N. M. Linke, S. Debnath, and C. Monroe. Complete 3-Qubit Grover search on a programmable quantum computer. *Nat. Comm.* 8:1918 (2017).
- [FMMC12] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86(3):032324, 2012.
- [FMP09] S. T. Flammia, N. C. Menicucci, and O. Pfister. The optical frequency comb as a one-way quantum computer. *J. Phys. B: At., Mol. Opt. Phys.*, 42(11):114009, 2009.
- [Föl14] S. Fölling. Quantum noise correlation experiments with ultracold atoms, 2014, arXiv:1403.6842.
- [Fow12] A. G. Fowler. Time-optimal quantum computation, 2012, arXiv:1210.4626.
- [Fow13a] A. G. Fowler. Analytic asymptotic performance of topological codes. *Phys. Rev. A*, 87(4):040301, 2013.
- [Fow13b] A. G. Fowler. Polyestimate: instantaneous open source surface code analysis, 2013, arXiv:1307.0689.
- [Fow15] A. G. Fowler. Minimum weight perfect matching of fault-tolerant topological quantum error correction in average $\mathcal{O}(1)$ parallel time. *Quantum Inf. Comput.*, 15:145–158, 2015.
- [Fox06] M. Fox. *Quantum Optics: An Introduction (Oxford Master Series in Physics, 6)*. Oxford University Press, USA, 2006.

- [FPC⁺00] J. R. Friedman, V. Patel, W. Chen, S. K. Tolpygo, and J. E. Lukens. Quantum superposition of distinct macroscopic states. *Nature*, 406(6791):43–46, 2000.
- [FSG09] A. G. Fowler, A. M. Stephens, and P. Groszkowski. High-threshold universal quantum computation on the surface code. *Phys. Rev. A*, 80(5):052312, 2009.
- [FSL⁺17] S. Freer, S. Simmons, A. Laucht, J. T. Muhonen, J. P. Dehollain, R. Kalra, F. A. Mohiyaddin, F. E. Hudson, K. M. Itoh, J. C. McCallum, D. N. Jamieson, A. S. Dzurak, and A. Morello. A single-atom quantum memory in silicon. *Quantum Sci. Technol.*, 2(1):015009, 2017.
- [FSLC97] P. Fisk, M. Sellars, M. Lawn, and G. Coles. Accurate measurement of the 12.6 GHz "clock" transition in trapped $^{171}\text{Yb}^+$ ions. *IEEE Trans. Ultrason., Ferroelect., Freq. Control*, 44(2):344–354, 1997.
- [FTCL09] J. Fischer, M. Trif, W. Coish, and D. Loss. Spin interactions, relaxation and decoherence in quantum dots. *Solid State Commun.*, 149(35):1443 – 1450, 2009.
- [FWMR12] A. G. Fowler, A. C. Whiteside, A. L. McInnes, and A. Rabbani. Topological code Autotune. *Phys. Rev. X*, 2(4):041003, 2012.
- [FWS⁺18] A. Fornieri, A.M. Whiticar, F. Setiawan, E. Portolés Marín, A.C.C. Drachmann, A. Keselman, S. Gronin, C. Thomas, T. Wang, R. Kallaher, G.C. Gardner, E. Berg, M.J. Manfra, A. Stern, C.M. Marcus, F. Nichele. Evidence of topological superconductivity in planar Josephson junctions. arXiv:1809.03037
- [FZ01] Rosario Fazio, Herre van der Zant, Quantum phase transitions and vortex dynamics in superconducting networks, Physics Report, Volume 355, Issue 4, December 2001, Pages 235-334
- [GBC⁺15] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. y, S. Schirmer, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm. Training Schrödinger's cat: quantum optimal control. *Eur. Phys. J. D*, 69(12):279, 2015.
- [GBY⁺11] S. Gustavsson, J. Bylander, F. Yan, W. D. Oliver, F. Yoshihara, and Y. Nakamura. Noise correlations in a flux qubit with tunable tunnel coupling. *Phys. Rev. B*, 84:014525, Jul 2011.
- [GCZ⁺17] S. Gazibegovic, D. Car, H. Zhang, S. C. Balk, J. A. Logan, M. W. A. de Moor, M. C. Cassidy, R. Schmits, D. Xu, G. Wang, P. Krogstrup, R. L. M. O. het Veld, K. Zuo, Y. Vos, J. Shen, D. Bouman, B. Shojaei, D. Pennachio, J. S. Lee, P. J. van Veldhoven, S. Koelling, M. A. Verheijen, L. P. Kouwenhoven, C. J. Palmstrøm, and E. P. A. M. Bakkers. Epitaxy of advanced nanowire quantum devices. *Nature*, 548(7668):434–438, 2017.
- [GD16] D. Greenbaum and Z. Dutton. Modeling coherent errors in quantum error correction. arXiv:1612.03908
- [GF15] J. Ghosh and A. G. Fowler. Leakage-resilient approach to fault-tolerant quantum computing with superconducting elements. *Phys. Rev. A*, 91(2):020302, 2015.
- [GF17] C. Gidney and A. Fowler. A slightly smaller surface code s gate, 2017, arXiv:1708.00054.
- [GFM10] O. Golubitsky, S. M. Falconer, and D. Maslov. Synthesis of the Optimal 4-bit Reversible Circuits. In Proc. Of the 47th Design Automation Conference DAC '10, pp. 653-656, ACM, 2010.
- [GGK⁺11] L. Gaudreau, G. Granger, A. Kam, G. C. Aers, S. A. Studenikin, P. Zawadzki, M. Pioro-Ladrière, Z. R. Wasilewski, and A. S. Sachrajda. Coherent control of three-spin states in a triple quantum dot. *Nat. Phys.*, 8(1):54–58, 2011.
- [GGZ⁺13] J. Ghosh, A. Galiutdinov, Z. Zhou, A. N. Korotkov, J. M. Martinis, and M. R. Geller. High-fidelity controlled- σ^Z gate for resonator-based superconducting quantum computers. *Phys. Rev. A*, 87(2):022309, 2013.
- [Gid18] C. Gidney: *Halving the cost of quantum addition*. Quantum 2, 74 (2018).
- [GiEk19] C. Gidney and M. Ekerå. *How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits*. arXiv:1905.09749v2, 2019.

- [Gir17] D. Giry. BlueKrypt | Cryptographic Key Length Recommendation, 2017.
- [GKP01] D. Gottesman, A. Kitaev, and J. Preskill. Encoding a qubit in an oscillator. *Phys. Rev. A*, 64(1):012310, 2001.
- [GKV06] S. Glancy, E. Knill, and H. M. Vasconcelos. Entanglement purification of any stabilizer state. *Phys. Rev. A*, 74(3):032319, 2006.
- [GLF⁺10] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert. Quantum state tomography via compressed sensing. *Phys. Rev. Lett.*, 105(15):150401, 2010.
- [GLRS16] M. Grassl, B. Langenberg, M. Roetteler, and R. Steinwandt. Applying Grover’s Algorithm to AES: Quantum Resource Estimates. In *Post-Quantum Cryptography*, volume 9606 of *Lecture Notes in Computer Science*. Springer, 2016.
- [GMT⁺12] J. P. Gaebler, A. M. Meier, T. R. Tan, R. Bowler, Y. Lin, D. Hanneke, J. D. Jost, J. P. Home, E. Knill, D. Leibfried, and D. J. Wineland. Randomized benchmarking of multiqubit gates. *Phys. Rev. Lett.*, 108(26):260503, 2012.
- [Got98] D. Gottesman. The heisenberg representation of quantum computers. 1998, arXiv:quant-ph/9807006.
- [GPP⁺15] L. C. G. Govia, E. J. Pritchett, B. L. T. Plourde, M. G. Vavilov, R. McDermott, and F. K. Wilhelm. Scalable two- and four-qubit parity measurement with a threshold photon counter. *Phys. Rev. A*, 92(2):022335, 2015.
- [GPX⁺14] L. C. G. Govia, E. J. Pritchett, C. Xu, B. L. T. Plourde, M. G. Vavilov, F. K. Wilhelm, and R. McDermott. High-fidelity qubit measurement with a microwave-photon counter. *Phys. Rev. A*, 90(6):062307, 2014.
- [Gre15] D. Greenbaum. Introduction to quantum gate set tomography, 2015, arXiv:1509.02921.
- [Gri11] W. P. Grice. Arbitrarily complete bell-state measurement using only linear optical elements. *Phys. Rev. A*, 84(4):042331, 2011.
- [Gro96] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings, 28th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 212–219, 1996.
- [GSVB13] M. Gutiérrez, L. Svec, A. Vargo, and K. R. Brown. Approximation of realistic errors by clifford channels and pauli measurements. *Phys. Rev. A*, 87(3), 2013.
- [GTL⁺16] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland. High-fidelity universal gate set for ⁹Be⁺ ion qubits. *Phys. Rev. Lett.*, 117(6):060505, 2016.
- [GZ13] M. R. Geller and Z. Zhou. Efficient error models for fault-tolerant architectures and the pauli twirling approximation. *Phys. Rev. A*, 88(1):012314, 2013.
- [HAB⁺14] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas. High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. *Phys. Rev. Lett.*, 113(22):220501, 2014.
- [Har13] S. Haroche. Nobel lecture: Controlling photons in a box and exploring the quantum to classical boundary. *Rev. Mod. Phys.*, 85(3):1083–1102, 2013.
- [HBD⁺15] B. Hensen, H. Bernien, A. E. Dreau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellan, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson. Loophole-free bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575):682–686, 2015.

- [HBR⁺16] A. Holleccek, O. Barter, A. Rubenok, J. Dilley, P. B. R. Nisbet-Jones, G. Langfahl-Klabes, G. D. Marshall, C. Sparrow, J. L. O'Brien, K. Poullos, A. Kuhn, and J. C. F. Matthews. Quantum logic with cavity photons from single atoms. *Phys. Rev. Lett.*, 117(2):023602, 2016.
- [HCS⁺06] D. Hong, D. Chang, J. Sung, S. Lee, S. Hong, J. Lee, D. Moon, and S. Chee. A New Dedicated 256-Bit Hash Function: FORK-256. In *Fast Software Encryption 2006 (FSE 2006)*, volume 4047 of *Lecture Notes in Computer Science*, pages 195–209. Springer, 2006.
- [HCW⁺12] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, K. S. McKay, D. P. Pappas, D. Leibfried, and D. J. Wineland. 100-fold reduction of electric-field noise in an ion trap cleaned with in situ argon-ion-beam bombardment. *Phys. Rev. Lett.*, 109(10):103001, 2012.
- [HDF18] E. Huang, A.C. Doherty, and S. Flammia, Performance of quantum error correction with coherent errors, 2019, arXiv:1805.08227
- [HDM⁺16] N. C. Harris, B. Darius, P. Mihir, G. R. Steinbrecher, M. Jacob, P. Mihika, B.-J. Tom, H. Michael, and E. Dirk. *nanoph*, volume 5, chapter Large-scale quantum photonic circuits in silicon, page 456. 2017 2016.
- [Hei03] T. Heinzel. *Mesoscopic Electronics in Solid State Nanostructures*. Wiley, 2003.
- [HF18] R. Harper and S. Flammia. Fault tolerance in the IBM Q Experience. ArXiv:1806.02359.
- [HFC⁺03] T. Hayashi, T. Fujisawa, H. D. Cheong, Y. H. Jeong, and Y. Hirayama. Coherent manipulation of electronic states in a double quantum dot. *Phys. Rev. Lett.*, 91(22):226804, 2003.
- [HFDvM12] C. Horsman, A.G. Fowler, S. Devitt, and R. van Meter. Surface code quantum computing by lattice surgery. *New. J. Phys.*, 14(12):123011, 2012.
- [HFJ⁺17] T. Hensgens, T. Fujita, L. Janssen, X. Li, C. J. Van Diepen, C. Reichl, W. Wegscheider, S. Das Sarma, and L. M. K. Vandersypen. Quantum simulation of a fermi-hubbard model using a semiconductor quantum dot array. *Nature*, 548(7665):70–73, August 2017.
- [HG14] S.-Y. Huang and H.-S. Goan. Optimal control for fast and high-fidelity quantum gates in coupled superconducting flux qubits. *Phys. Rev. A*, 90(1):012318, Jul 2014.
- [HGL⁺98] M. Haake, B. M. Goodson, D. D. Laws, E. Brunner, M. C. Cyrier, R. H. Havlin, and A. Pines. Nmr of supercritical laser-polarized xenon. *Chem. Phys. Lett.*, 292(4–6):686 – 690, 1998.
- [HHJ⁺09] J. P. Home, D. Hanneke, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland. Complete methods set for scalable ion trap quantum information processing. *Science*, 325(5945):1227–1230, 2009.
- [HHL09] A. W. Harrow, A. Hassidim, and S. Lloyd. “Quantum algorithm for linear systems of equations”. In: *Physical review letters* 103.15 (2009), p. 150502.
- [HJNRS20a] T. Häner, S. Jaques, M. Naehrig, M. Roetteler, and M. Soeken. *Improved Quantum Circuits for Elliptic Curve Discrete Logarithms*. In *Post-Quantum Cryptography – PQCrypto 2020*, Springer LNCS vol. 12100, pp. 425–444, 2020.
- [HKD⁺09] A. A. Houck, J. Koch, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Life after charge noise: recent results with transmon qubits. *Quantum Inf. Process.*, 8(2):105–115, 2009.
- [HKEG19] D. Hangleiter, M. Kliesch, J. Eisert, and C. Gogolin, Sample Complexity of Device-Independently Certified “Quantum Supremacy”, *Phys. Rev. Lett.* 122, 210502
- [HLBH11] M. D. Hughes, B. Lekitsch, J. A. Broersma, and W. K. Hensinger. Microfabricated ion traps. *Contemp. Phys.*, 52(6):505–529, 2011.
- [HMB⁺15] T. Huang, A. Malmros, J. Bergsten, S. Gustafsson, O. Axelsson, M. Thorsell, and N. Rorsman. Suppression of dispersive effects in algan/gan high-electron-mobility transistors using bilayer sinx grown by low pressure chemical vapor deposition. *IEEE Electron Device Lett.*, 36(6):537–539, 2015.

- [HPZ⁺17] G. Higgins, F. Pokorny, C. Zhang, Q. Bodart, and M. Hennrich. Coherent control of a single trapped Rydberg ion, 2017, arXiv:1708.06387.
- [HR06] S. Haroche and J. Raimond. *Exploring the Quantum: Atoms, Cavities, and Photons*. Oxford Graduate Texts. OUP Oxford, 2006.
- [HRC02] A. W. Harrow, B. Recht, and I. L. Chuang. Efficient discrete approximations of quantum gates. *J. Math. Phys.*, 43(9):4445–4451, 2002.
- [HRO⁺16] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf. Implementing a universal gate set on a logical qubit encoded in an oscillator, 2016, arXiv:1608.02430.
- [HRP⁺06] T. Hime, P. A. Reichardt, B. L. T. Plourde, T. L. Robertson, C.-E. Wu, A. V. Ustinov, and J. Clarke. Solid-state qubits with current-controlled coupling. *Science*, 314(5804):1427–1429, 2006.
- [HJNRS20] T. Häner, S. Jaques, M. Naehrig, M. Roetteler, and M. Soeken. *Improved Quantum Circuits for Elliptic Curve Discrete Logarithms*, Cryptology ePrint Archive: Report 2020/077, 2020.
- [HRS17] T. Häner, M. Roetteler, and K. M. Svore. Factoring using $2n + 2$ qubits with Toffoli-based modular multiplication, *Quantum Information & Computation* 17 (7&8), pp. 673–684, 2017.
- [HS17a] T. Häner and D. S. Steiger. 0.5 petabyte simulation of a 45-qubit quantum circuit, 2017, Proc. of the International Conference for High Performance Computing, Networking, Storage and Analysis SC 2017: 33: 1–33: 10, 2017.
- [HSA⁺16] T. P. Harty, M. A. Sepiol, D. T. C. Allcock, C. J. Ballance, J. E. Tarlton, and D. M. Lucas. High-fidelity trapped-ion quantum logic using near-field microwaves. *Phys. Rev. Lett.*, 117(14):140501, 2016.
- [HSG⁺07] A. A. Houck, D. I. Schuster, J. M. Gambetta, J. A. Schreier, B. R. Johnson, J. M. Chow, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Generating single microwave photons in a circuit. *Nature*, 449(7160):328–331, Sep 2007.
- [HSM06] J. Harrison, M. Sellars, and N. Manson. Measurement of the optically induced spin polarisation of n-v centres in diamond. *Diamond Relat. Mater.*, 15(4–8):586 – 588, 2006.
- [HVS⁺11] M. Hatridge, R. Vijay, D. H. Slichter, J. Clarke, and I. Siddiqi. Dispersive magnetometry with a quantum limited squid parametric amplifier. *Phys. Rev. B*, 83(13):134501, 2011.
- [HYC⁺18] W. Huang, C. H. Yang, K. W. Chan, T. Tanttu, B. Hensen, R.C.C. Leon, M. A. Fogarty, J.C.C. Hwang, F. E. Hudson, K. M. Itoh, A. Morello, A. Laucht, and A. S. Dzurak, Fidelity benchmarks for two-qubit gates in silicon. ArXiv:1805.05027.
- [IBM18] IBM. Quantum Information Science Kit, 2018. <https://qiskit.org/>.
- [IC]⁺05] G. Ithier, E. Collin, P. Joyez, P. J. Meeson, D. Vion, D. Esteve, F. Chiarello, A. Shnirman, Y. Makhlin, J. Schrieffer, and G. Schön. Decoherence in a superconducting quantum bit circuit. *Phys. Rev. B*, 72(13):134519, 2005.
- [IFI⁺02] L. B. Ioffe, M. V. Feigel'man, A. Ioselevich, D. Ivanov, M. Troyer, and G. Blatter. Topologically protected quantum bits using josephson junction arrays. *Nature*, 415(6871):503–506, 2002.
- [IIS⁺13] V. V. Ivanov, J. A. Isaacs, M. Saffman, S. Kemme, A. Ellis, G. Brady, J. Wendt, G. W. Biedermann, and S. Samora. Atom trapping in a bottle beam created by a diffractive optical element, 2013, arXiv:1305.5309, <https://arxiv.org/abs/1305.5309>.
- [Il[']16] E. V. Il'ichev. A microwave photon detector. *Phys. Solid State*, 58(11):2160–2164, 2016.
- [JAG⁺11] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose. Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, 2011.

- [JBM⁺18a] S. Jiang, K.A. Britt, A.J. McCaskey, T.S. Humble, and Sabre Kais. Quantum Annealing for Prime Factorization, 2018, arXiv:1804.02733.
- [JdSR⁺15] B. R. Johnson, M. P. da Silva, C. A. Ryan, S. Kimmel, J. M. Chow, and T. A. Ohki. Demonstration of robust quantum gate tomography via randomized benchmarking. *New J. Phys.*, 17(11):113019, 2015.
- [JF08] S. P. Jordan and E. Farhi. Perturbative gadgets at arbitrary orders. *Phys. Rev. A*, 77(6):062329, 2008.
- [JFV⁺11] J. Jang, D. G. Ferguson, V. Vakaryuk, R. Budakian, S. B. Chung, P. M. Goldbart, and Y. Maeno. Observation of half-height magnetization steps in sr2ruo4. *Science*, 331(6014):186–188, 2011.
- [JGBW⁺16] G. Jacob, K. Groot-Berning, S. Wolf, S. Ulm, L. Couturier, S. T. Dawkins, U. G. Poschinger, F. Schmidt-Kaler, and K. Singer. Transmission microscopy with nanometer resolution using a deterministic single ion source. *Phys. Rev. Lett.*, 117(4):043001, 2016.
- [JGLM19] D. Joseph, A. Ghionis, C. Ling, F. Mintert, Not-so-adiabatic quantum computation for the shortest vector problem, 23rd October, 2019
- [JGP⁺04] F. Jelezko, T. Gaebel, I. Popa, M. Domhan, A. Gruber, and J. Wrachtrup. Observation of coherent oscillation of a single nuclear spin and realization of a two-qubit conditional quantum gate. *Phys. Rev. Lett.*, 93(13):130501, 2004.
- [JHK⁺16] Y.-Y. Jau, A. M. Hankin, T. Keating, I. H. Deutsch, and G. W. Biedermann. Entangling atomic spins with a rydberg-dressed spin-flip blockade. *Nat. Phys.*, 12(1):71–74, 2016.
- [JKS⁺15] X. Y. Jin, A. Kamal, A. P. Sears, T. Gudmundsen, D. Hover, J. Miloshi, R. Slattery, F. Yan, J. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver. Thermal and residual excited-state population in a 3d transmon qubit. *Phys. Rev. Lett.*, 114(24):240501, 2015.
- [JLH⁺14] P. Jurcevic, B. P. Lanyon, P. Hauke, C. Hempel, P. Zoller, R. Blatt, and C. F. Roos. Quasiparticle engineering and entanglement propagation in a quantum many-body system. *Nature*, 511(7508):202–205, 2014.
- [JMB+18a] Jiang, K.A. Britt, A.J. McCaskey, T.S. Humble, and Sabre Kais. Quantum Annealing for Prime Factorization, Scientific Reports vol. 8, Article no. 17667, 2018.
- [JNRV20] S. Jaques, M. Naehrig, M. Roetteler, and F. Virdia. *Implementing Grover oracles for quantum key search on AES and LowMC*. In Advances in Cryptology - EUROCRYPT 2020, Springer LNCS vol. 12106, pp. 280-310, 2020.
- [JOB16] T. Jochym-O'Connor and S. D. Bartlett. Stacked codes: Universal fault-tolerant quantum computation in a two-dimensional layout. *Phys. Rev. A*, 93(2):022323, 2016.
- [Joh28] J. B. Johnson. Thermal agitation of electricity in conductors. *Phys. Rev.*, 32(1):97–109, 1928.
- [Joh17] A. M. Johnston. Shor's Algorithm and Factoring: Don't Throw Away the Odd Orders. Cryptology ePrint Archive: Report 2017/083, 2017, <http://eprint.iacr.org/2017/083>.
- [Jon13] C. Jones. Low-overhead constructions for the fault-tolerant toffoli gate. *Phys. Rev. A*, 87:022328, Feb 2013.
- [JSM⁺14] E. Jeffrey, D. Sank, J. Y. Mutus, T. C. White, J. Kelly, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Megrant, P. J. J. O'Malley, C. Neill, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis. Fast accurate state measurement with superconducting qubits. *Phys. Rev. Lett.*, 112(19):190504, 2014.
- [Kan98] B. Kane. A silicon-based nuclear spin quantum computer. *Nature*, 393(6681):133–137, 1998.
- [Kas50] A. Kastler. Applications of polarimetry to infra-red and micro-wave spectroscopy. *Nature*, 166(4211):113–113, 1950.
- [Kas67] A. Kastler. Optical methods for studying hertzian resonances. *Science*, 158(3798):214–221, 1967.

- [KBF⁺15] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis. State preservation by repetitive error detection in a superconducting quantum circuit. *Nature*, 519(7541):66–69, 2015.
- [KBLW01] J. Kempe, D. Bacon, D. A. Lidar, and K. B. Whaley. Theory of decoherence-free fault-tolerant universal quantum computation. *Phys. Rev. A*, 63(4):042307, 2001.
- [KC11] D. Kinion and J. Clarke. Superconducting quantum interference device as a near-quantum-limited amplifier for the axion dark-matter experiment. *Appl. Phys. Lett.*, 98(20):202503, 2011.
- [KCK⁺10] K. Kim, M.-S. Chang, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G.-D. Lin, L.-M. Duan, and C. Monroe. Quantum simulation of frustrated ising spins with trapped ions. *Nature*, 465(7298):590–593, 2010.
- [KdSR⁺14] S. Kimmel, M. P. da Silva, C. A. Ryan, B. R. Johnson, and T. Ohki. Robust extraction of tomographic information via randomized benchmarking. *Phys. Rev. X*, 4(1), 2014.
- [KGPUK16] F. Kaneda, K. Garay-Palmett, A. B. U'Ren, and P. G. Kwiat. Heralded single-photon source utilizing highly nondegenerate, spectrally factorable spontaneous parametric downconversion. *Opt. Express*, 24(10):10733–10747, May 2016.
- [KHDS01] B. Koiller, X. Hu, and S. Das Sarma. Exchange in silicon-based quantum computer architecture. *Phys. Rev. Lett.*, 88(2):027903, 2001.
- [KHJ18] P. Kim, D. Han, and K. C. Jeong. Time-space complexity of quantum search algorithms in symmetric cryptanalysis: applying to AES and SHA-2. *Quantum Information Processing* 17:339, 2018.
- [Kie01] D. Kielpinski. A decoherence-free quantum memory using trapped ions. *Science*, 291(5506):1013–1015, 2001.
- [Kim08] H. J. Kimble. The quantum internet. *Nature*, 453(7198):1023–1030, 2008.
- [Kit97a] A. Y. Kitaev. Quantum error correction with imperfect gates. In O. Hirota, A. Holevo, and C. Caves, editors, *Proceedings of the Third International Conference on Quantum Communication, Computing and Measurement*. Plenum Press, New York, 1997.
- [Kit97b] A. Y. Kitaev. Quantum computations: algorithms and error correction. *Russ. Math. Surv.*, 52(6):1191–1249, 1997.
- [Kit06] A. Kitaev. Protected qubit based on a superconducting current mirror, 2006, arXiv:cond-mat/0609441.
- [KJS⁺16] E. Kawakami, T. Jullien, P. Scarlino, D. R. Ward, D. E. Savage, M. G. Lagally, V. V. Dobrovitski, M. Friesen, S. N. Coppersmith, M. A. Eriksson, , and L. M. K. Vandersypen. Gate fidelity and coherence of an electron spin in an si/sige quantum dot with micromagnet. *Proc. Natl. Acad. Sci. U.S.A.*, 113(42):11738–11743, 2016.
- [KKL⁺17] T. Karzig, C. Knapp, R. M. Lutchyn, P. Bonderson, M. B. Hastings, C. Nayak, J. Alicea, K. Flensberg, S. Plugge, Y. Oreg, C. M. Marcus, and M. H. Freedman. Scalable designs for quasiparticle-poisoning-protected topological quantum computation with majorana zero modes. *Phys. Rev. B*, 95(23):235305, 2017.
- [KKR06] J. Kempe, A. Kitaev, and O. Regev. The complexity of the local Hamiltonian problem. *SIAM J. Comput.*, 35(5):1070–1097, 2006.
- [KLC⁺15] J. Kerckhoff, K. Lalumière, B. J. Chapman, A. Blais, and K. W. Lehnert. On-chip superconducting microwave circulator from synthetic rotation. *Phys. Rev. Applied*, 4(3):034002, 2015.
- [KLFF⁺15] A. M. Kaufman, B. J. Lester, M. Foss-Feig, M. L. Wall, A. M. Rey, and C. A. Regal. Entangling two transportable neutral atoms via local spin exchange. *Nature*, 527(7577):208–211, 2015.

- [KLH⁺04] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther. Continuous generation of single photons with controlled waveform in an ion-trap cavity system. *Nature*, 431(7012):1075–1078, Oct 2004.
- [KLHM14] R. Kalra, A. Laucht, C. D. Hill, and A. Morello. Robust two-qubit gates for donors in silicon controlled by hyperfine interactions. *Phys. Rev. X*, 4(2):021044, 2014.
- [Kli13] V. Kliuchnikov. Synthesis of unitaries with Clifford+T circuits. arXiv:1306.3200, 2013.
- [KLLNP16] M. Kaplan, G. Leurent, A. Leverrier, and M. Naya-Plasencia. Quantum differential and linear cryptanalysis. *IACR Transactions on Symmetric Cryptology*, 2016(1), 2016.
- [KLM01] E. Knill, R. Laflamme, and G. J. Milburn. A scheme for efficient quantum computation with linear optics. *Nature*, 409(6816):46–52, 2001.
- [KLNP16] M. Kaplan, A. Leverrier, and M. Naya-Plasencia. Breaking symmetric cryptosystems using quantum period finding. In *Advances in Cryptology – CRYPTO 2016*, volume 9815 of *Lecture Notes in Computer Science*, pages 207–237. Springer, 2016.
- [KLR⁺08] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Randomized benchmarking of quantum gates. *Phys. Rev. A*, 77(1):012307, 2008.
- [KMBD06] C. Kruszynska, A. Miyake, H. J. Briegel, and W. Dür. Entanglement purification protocols for all graph states. *Phys. Rev. A*, 74(5):052316, 2006.
- [KMW02] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417(6890):709–711, 2002.
- [Kni05] E. Knill. Quantum computing with realistically noisy devices. *Nature*, 434(7029):39–44, 2005.
- [Kor08] A. N. Korotkov. Quantum efficiency of binary-outcome detectors of solid-state qubits. *Phys. Rev. B*, 78(17):174512, 2008.
- [KRK⁺05] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser. Optimal control of coupled spin dynamics: design of {NMR} pulse sequences by gradient ascent algorithms. *J. Magn. Reson.*, 172(2):296 – 305, 2005.
- [KRS⁺17] H. Kaufmann, T. Ruster, C. T. Schmiegelow, M. A. Luda, V. Kaushal, J. Schulz, D. von Lindenfels, F. Schmidt-Kaler, and U. G. Poschinger. Scalable creation of long-lived multipartite entanglement. *Phys. Rev. Lett.*, 119:150503, Oct 2017.
- [KSB⁺16] P. Kumar, S. Sendelbach, M. A. Beck, J. W. Freeland, Z. Wang, H. Wang, C. C. Yu, R. Q. Wu, D. P. Pappas, and R. McDermott. Origin and reduction of $1/f$ magnetic flux noise in superconducting devices. *Phys. Rev. Applied*, 6(4):041001, 2016.
- [KSS⁺14] D. Kim, Z. Shi, C. B. Simmons, D. R. Ward, J. R. Prance, T. S. Koh, J. K. Gamble, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, and M. A. Eriksson. Quantum control and process tomography of a semiconductor quantum dot hybrid qubit. *Nature*, 511(7507):70–74, July 2014.
- [KXB⁺16] D. Korenkevych, Y. Xue, Z. Bian, F. Chudak, W. G. Macready, J. Rolfe, and E. Andriyash. Benchmarking quantum hardware for training of fully visible boltzmann machines, 2016, arXiv:1611.04528.
- [KYG⁺07] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Charge-insensitive qubit design derived from the Cooper pair box. *Phys. Rev. A*, 76:042319, 2007.
- [KYP15] A. Kubica, B. Yoshida, and F. Pastawski. Unfolding the color code. *New J. Phys.*, 17(8):083026, 2015.
- [KZ04] P. Kaye and C. Zalka. Optimized quantum implementation of elliptic curve arithmetic over binary fields, 2004.
- [Lan95] R. Landauer. Is quantum mechanics useful? *Phil. Trans. R. Soc. A*, 353(1703):367–376, 1995.

- [LB13] D. A. Lidar and T. A. Brun. *Quantum Error Correction*, chapter Introduction to decoherence and noise in open quantum systems, pages 3–45. Cambridge University Press, 2013.
- [LBH⁺16] J. C. Loredo, M. A. Broome, P. Hilaire, O. Gazzano, I. Sagnes, A. Lemaitre, M. P. Almeida, P. Senellart, and A. G. White. Bosonsampling with single-photon fock states from a bright solid-state source, 2016, arXiv:1603.00054, PRL to appear, <https://arxiv.org/abs/1603.00054v1>.
- [LBM⁺16] J. Lisenfeld, A. Bilmes, S. Matityahu, S. Zanker, M. Marthaler, M. Schechter, G. Schön, G.n, A. Shnirman, G. Weiss, and A. V. Ustinov. Decoherence spectroscopy with individual two-level tunneling defects. *Sci. Rep.*, 6:23786–, 2016.
- [LCYY00] D. W. Leung, I. L. Chuang, F. Yamaguchi, and Y. Yamamoto. Efficient implementation of coupled logic gates for quantum computation. *Phys. Rev. A*, 61(4):042310, 2000.
- [LD98] D. Loss and D. P. DiVincenzo. Quantum computation with quantum dots. *Phys. Rev. A*, 57(1):120–126, 1998.
- [LDP⁺17] X. Lin, X. Dai, C. Pu, Y. Deng, Y. Niu, L. Tong, W. Fang, Y. Jin, and X. Peng. Electrically-driven single-photon sources based on colloidal quantum dots with near-optimal antibunching at room temperature. *Nature Communications*, 8(1):1132, 2017.
- [Leg80] A. J. Leggett. Macroscopic quantum systems and the quantum theory of measurement. *Progr. Theor. Phys.*, 69:80, 1980.
- [Lev01] M. Levitt. *Spin Dynamics: Basics of Nuclear Magnetic Resonance*. Spin Dynamics: Basics of Nuclear Magnetic Resonance. Wiley, 2001.
- [LGA⁺08] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibbrandt, K. R. Brown, and I. L. Chuang. Suppression of heating rates in cryogenic surface-electrode ion traps. *Phys. Rev. Lett.*, 100(1):013001, 2008.
- [LH10] M. Lieb, M. J. Hartmann, Bose–Hubbard dynamics of polaritons in a chain of circuit quantum electrodynamics cavities, *New Journal of Physics*, Volume 12, September 2010.
- [LHA⁺08] E. Lucero, M. Hofheinz, M. Ansmann, R. C. Bialczak, N. Katz, M. Neeley, A. D. O’Connell, H. Wang, A. N. Cleland, and J. M. Martinis. High-fidelity gates in a single josephson qubit. *Phys. Rev. Lett.*, 100(24):247001, 2008.
- [LHZ15] W. Lechner, P. Hauke, and P. Zoller. A quantum annealing architecture with all-to-all connectivity from local interactions. *Sci. Adv.*, 1(9):e1500838–e1500838, 2015.
- [Li15] Y. Li. A magic state’s fidelity can be superior to the operations that created it. *New J. Phys.*, 17(2):023037, 2015.
- [Lin76] G. Lindblad. On the generators of quantum dynamical semigroups. *Commun. Math. Phys.*, 48:119–130, 1976.
- [LJH⁺13] E. J. H. Lee, X. Jiang, M. Houzet, R. Aguado, C. M. Lieber, and S. D. Franceschi. Spin-resolved andreev levels and parity crossings in hybrid superconductor–semiconductor nanostructures. *Nat. Nanotechnol.*, 9(1):79–84, 2013.
- [LKB⁺10] E. Lucero, J. Kelly, R. C. Bialczak, M. Lenander, M. Mariantoni, M. Neeley, A. D. O’Connell, D. Sank, H. Wang, M. Weides, J. Wenner, T. Yamamoto, A. N. Cleland, and J. M. Martinis. Reduced phase error through optimized control of a superconducting qubit. *Phys. Rev. A*, 82(4):042339, 2010.
- [LKEH17] T. Lanting, A. D. King, B. Evert, and E. Hoskinson. Experimental demonstration of perturbative anticrossing mitigation using nonuniform driver hamiltonians. *Phys. Rev. A*, 96:042322, Oct 2017.
- [LKV⁺13] Z. Leghtas, G. Kirchmair, B. Vlastakis, R. J. Schoelkopf, M. H. Devoret, and M. Mirrahimi. Hardware-efficient autonomous quantum memory protection. *Phys. Rev. Lett.*, 111(12):120501, 2013.

- [LL01] M. N. Leuenberger and D. Loss. Quantum computing in molecular magnets. *Nature*, 410(6830):789–793, 2001.
- [LLF⁺18] P.H. Leung, K.A. Landsma, C. Figgatt, N.M. Linke, C. Monroe, and K.R. Brown. Robust 2-Qubit Gates in a Linear Ion Crystal Using a Frequency-Modulated Driving Force. *Phys. Rev. Lett.*, 120(02):02001, 2018.
- [LmcHM05] A. Lupaşcu, C. J. P. M. Harmans, and J. E. Mooij. Quantum state detection of a superconducting flux qubit using a dc-SQUID in the inductive mode. *Phys. Rev. B*, 71(18):184506, May 2005.
- [LMR⁺17] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe. Experimental comparison of two quantum computing architectures. 2017, arXiv:1702.01852.
- [LPK⁺15] T. W. Larsen, K. D. Petersson, F. Kuemmeth, T. S. Jespersen, P. Krogstrup, J. Nygård, and C. M. Marcus. Semiconductor-nanowire-based superconducting qubit. *Phys. Rev. Lett.*, 115(12):127001, 2015.
- [LPS⁺14] T. Lanting, A. J. Przybysz, A. Y. Smirnov, F. M. Spedalieri, M. H. Amin, A. J. Berkley, R. Harris, F. Altomare, S. Boixo, P. Bunyk, N. Dickson, C. Enderud, J. P. Hilton, E. Hoskinson, M. W. Johnson, E. Ladizinsky, N. Ladizinsky, R. Neufeld, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, S. Uchaikin, A. B. Wilson, and G. Rose. Entanglement in a quantum annealing processor. *Phys. Rev. X*, 4(2):021041, 2014.
- [LPS20] B. Langenberg, H. Pham, and R. Steinwandt. *Reducing the Cost of Implementing AES as a Quantum Circuit*. IEEE Transactions on Quantum Engineering, 2020. Preprint available as Cryptology ePrint Archive: Report 2019/854.
- [LSDS10] R. M. Lutchyn, J. D. Sau, and S. Das Sarma. Majorana fermions and a topological phase transition in semiconductor-superconductor heterostructures. *Phys. Rev. Lett.*, 105(7):077001, 2010.
- [LTD⁺10] E. A. Laird, J. M. Taylor, D. P. DiVincenzo, C. M. Marcus, M. P. Hanson, and A. C. Gossard. Coherent spin manipulation in an exchange-only qubit. *Phys. Rev. B*, 82(7):075403, 2010.
- [LW03] D. A. Lidar and K. B. Whaley. *Irreversible Quantum Dynamics*, volume 622 of *Springer Lecture Notes in Physics*, chapter Decoherence-Free Subspaces and Subsystems, pages 83–120. Springer, 2003.
- [LWF⁺17] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger. Blueprint for a microwave trapped ion quantum computer. *Sci. Adv.*, 3(2):e1601540, 2017.
- [LZG⁺07] C.-Y. Lu, X.-Q. Zhou, O. Gühne, W.-B. Gao, J. Zhang, Z.-S. Yuan, A. Goebel, T. Yang, and J.-W. Pan. Experimental entanglement of six photons in graph states. 3:91–95, Jan 2007.
- [Mak02] Y. Makhlin. Nonlocal properties of two-qubit gates and mixed states, and the optimization of quantum computations. *Quantum Inf. Process.*, 1(4):243–252, 2002.
- [Mar15] J. M. Martinis. Qubit metrology for building a fault-tolerant quantum computer, 2015, arxiv:1510.01406v1, <http://arxiv.org/abs/1510.01406v1>.
- [Mar17] J. M. Martinis, 2017, <http://www.acm.org/articles/people-of-acm/2017/john-martinis>.
- [Mau07] P. Maunz. High-fidelity operations in microfabricated surface ion traps. 2007.
- [MBB⁺16] T. Meany, D. N. Biggerstaff, M. A. Broome, A. Fedrizzi, M. Delanty, M. Steel, A. Gilchrist, G. D. Marshall, A. G. White, and M. J. Withford. Engineering integrated photonics for heralded quantum gates. *Sci. Rep.*, 6:25126, 2016.
- [MCP⁺17] A. C. Mahoney, J. I. Colless, S. J. Pauka, J. M. Hornibrook, J. D. Watson, G. C. Gardner, M. J. Manfra, A. C. Doherty, and D. J. Reilly. On-chip microwave quantum hall circulator. *Phys. Rev. X*, 7(1):011007, 2017.
- [MdARC⁺14] R. Medeiros de Araújo, J. Roslund, Y. Cai, G. Ferrini, C. Fabre, and N. Treps. Full characterization of a highly multimode entangled state embedded in an optical frequency comb using pulse shaping. *Phys. Rev. A*, 89:053828, May 2014.
- [McKWS⁺17] D.C. McKay, C.J.W Wood, S. Sheldon, J.M. Chow, and J.M. Gambetta. Efficient Z gates for quantum computing. *Phys. Rev. A*. 96(2):022330, 2017.

- [MDL⁺14] J. T. Muhonen, J. P. Dehollain, A. Laucht, F. E. Hudson, R. Kalra, T. Sekiguchi, K. M. Itoh, D. N. Jamieson, J. C. McCallum, A. S. Dzurak, and A. Morello. Storing quantum information for 30 seconds in a nanoelectronic device. *Nat. Nano.*, 9(12):986–991, 2014.
- [Men14] N. C. Menicucci. Fault-tolerant measurement-based quantum computing with continuous-variable cluster states. *Phys. Rev. Lett.*, 112(12):120504, 2014.
- [MFZP07] N. C. Menicucci, S. T. Flammia, H. Zaidi, and O. Pfister. Ultracompact generation of continuous-variable cluster states. *Phys. Rev. A*, 76:010302, Jul 2007.
- [MGE11] E. Magesan, J. M. Gambetta, and J. Emerson. Scalable and robust randomized benchmarking of quantum processes. *Phys. Rev. Lett.*, 106(18):180504, 2011.
- [MGE12] E. Magesan, J. M. Gambetta, and J. Emerson. Characterizing quantum gates via randomized benchmarking. *Phys. Rev. A*, 85(4):042311, 2012.
- [MGJ⁺12] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. da Silva, G. A. Keefe, M. B. Rothwell, T. A. Ohki, M. B. Ketchen, and M. Steffen. Efficient measurement of quantum gate error by interleaved randomized benchmarking. *Phys. Rev. Lett.*, 109(8):080505, 2012.
- [MGRW09] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm. Simple pulses for elimination of leakage in weakly nonlinear qubits. *Phys. Rev. Lett.*, 103(11):110501, 2009.
- [MGS⁺13] S. T. Merkel, J. M. Gambetta, J. A. Smolin, S. Poletto, A. D. Córcoles, B. R. Johnson, C. A. Ryan, and M. Steffen. Self-consistent quantum process tomography. *Phys. Rev. A*, 87(6), 2013.
- [MGW⁺03] O. Mandel, M. Greiner, A. Widera, T. Rom, T. W. Hansch, and I. Bloch. Controlled collisions for multi-particle entanglement of optically trapped atoms. *Nature*, 425(6961):937–940, 2003.
- [Mic18] Microsoft. Quantum Development Kit, 2018, <https://www.microsoft.com/en-us/quantum/development-kit>.
- [Mir16] M. Mirrahimi. Cat-qubits for quantum computation. *Comptes Rendus Physique*, 17(7):778 – 787, 2016.
- [MK13] C. Monroe and J. Kim. Scaling the ion trap quantum processor. *Science*, 339(6124):1164–1169, 2013.
- [MKC⁺15] E. Mount, C. Kabytayev, S. Crain, R. Harper, S.-Y. Baek, G. Vrijsen, S. T. Flammia, K. R. Brown, P. Maunz, and J. Kim. Error compensation of single-qubit gates in a surface-electrode ion trap using composite pulses. *Phys. Rev. A*, 92(6):060301, 2015.
- [MKT⁺00] C. J. Myatt, B. E. King, Q. A. Turchette, C. A. Sackett, D. Kielpinski, W. M. Itano, C. Monroe, and D. J. Wineland. Decoherence of quantum superpositions through coupling to engineered reservoirs. *Nature*, 403(6767):269–273, 2000.
- [MLA⁺14] M. Mirrahimi, Z. Leghtas, V. V. Albert, S. Touzard, R. J. Schoelkopf, L. Jiang, and M. H. Devoret. Dynamically protected cat-qubits: a new paradigm for universal quantum computation. *New J. Phys.*, 16(4):045014, 2014.
- [MLS⁺15] J. Muhonen, A. Laucht, S. Simmons, J. Dehollain, R. Kalra, F. Hudson, S. Freer, K. Itoh, D. Jamieson, J. McCallum, A. Dzurak, and A. Morello. Quantifying the quantum gate fidelity of single-atom spin qubits in silicon by randomized benchmarking. *J. Phys.: Condens. Matter*, 27(15):154205, 2015.
- [MLX⁺15] K. M. Maller, M. T. Lichtman, T. Xia, Y. Sun, M. J. Piotrowicz, A. W. Carr, L. Isenhower, and M. Saffman. Rydberg-blockade controlled-not gate and entanglement in a two-dimensional array of neutral-atom qubits. *Phys. Rev. A*, 92:022336, Aug 2015.
- [MMCP09] D. Maslov, J. Mathew, D. Cheung, and D. K. Pradhan. An $O(m^2)$ -depth quantum algorithm for the elliptic curve discrete logarithm problem over $GF(2^m)$. *Quantum Inf. Comput.*, 9(7):610–621, 2009.

- [MMK⁺95] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. Demonstration of a fundamental quantum logic gate. *Phys. Rev. Lett.*, 75(25):4714–4717, 1995.
- [MMK⁺12] N. Mizuochi, T. Makino, H. Kato, D. Takeuchi, M. Ogura, H. Okushi, M. Nothaft, P. Neumann, A. Gali, F. Jelezko, J. Wrachtrup, and S. Yamasaki. Electrically driven single-photon source at room temperature in diamond. 6:299 EP –, Apr 2012.
- [MMN⁺17a] F. K. Malinowski, F. Martins, P. D. Nissen, E. Barnes, C. Łukasz, M. S. Rudner, S. Fallahi, G. C. Gardner, M. J. Manfra, C. M. Marcus, and F. Kuemmeth. Notch filtering the nuclear environment of a spin qubit. *Nat. Nanotechnol.*, 12(1):16–20, 2017.
- [MMN⁺17b] F. K. Malinowski, F. Martins, P. D. Nissen, S. Fallahi, G. C. Gardner, M. J. Manfra, C. M. Marcus, and F. Kuemmeth. Symmetric operation of the resonant exchange qubit. *Phys. Rev. B*, 96:045443, Jul 2017.
- [MOH⁺15] C. Macklin, K. O’Brien, D. Hover, M. E. Schwartz, V. Bolkhovskiy, X. Zhang, W. D. Oliver, and I. Siddiqi. A near-quantum-limited josephson traveling-wave parametric amplifier. *Science*, 2015.
- [MOL⁺99] J. E. Mooij, T. P. Orlando, L. Levitov, L. Tian, C. H. van der Wal, and S. Lloyd. Josephson persistent-current qubit. *Science*, 285(5430):1036–1039, 1999.
- [MPB⁺15] R. Marx, N. Pomplun, W. Bermel, H. Zeiger, F. Engelke, A. F. Fahmy, and S. J. Glaser. Engineering of an all-heteronuclear 5-qubit nmr quantum computer. *Magn. Reson. Chem.*, 53(6):442–447, 2015.
- [MPZ⁺10] A. Morello, J. J. Pla, F. A. Zwanenburg, K. W. Chan, K. Y. Tan, H. Huebl, M. Mottonen, C. D. Nugroho, C. Yang, J. A. van Donkelaar, A. D. C. Alves, D. N. Jamieson, C. C. Escott, L. C. L. Hollenberg, R. G. Clark, and A. S. Dzurak. Single-shot readout of an electron spin in silicon. *Nature*, 467(7316):687–691, 2010.
- [MRR⁺14] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim. Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Phys. Rev. A*, 89(2):022317, 2014.
- [MS04] Y. Makhlin and A. Shnirman. Dephasing of solid-state qubits at optimal points. *Phys. Rev. Lett.*, 92(17):178301, 2004.
- [MSB⁺11] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt. 14-qubit entanglement: Creation and coherence. *Phys. Rev. Lett.*, 106(13):130506, 2011.
- [MSS00a] Y. Makhlin, G. Schön, and A. Shnirman. Nano-electronic realizations of quantum bits. *J. Low Temp. Phys.*, 118(5):751–763, 2000.
- [MSS00b] Y. Makhlin, G. Schön, G.n, and A. Shnirman. Josephson-junction qubits. *Fortschr. Phys.*, 48(9-11):1043–1054, 2000.
- [MSW⁺08] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-fidelity readout of trapped-ion qubits. *Phys. Rev. Lett.*, 100(20):200502, 2008.
- [MVP⁺17] R. McDermott, M. G. Vavilov, B. L. T. Plourde, F. K. Wilhelm, P. J. Liebermann, O. A. Mukhanov, and T. A. Ohki. Quantum-classical interface based on single flux quantum digital logic, 2017, arXiv:1710.04645.
- [MvTA⁺07] G. W. Morley, J. van Tol, A. Ardavan, K. Porfyrikis, J. Zhang, and G. A. D. Briggs. Efficient dynamic nuclear polarization at high magnetic fields. *Phys. Rev. Lett.*, 98(22):220501, 2007.
- [MW01] F. Mintert and C. Wunderlich. Ion-trap quantum logic using long-wavelength radiation. *Phys. Rev. Lett.*, 87(25):257904, 2001.
- [MZ04] M. Mosca and C. Zalka. Exact Quantum Fourier Transforms and Discrete Logarithm Algorithms. *Int. J. Quantum Inf.*, 2(1):91–100, 2004.

- [MZ⁺12] V. Mourik, K. Zuo, S. M. Frolov, S. R. Plissard, E. P. A. M. Bakkers, and L. P. Kouwenhoven. Signatures of majorana fermions in hybrid superconductor-semiconductor nanowire devices. *Science*, 336(6084):1003–1007, 2012, <http://science.sciencemag.org/content/336/6084/1003.full.pdf>.
- [Nam17] Y. Nam. Running Shor’s Algorithm on a complete, gate-by-gate implementation of a virtual, universal quantum computer, April 2017.
- [NB15a] Y. S. Nam and R. Blümel. Analytical formulas for the performance scaling of quantum processors with a large number of defective gates. *Phys. Rev. A*, 92, 2015.
- [NB15b] Y. Nam and R. Blümel. Performance scaling of the quantum Fourier transform with defective rotation gates. *Quantum Inf. Comput.*, 15(9 & 10):0721–0746, 2015.
- [NB17] Y. S. Nam and R. Blümel. Optimal length of decomposition sequences composed of imperfect gates. *Quantum Inf. Process.*, 16(5), 2017.
- [NBS⁺10] P. Neumann, J. Beck, M. Steiner, F. Rempp, H. Fedder, P. R. Hemmer, J. Wrachtrup, and F. Jelezko. Single-shot readout of a single nuclear spin. *Science*, 329(5991):542–544, 2010.
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [NER16] G. J. K. Nielsen E, Rudinger K and B.-K. R. pygsti: A python implementation of gate set tomography. <http://github.com/pygstio>, 2016.
- [NFC09] J. Niset, J. Fiurášek, and N. J. Cerf. No-go theorem for gaussian quantum error correction. *Phys. Rev. Lett.*, 102(12):120501, 2009.
- [NIS99] NIST. Data Encryption Standard (DES). Federal Information Processing Standards Publication 46-3, 1999.
- [NIS01] NIST. Advanced Encryption Standard (AES). Federal Information Processing Standards Publication 197, 2001.
- [NIS15a] NIST. Secure Hash Standard (SHS). Federal Information Processing Standards Publication 180-4, 2015.
- [NIS15b] NIST. SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions. Federal Information Processing Standards Publication 202, 2015.
- [NIS16] NIST. Advancing quantum information science: National challenges and opportunities, 2016.
- [NMM⁺14] D. Nigg, M. Müller, E. A. Martinez, P. Schindler, M. Hennrich, T. Monz, M. A. Martin-Delgado, and R. Blatt. Quantum computations on a topologically encoded qubit. *Science*, 345(6194):302–305, 2014.
- [NMR⁺06] C. Negrevergne, T. S. Mahesh, C. A. Ryan, M. Ditty, F. Cyr-Racine, W. Power, N. Boulant, T. Havel, D. G. Cory, and R. Laflamme. Benchmarking quantum control methods on a 12-qubit system. *Phys. Rev. Lett.*, 96(17):170501, 2006.
- [NMR⁺08] P. Neumann, N. Mizuochi, F. Rempp, P. Hemmer, H. Watanabe, S. Yamasaki, V. Jacques, T. Gaebel, F. Jelezko, and J. Wrachtrup. Multipartite entanglement among single spins in diamond. *Science*, 320(5881):1326–1329, 2008, <http://science.sciencemag.org/content/320/5881/1326.full.pdf>.
- [NPT99] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai. Coherent control of macroscopic quantum states in a single-cooper-pair box. *Nature*, 398(6730):786–788, 1999.
- [NRK⁺17] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, R. Barends, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Fowler, B. Foxen, R. Graff, E. Jeffrey, J. Kelly, E. Lucero, A. Megrant, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, H. Neven, and J. M. Martinis. A blueprint for demonstrating quantum supremacy with superconducting qubits, 2017, arXiv:1709.06678.

- [NSL⁺11] K. C. Nowack, M. Shafiei, M. Laforest, G. E. D. K. Prawiroatmodjo, L. R. Schreiber, C. Reichl, W. Wegscheider, and L. M. K. Vandersypen. Single-shot correlations and two-qubit gate of solid-state spins. *Science*, 333(6047):1269–1272, 2011, <http://science.sciencemag.org/content/333/6047/1269.full.pdf>.
- [NSM20] Y. Nam, Y. Su, and D. Maslov. Approximate Quantum Fourier Transform with $O(n \log(n))$ T gates, *npj Quantum Information* vol. 6, Article no. 26, 2020.
- [NSS⁺08] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma. Non-abelian anyons and topological quantum computation. *Rev. Mod. Phys.*, 80(3):1083–1159, 2008.
- [NTC11] A. Negretti, P. Treutlein, and T. Calarco. Quantum computing implementations with neutral particles. *Quantum Inf. Process.*, 10(6):721, 2011.
- [NTD⁺14] K. Nemoto, M. Trupke, S. J. Devitt, A. M. Stephens, B. Scharfenberger, K. Buczak, T. Nöbauer, M. S. Everitt, J. Schmiedmayer, and W. J. Munro. Photonic architecture for scalable quantum information processing in diamond. *Phys. Rev. X*, 4(3):031022, 2014.
- [Nyq28] H. Nyquist. Thermal agitation of electric charge in conductors. *Phys. Rev.*, 32(1):110–113, 1928.
- [OBK⁺16] P. J. J. O’Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis. Scalable quantum simulation of molecular energies. *Phys. Rev. X*, 6(3):031007, 2016.
- [OC17] J. O’Gorman and E. T. Campbell. Quantum computation with realistic magic-state factories. *Phys. Rev. A*, 95(3):032338, 2017.
- [OFV09] J. L. O’Brien, A. Furusawa, and J. Vuckovic. Photonic quantum technologies. *Nat. Photon.*, 3(12):687–695, 2009.
- [OK91] W. Ogata and K. Kurosawa. On Claw Free Families. In *International Conference on the Theory and Applications of Cryptology (ASIACRYPT ’91)*, pages 111–123, 1991.
- [OMM⁺09] S. Olmschenk, D. N. Matsukevich, P. Maunz, D. Hayes, L.-M. Duan, and C. Monroe. Quantum teleportation between distant matter qubits. *Science*, 323(5913):486–489, 2009.
- [OMT⁺99] T. P. Orlando, J. E. Mooij, L. Tian, C. H. van der Wal, L. S. Levitov, S. Lloyd, and J. J. Mazo. Superconducting persistent-current qubit. *Phys. Rev. B*, 60(22):15398–15413, 1999.
- [OOHT11] R. Okamoto, J. L. O’Brien, H. F. Hofmann, and S. Takeuchi. Realization of a knill-laflamme-milburn controlled-not photonic quantum circuit combining effective optical nonlinearities. *Proc. Natl. Acad. Sci. U.S.A.*, 108(25):10067–10071, 2011.
- [OOT⁺17] T. Ono, R. Okamoto, M. Tanida, H. F. Hofmann, and S. Takeuchi. Implementation of a quantum controlled-swap gate with photonic circuits. 7:45353, Mar 2017.
- [OPH⁺16] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf. Extending the lifetime of a quantum bit with error correction in superconducting circuits. *Nature*, 536(7617):441–445, 2016.
- [OPLTT12] T. Obata, M. Pioro-Ladrière, Y. Tokura, and S. Tarucha. The photon-assisted dynamic nuclear polarization effect in a double quantum dot. *New J. Phys.*, 14(12):123013–, 2012.
- [OPW⁺03] J. L. O’Brien, G. J. Pryde, A. G. White, T. C. Ralph, and D. Branning. Demonstration of an all-optical quantum controlled-not gate. *Nature*, 426(6964):264–267, 2003.
- [ORvO10] Y. Oreg, G. Refael, and F. von Oppen. Helical liquids and majorana bound states in quantum wires. *Phys. Rev. Lett.*, 105(17):177002, 2010.

- [OWC⁺11] C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland. Microwave quantum logic gates for trapped ions. *Nature*, 476(7359):181–184, 2011.
- [Oxf17] Oxford Instruments, 2017, <https://www.oxford-instruments.com/products/cryogenic-environments/dilution-refrigerator/cryogen-free-dilution-refrigerators/tritonxl>.
- [OYL⁺05] W. D. Oliver, Y. Yu, J. C. Lee, K. K. Berggren, L. S. Levitov, and T. P. Orlando. Mach-zehnder interferometry in a strongly driven superconducting qubit. *Science*, 310(5754):1653–1657, 2005.
- [PAL14] K. L. Pudenz, T. Albash, and D. A. Lidar. Error-corrected quantum annealing with hundreds of qubits. *Nat. Commun.*, 5, 2014.
- [PAS18] PASQuanS <https://pasquans.eu/>
- [Pau90] W. Paul. Electromagnetic traps for charged and neutral particles. *Rev. Mod. Phys.*, 62(3):531–540, 1990.
- [PBA14] D. Paredes-Barato and C. S. Adams. All-optical quantum information processing using rydberg gates. *Phys. Rev. Lett.*, 112(4):040501, 2014.
- [PCR⁺08] A. Politi, M. J. Cryan, J. G. Rarity, S. Yu, and J. L. O’Brien. Silica-on-silicon waveguide quantum circuits. *Science*, 320(5876):646–649, 2008, <http://science.sciencemag.org/content/320/5876/646.full.pdf>.
- [PdGHM07] J. H. Plantenberg, P. C. de Groot, C. J. P. M. Harmans, and J. E. Mooij. Demonstration of controlled-not quantum gates on a pair of superconducting quantum bits. *Nature*, 447(7146):836–839, 2007.
- [PF13] A. Paetznick and A. G. Fowler. Quantum circuit optimization by topological compaction in the surface code, 2013, arXiv:1304.2807.
- [PFR⁺15] D. K. Park, G. Feng, R. Rahimi, S. Labruyère, T. Shibata, S. Nakazawa, K. Sato, T. Takui, R. Laflamme, and J. Baugh. Hyperfine spin qubits in irradiated malonic acid: heat-bath algorithmic cooling. *Quantum Inf. Process.*, 14(7):2435–2461, 2015.
- [PGC⁺14] I. M. Pop, K. Geerlings, G. Catelani, R. J. Schoelkopf, L. I. Glazman, and M. H. Devoret. Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles. *Nature*, 508(7496):369–372, 2014.
- [PGMG19] E Pednault, J Gunnels, D Maslov, J Gambetta. On Quantum Supremacy, IBM Blog, 2019
- [Phi98] W. D. Phillips. Nobel lecture: Laser cooling and trapping of neutral atoms. *Rev. Mod. Phys.*, 70(3):721–741, 1998.
- [PJM⁺04] J. R. Petta, A. C. Johnson, C. M. Marcus, M. P. Hanson, and A. C. Gossard. Manipulation of a single charge in a double quantum dot. *Phys. Rev. Lett.*, 93:186802, Oct 2004.
- [PJT⁺05] J. R. Petta, A. C. Johnson, J. M. Taylor, E. A. Laird, A. Yacoby, M. D. Lukin, C. M. Marcus, M. P. Hanson, and A. C. Gossard. Coherent manipulation of coupled electron spins in semiconductor quantum dots. *Science*, 309(5744):2180–2184, 2005, <http://science.sciencemag.org/content/309/5744/2180.full.pdf>.
- [PJY⁺05] J. R. Petta, A. C. Johnson, A. Yacoby, C. M. Marcus, M. P. Hanson, and A. C. Gossard. Pulsed-gate measurements of the singlet-triplet relaxation time in a two-electron double quantum dot. *Phys. Rev. B*, 72(16):161301, 2005.
- [PKS17] M. D. Penny, D. E. Koh, R. W. Spekkens, Quantum circuit dynamics via path integrals: Is there a classical action for discrete-time paths . *New J. Phys.* 19 (2017) 073006
- [Pla19] R. Player. *On the condition number of Macaulay matrices*. Presentation at Dagstuhl Seminar 19421 Quantum Cryptanalysis, October 2019.
- [PLOT⁺08] M. Pioro-Ladrière, T. Obata, Y. Tokura, Y.-S. Shin, T. Kubo, K. Yoshida, T. Taniyama, and S. Tarucha. Electrically driven single-electron spin resonance in a slanting zeeman field. *Nat. Phys.*, 4(10):776–779, 2008.

- [PLP⁺11] A. Peruzzo, A. Laing, A. Politi, T. Rudolph, and J. L. O'Brien. Multimode quantum interference of photons in multiport integrated devices. *Nat. Commun.*, 2:224 EP –, 2011.
- [PMS⁺16] H. Paik, A. Mezzacapo, M. Sandberg, D. T. McClure, B. Abdo, A. D. Córcoles, O. Dial, D. F. Bogorin, B. L. T. Plourde, M. Steffen, A. W. Cross, J. M. Gambetta, and J. M. Chow. Experimental demonstration of a resonator-induced phase gate in a multiqubit circuit-qed system. *Phys. Rev. Lett.*, 117(25):250502, 2016.
- [Pob96] F. Pobell. *Matter and Methods at Low Temperatures*. Springer-Verlag, 1996.
- [Poz12] D. Pozar. *Microwave Engineering*. Wiley, 2012.
- [Pre97a] J. Preskill. Fault-tolerant quantum computation, 1997, arXiv:quant-ph/9712048v1, <https://arxiv.org/abs/quant-ph/9712048v1>.
- [Pre97b] J. Preskill. Reliable quantum computers. 1997, arXiv:quant-ph/9705031.
- [Pre12] J. Preskill. Quantum computing and the entanglement frontier, 2012, arXiv:1203.5813.
- [PREF17] S. Plugge, A. Rasmussen, R. Egger, and K. Flensberg. Majorana box qubits. *New J. Phys.*, 19(1):012001, 2017.
- [PRY⁺17] T. Proctor, K. Rudinger, K. Young, M. Sarovar, and R. Blume-Kohout. What randomized benchmarking actually measures, 2017, arXiv:1702.01853.
- [PS53] W. Paul and H. Steinwedel. Ein neues massenspektrometer ohne magnetfeld. *Z. Naturforsch. A*, 8:448–450, 1953.
- [PS13] P. Pham and K. M. Svore. A 2D Nearest-Neighbor Quantum Architecture for Factoring in Polylogarithmic Depth. *Quantum Inf. Comput.*, 13(11 & 12):0937–0962, 2013.
- [PSB⁺11] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Observation of high coherence in josephson junction qubits measured in a three-dimensional circuit qed architecture. *Phys. Rev. Lett.*, 107(24):240501, 2011.
- [PSL13] G. A. Paz-Silva and D. A. Lidar. Optimally combining dynamical decoupling and quantum error correction. *Sci. Rep.*, 3(1), 2013.
- [PSS⁺12] J. R. Prance, Z. Shi, C. B. Simmons, D. E. Savage, M. G. Lagally, L. R. Schreiber, L. M. K. Vandersypen, M. Friesen, R. Joynt, S. N. Coppersmith, and M. A. Eriksson. Single-shot measurement of triplet-singlet relaxation in a Si/SiGe double quantum dot. *Phys. Rev. Lett.*, 108(4):046808, 2012.
- [PSV⁺12] S. R. Plissard, D. R. Slapak, M. A. Verheijen, M. Hocevar, G. W. G. Immink, I. van Weperen, S. Nadj-Perge, S. M. Frolov, L. P. Kouwenhoven, and E. P. A. M. Bakkers. From insb nanowires to nanocubes: Looking for the sweet spot. *Nano Lett.*, 12(4):1794–1798, 2012, <http://dx.doi.org/10.1021/nl203846g>.
- [PTD⁺12] J. J. Pla, K. Y. Tan, J. P. Dehollain, W. H. Lim, J. J. L. Morton, D. N. Jamieson, A. S. Dzurak, and A. Morello. A single-atom electron spin qubit in silicon. *Nature*, 489(7417):541–545, 2012.
- [PTD⁺13] J. J. Pla, K. Y. Tan, J. P. Dehollain, W. H. Lim, J. J. L. Morton, F. A. Zwanenburg, D. N. Jamieson, A. S. Dzurak, and A. Morello. High-fidelity readout and control of a nuclear spin qubit in silicon. *Nature*, 496(7445):334–338, 2013.
- [PTR⁺17] S. Peters, L. Tiemann, C. Reichl, S. F. t, W. Dietsche, and W. Wegscheider. Improvement of the transport properties of a high-mobility electron system by intentional parallel conduction. *Appl. Phys. Lett.*, 110(4):042106, 2017, <http://dx.doi.org/10.1063/1.4975055>.
- [PvWC⁺13] S. R. Plissard, I. van Weperen, D. Car, M. A. Verheijen, G. W. G. Immink, J. Kammhuber, L. J. Cornelissen, D. B. Szombati, A. Geresdi, S. M. Frolov, L. P. Kouwenhoven, and E. P. A. M. Bakkers. Formation and electronic properties of InSb nanocrosses. *Nat. Nanotechnol.*, 8(11):859–864, 2013.

- [PWW⁺05] L. N. Pfeiffer, K. W. West, R. L. Willett, H. Akiyama, and L. P. Rokhinson. Nanostructures in gas fabricated by molecular beam epitaxy. *Bell Labs Tech. J.*, 10(3):151–159, 2005.
- [PZ03] J. Proos and C. Zalka. Shor’s discrete logarithm quantum algorithm for elliptic curves. *Quantum Inf. Comput.*, 3(4):317–344, 2003.
- [QE16] QUTE-Europe. Qt roadmap 2016, 2016, <http://quope.eu/h2020/qtflagship/roadm016>.
- [QL12] G. Quiroz and D. A. Lidar. High-fidelity adiabatic quantum computation via dynamical decoupling. *Phys. Rev. A*, 86(4):042333, 2012.
- [RB01] R. Raussendorf and H. J. Briegel. A one-way quantum computer. *Phys. Rev. Lett.*, 86(22):5188–5191, 2001.
- [RBB03] R. Raussendorf, D. E. Browne, and H. J. Briegel. Measurement-based quantum computation on cluster states. *Phys. Rev. A*, 68(2):022312, 2003.
- [RBLD12] D. Ristè, C. C. Bultink, K. W. Lehnert, and L. DiCarlo. Feedback control of a solid-state qubit using high-fidelity projective measurement. *Phys. Rev. Lett.*, 109(24):240502, 2012.
- [RCB⁺11] L. Robledo, L. Childress, H. Bernien, B. Hensen, P. F. A. Alkemade, and R. Hanson. High-fidelity projective read-out of a solid-state spin quantum register. *Nature*, 477(7366):574–578, 2011.
- [RD15] D. Ristè and L. DiCarlo. Digital feedback in superconducting quantum circuits, 2015, arXiv:1508.01385.
- [Res16] M. Research. Language-integrated quantum operations: LiqUi]. <https://www.microsoft.com/en-us/research/project/language-integrated-quantum-operations-liqui/>, 2016.
- [RFB⁺10] S. Rihani, R. Faulks, H. E. Beere, I. Farrer, M. Evans, D. A. Ritchie, and M. Pepper. Enhanced terahertz emission from a multilayered low temperature grown gas structure. *Appl. Phys. Lett.*, 96(9):091101, 2010, <http://dx.doi.org/10.1063/1.3332587>.
- [RGM⁺03] T. C. Ralph, A. Gilchrist, G. J. Milburn, W. J. Munro, and S. Glancy. Quantum computation with optical coherent states. *Phys. Rev. A*, 68(4):042319, 2003.
- [RGP⁺12] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Córcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen. Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms. *Phys. Rev. B*, 86(10):100506, 2012.
- [RGR⁺17] S. Rosenblum, Y. Gao, P. Reinhold, C. Wang, C. Axline, L. Frunzio, S. Girvin, L. Jiang, M. Mirrahimi, M. Devoret, and R. Schoelkopf. A cnot gate between multiphoton qubits encoded in two cavities, 2017, arXiv:1709.05425.
- [RH07] R. Raussendorf and J. Harrington. Fault-tolerant quantum computation with high threshold in two dimensions. *Phys. Rev. Lett.*, 98(19):190504, 2007.
- [RHG06] R. Raussendorf, J. Harrington, and K. Goyal. A fault-tolerant one-way quantum computer. *Ann. Phys.*, 321(9):2242–2270, 2006.
- [RKB⁺16] A. Reiserer, N. Kalb, M. S. Blok, K. J. M. van Bemmelen, T. H. Taminiau, R. Hanson, D. J. Twitchen, and M. Markham. Robust quantum-network memory using decoherence-protected subspaces of nuclear spins. *Phys. Rev. X*, 6(2):021040, 2016.
- [RLB⁺14] S. Ravets, H. Labuhn, D. Barredo, L. Beguin, T. Lahaye, and A. Browaeys. Coherent dipole-dipole coupling between two single rydberg atoms at an electrically-tuned forster resonance. *Nat. Phys.*, 10(12):914–917, 2014.
- [RLF12] L. P. Rokhinson, X. Liu, and J. K. Furdyna. The fractional a.c. josephson effect in a semiconductor–superconductor nanowire as a signature of majorana particles. *Nat. Phys.*, 8(11):795–799, 2012.

- [RLL09] C. A. Ryan, M. Laforest, and R. Laflamme. Randomized benchmarking of single- and multi-qubit control in liquid-state nmr quantum information processing. *New J. Phys.*, 11(1):013034, 2009.
- [RMBL08] C. A. Ryan, O. Moussa, J. Baugh, and R. Laflamme. Spin based heat engine: Demonstration of multiple rounds of algorithmic cooling. *Phys. Rev. Lett.*, 100(14):140501, 2008.
- [RMR⁺07] K. D. Raedt, K. Michielsen, H. D. Raedt, B. Trieu, G. Arnold, M. Richter, T. Lippert, H. Watanabe, and N. Ito. Massively parallel quantum computer simulator. *Comput. Phys. Commun.*, 176(2):121 – 136, 2007.
- [RNSL17c] M. Roetteler, M. Naehrig, K. M. Svore, and K. Lauter. Quantum Resource Estimates for Computing Elliptic Curve Discrete Logarithms. In *Advances in Cryptology - ASIACRYPT 2017 (Part 2)*, Springer LNCS 10625, pp. 241-270, 2017.
- [Roe17] M. Roetteler. (private email communication), October 2017.
- [ROT⁺17] M. Reagor, C. B. Osborn, N. Tezak, A. Staley, G. Prawiroatmodjo, M. Scheer, N. Alidoust, E. A. Sete, N. Didier, M. P. da Silva, E. Acala, J. Angeles, A. Bestwick, M. Block, B. Bloom, A. Bradley, C. Bui, S. Caldwell, L. Capelluto, R. Chilcott, J. Cordova, G. Crossman, M. Curtis, S. Deshpande, T. E. Bouayadi, D. Girshovich, S. Hong, A. Hudson, P. Karalekas, K. Kuang, M. Lenihan, R. Manenti, T. Manning, J. Marshall, Y. Mohan, W. O'Brien, J. Otterbach, A. Papageorge, J. P. Paquette, M. Pelstring, A. Polloreno, V. Rawat, C. A. Ryan, R. Renzas, N. Rubin, D. Russell, M. Rust, D. Scarabelli, M. Selvanayagam, R. Sinclair, R. Smith, M. Suska, T. W. To, M. Vahidpour, N. Vodrahalli, T. Whyland, K. Yadav, W. Zeng, and C. T. Rigetti. Demonstration of universal parametric entangling gates on a multi-qubit lattice, 2017, arXiv:1706.06570.
- [RS14] M. Roetteler and R. Steinwandt. A quantum circuit to find discrete logarithms on ordinary binary elliptic curves in depth $O(\log^2 n)$. *Quantum Inf. Comput.*, 14:888–900, 2014.
- [RS15] M. Roetteler and R. Steinwandt. A note on quantum related-key attacks. *Inf. Process. Lett.*, 115(1):40–44, 2015.
- [RS16a] N. J. Ross and P. Selinger. Optimal ancilla-free Clifford+T approximation of z-rotations, *Quantum Information and Computation* 16(11-12):901-953, 2016.
- [RvLK⁺12] D. Ristè, J. G. van Leeuwen, H.-S. Ku, K. W. Lehnert, and L. DiCarlo. Initialization by measurement of a superconducting quantum bit circuit. *Phys. Rev. Lett.*, 109(5):050507, 2012.
- [RWJ⁺14] T. F. Ronnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer. Defining and detecting quantum speedup. *Science*, 345(6195):420–424, 2014.
- [RWL⁺18] J. Randall, A. M. Lawrence, S. C. Webster, S. Weidt, N. V. Vitanov, and W. K. Hensinger. Generation of high-fidelity quantum control methods for multilevel systems. *Phys. Rev. A* 98(4):043414, 2018.
- [RWN⁺17] M. Radulaski, M. Widmann, M. Niethammer, J. L. Zhang, S.-Y. Lee, T. Rendler, K. G. Lagoudakis, N. T. Son, E. Janzen, T. Ohshima, J. Wrachtrup, and J. Vuckovic. Scalable quantum photonics with single color centers in silicon carbide. *Nano Lett.*, 17(3):1782–1786, 2017, <http://dx.doi.org/10.1021/acs.nanolett.6b05102>.
- [Saf16] M. Saffman. Quantum computing with atomic qubits and rydberg interactions: progress and challenges. *J. Phys. B: At., Mol. Opt. Phys.*, 49(20):202001, 2016.
- [SBK⁺17] M. Stevens, E. Bursztein, P. Karpman, A. Albertini, and Y. Markov. The first collision for full SHA-1, 2017.
- [SBM⁺11] P. Schindler, J. T. Barreiro, T. Monz, V. Nebendahl, D. Nigg, M. Chwalla, M. Hennrich, and R. Blatt. Experimental repetitive quantum error correction. *Science*, 332(6033):1059–1061, 2011.
- [SBM⁺16] S. Sheldon, L. S. Bishop, E. Magesan, S. Filipp, J. M. Chow, and J. M. Gambetta. Characterizing errors on qubit operations via iterative randomized benchmarking. *Phys. Rev. A*, 93(1):012301, 2016.

- [SBT⁺17] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas. Fast quantum logic gates with trapped-ion qubits, 2017, arXiv:1709.06952.
- [SCB⁺13] A. Stute, B. Casabone, B. Brandstatter, K. Friebe, T. Northup, and R. Blatt. Quantum-state transfer from an ion to a photon. *Nat. Photon.*, 7(3):219–222, 2013.
- [SCS⁺16] I. Schwartz, D. Cogan, E. R. Schmidgall, Y. Don, L. Gantz, O. Kenneth, N. H. Lindner, and D. Gershoni. Deterministic generation of a cluster state of entangled photons. *Science*, 354(6311):434–437, 2016, <http://science.sciencemag.org/content/354/6311/434.full.pdf>.
- [SDEW13] R. Schutjens, F. A. Dagg, D. J. Egger, and F. K. Wilhelm. Single-qubit gates in frequency-crowded transmon systems. *Phys. Rev. A*, 88:052330, 2013.
- [SDH⁺12] M. D. Shulman, O. E. Dial, S. P. Harvey, H. Bluhm, V. Umansky, and A. Yacoby. Demonstration of entanglement of electrostatically coupled singlet-triplet qubits. *Science*, 336(6078):202–205, 2012, <http://science.sciencemag.org/content/336/6078/202.full.pdf>.
- [SDP⁺12] J. Stehlik, Y. Dovzhenko, J. R. Petta, J. R. Johansson, F. Nori, H. Lu, and A. C. Gossard. Landau-Zener-Stückelberg interferometry of a single electron charge qubit. *Phys. Rev. B*, 86(12):121303, 2012.
- [Sei00] J.-P. Seifert. Using fewer Qubits in Shor’s Factorization Algorithm via Simultaneous Diophantine Approximation. Electronic Colloquium on Computational Complexity, Report 78, 2000.
- [Sel16] P. Selinger. The Quipper Language. <http://www.mathstat.dal.ca/~selinger/quipper/>, 2016.
- [SFD⁺10] D. I. Schuster, A. Fragner, M. I. Dykman, S. A. Lyon, and R. J. Schoelkopf. Proposal for manipulating and detecting spin and orbital states of trapped electrons on helium using cavity quantum electrodynamics. *Phys. Rev. Lett.*, 105(4):040503, 2010.
- [SFV⁺02] C. Santori, D. Fattal, J. Vuckovic, G. S. Solomon, and Y. Yamamoto. Indistinguishable photons from a single-photon device. *Nature*, 419(6907):594–597, 2002.
- [SGJ⁺13] M. J. Schwarz, J. Goetz, Z. Jiang, T. Niemczyk, F. Deppe, A. Marx, and R. Gross. Gradiometric flux qubits with a tunable gap. *New Journal of Physics*, 15(4):045001, 2013.
- [SHD⁺13] X. Su, S. Hao, X. Deng, L. Ma, M. Wang, X. Jia, C. Xie, and K. Peng. Gate sequence for continuous variable one-way quantum computation. *Nat. Commun.*, 4:2828, 2013.
- [SHK⁺08] J. A. Schreier, A. A. Houck, J. Koch, D. I. Schuster, B. R. Johnson, J. M. Chow, J. M. Gambetta, J. Majer, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Suppressing charge noise decoherence in superconducting charge qubits. *Phys. Rev. B*, 77(18):180502, 2008.
- [SHKW05] M. J. Storz, U. Hartmann, S. Kohler, and F. K. Wilhelm. Intrinsic phonon decoherence and quantum gates in coupled lateral quantum-dot charge qubits. *Phys. Rev. B*, 72(23):235321, 2005.
- [Sho94] P. W. Shor. Algorithms for Quantum Computation: Discrete Logarithms and Factoring. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, 1994.
- [SHT18a] D. S. Steiger, T. Häner, and M. Troyer. ProjectQ: An Open Source Software Framework for Quantum Computing, *Quantum* 2, 49, 2018.
- [Sho97] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal of Computing*, 26(5):1484–1509, 1997.
- [SJ09] R. Stock and D. F. V. James. Scalable, high-speed measurement-based quantum computer using trapped ions. *Phys. Rev. Lett.*, 102:170501, Apr 2009.
- [SJD06] L. C. H. Simon J. Devitt, Austin G. Fowler. Robustness of Shor’s algorithm. *Quant. Inf. Comp.*, 6:616–629, 2006.

- [SKD⁺10] M. Steffen, S. Kumar, D. P. DiVincenzo, J. R. Rozen, G. A. Keefe, M. B. Rothwell, and M. B. Ketchen. High-coherence hybrid superconducting qubit. *Phys. Rev. Lett.*, 105(10):100502, 2010.
- [SKS⁺14] P. Scarlino, E. Kawakami, P. Stano, M. Shafiei, C. Reichl, W. Wegscheider, and L. M. K. Vandersypen. Spin-relaxation anisotropy in a gaas quantum dot. *Phys. Rev. Lett.*, 113:256802, Dec 2014.
- [SLH⁺04] R. W. Simmonds, K. M. Lang, D. A. Hite, S. Nam, D. P. Pappas, and J. M. Martinis. Decoherence in josephson phase qubits from junction resonators. *Phys. Rev. Lett.*, 93(7):077003, 2004.
- [SLTDS10] J. D. Sau, R. M. Lutchyn, S. Tewari, and S. Das Sarma. Generic new platform for topological quantum computation using semiconductor heterostructures. *Phys. Rev. Lett.*, 104(4):040502, 2010.
- [SM00] A. Sørensen and K. Mølmer. Entanglement and quantum computation with ions in thermal motion. *Phys. Rev. A*, 62(2):022311, 2000.
- [SMCG16] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta. Procedure for systematically tuning up cross-talk in the cross-resonance gate. *Phys. Rev. A*, 93(6):060302, 2016.
- [Smi16] G. Smith. private communication, 2016.
- [SMS02] A. Shnirman, Y. Makhlin, and G. Schön. Noise and decoherence in quantum two-level systems. *Phys. Scr.*, 2002(T102):147, 2002.
- [SMS+19] M. Soeken, F. Mozafari, B. Schmitt, G. De Micheli. Compiling permutations for superconducting QPUs. Design, Automation and Test in Europe (DATE) 2019, IEEE, preprint available at https://msoeken.github.io/papers/2019_date_4.pdf.
- [SMSS06] J. Schrieffer, Y. Makhlin, A. Shnirman, and G. Schön. Decoherence from ensembles of two-level fluctuators. *New J. Phys.*, 8(1):1, 2006.
- [SNS⁺13] V. Srinivasa, K. C. Nowack, M. Shafiei, L. M. K. Vandersypen, and J. M. Taylor. Simultaneous spin-charge relaxation in double quantum dots. *Phys. Rev. Lett.*, 110(19):196803, 2013.
- [Sol00] R. Solovay. Lie groups and quantum circuits, 2000, <http://www.msri.org/publications/ln/msri/2000/qcomputing/solovay/1/>.
- [SRDR17] E. A. Sete, M. J. Reagor, N. Didier, and C. T. Rigetti. Charge- and flux-insensitive tunable superconducting qubit. *Phys. Rev. Applied*, 8(2):024004, 2017.
- [SRW⁺14] R. C. Sterling, H. Rattanasonti, S. Weidt, K. Lake, P. Srinivasan, S. C. Webster, M. Kraft, and W. K. Hensinger. Fabrication and operation of a two-dimensional ion-trap lattice on a high-voltage microchip. *Nat. Commun.*, 5:3637, 2014.
- [SS07] G. Smith and J. A. Smolin. Degenerate quantum codes for pauli channels. *Phys. Rev. Lett.*, 98(3):030501, 2007.
- [SS10] G. Schaller and R. Schützhold. The role of symmetries in adiabatic quantum algorithms. *Quantum Inf. Comput.*, 10:0109–0140, 2010.
- [SS17] T. Santoli and C. Schaffner. Using Simon’s Algorithm to Attack Symmetric-Key Cryptographic Primitives. *Quantum Inf. Comput.*, 17(1&2):65–78, 2017.
- [SSA⁺15] A. Srivastava, M. Sidler, A. V. Allain, D. S. Lembke, A. Kis, and Imamoglu. Optically active quantum dots in monolayer wse₂. *Nat. Nanotechnol.*, 10(6):491–496, 2015.
- [SSAG16] M. Smelyanskiy, N. P. D. Sawaya, and A. Aspuru-Guzik. qhipster: The quantum high performance software testing environment, 2016, arXiv:1601.07195.
- [SSK⁺18] A. Stockklauser, P. Scarlino, J. V. Koski, S. Gasparinetti, C. K. Andersen, C. Reichl, W. Wegscheider, T. Ihn, K. Ensslin, and A. Wallraff. Strong Coupling Cavity QED with Gate-Defined Double Quantum Dots Enabled by a High Impedance Resonator. *Phys. Rev. X*, 7(1):011030, 2018.

- [SSO92] A. Suarez, R. Silbey, and I. Oppenheim. Memory effects in the relaxation of quantum open systems. *J. Chem. Phys.*, 97(7):5101–5107, 1992, <http://dx.doi.org/10.1063/1.463831>.
- [Ste96] A. Steane. Multiple-particle interference and quantum error correction. *Proc. Royal Soc. A*, 452(1954):2551–2577, 1996.
- [Ste03] A. Steane. Overhead and noise threshold of fault-tolerant quantum error correction. *Phys. Rev. A*, 68:042322, 2003.
- [SVMA+17] A. Scherer, B. Valiron, S.-C. Mau, S. Alexander, E. van den Berg, and T. E. Chapuran. Concrete resource analysis of the quantum linear system algorithm used to compute the electromagnetic scattering cross section of a 2D target, *Quantum Inf. Process.* (2017) 16: 60. Preprint available at arXiv:1505.06552v2.
- [SW03] M. J. Storcz and F. K. Wilhelm. Decoherence and gate performance of coupled solid-state qubits. *Phys. Rev. A*, 67(4):042319, 2003.
- [SWM10] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with Rydberg atoms. *Rev. Mod. Phys.*, 82(3):2313–2363, 2010.
- [SWS16] Y. R. Sanders, J. J. Wallman, and B. C. Sanders. Bounding quantum gate error rate based on reported average fidelity. *New J. Phys.*, 18(1):012002, 2016.
- [SZH⁺16] H.-H. Sun, K.-W. Zhang, L.-H. Hu, C. Li, G.-Y. Wang, H.-Y. Ma, Z.-A. Xu, C.-L. Gao, D.-D. Guan, Y.-Y. Li, C. Liu, D. Qian, Y. Zhou, L. Fu, S.-C. Li, F.-C. Zhang, and J.-F. Jia. Majorana zero mode detected with spin selective andreev reflection in the vortex of a topological superconductor. *Phys. Rev. Lett.*, 116(25):257003, 2016.
- [SZK⁺18] N. Samkharadze, G. Zheng, N. Kalhor, D. Brousse, A. Sammak, U.C. Mendes, A. Blais, G. Scappucci, L.M.K. Vandersypen. Strong spin-photon coupling in silicon. *Science* 359:1123, 2018.
- [SZW93] R. I. Shekhter, A. M. Zagoskin, and G. Wendin. Oxygen diffusion and dynamical disorder in high- t_c superconductors: low frequency noise in superconducting tunnel junctions. *Z. Phys. B*, 91(3):277–284, 1993.
- [TB05] B. M. Terhal and G. Burkard. Fault-tolerant quantum computation for local non-markovian noise. *Phys. Rev. A*, 71(1):012336, 2005.
- [TBF18] D.K. Tuckett, S.D. Bartlett, and S.T. Flammia. Ultrahigh Error Threshold for Surface Codes with Biased Noise. *Phys. Rev. Lett.* 120(5):050505, 2018.
- [TBJ⁺11] N. Timoney, I. Baumgart, M. Johanning, A. F. Varon, M. B. Plenio, A. Retzker, and C. Wunderlich. Quantum gates and memory using microwave-dressed states. *Nature*, 476(7359):185–188, 2011.
- [TCT⁺10] E. Togan, Y. Chu, A. S. Trifonov, L. Jiang, J. Maze, L. Childress, M. V. G. Dutt, A. S. Sorensen, P. R. Hemmer, A. S. Zibrov, and M. D. Lukin. Quantum entanglement between an optical photon and a solid-state spin qubit. *Nature*, 466(7307):730–734, 2010.
- [TDL⁺11] A. Tipsmark, R. Dong, A. Laghaout, P. Marek, M. Ježek, and U. L. Andersen. Experimental demonstration of a hadamard gate for coherent state qubits. *Phys. Rev. A*, 84(5):050301, 2011.
- [TdVR02] C. M. Tesch and R. de Vivie-Riedle. Quantum computation with vibrationally excited molecules. *Phys. Rev. Lett.*, 89(15):157901, 2002.
- [TED⁺05] J. M. Taylor, H.-A. Engel, W. Dur, A. Yacoby, C. M. Marcus, P. Zoller, and M. D. Lukin. Fault-tolerant architecture for quantum computation using electrically controlled semiconductor spins. *Nat. Phys.*, 1(3):177–183, 2005.
- [TGA⁺05] F. Troiani, A. Ghirri, M. Affronte, S. Carretta, P. Santini, G. Amoretti, S. Piligkos, G. Timco, and R. E. P. Winpenny. Molecular engineering of antiferromagnetic rings for quantum computation. *Phys. Rev. Lett.*, 94(20):207208, 2005.

- [THWZ16] A. Tayebi, T. N. Hoatson, J. Wang, and V. Zelevinsky. Environment-protected solid-state-based distributed charge qubit. *Phys. Rev. B*, 94(23):235150, 2016.
- [TKL⁺04] F. Tafuri, J. R. Kirtley, F. Lombardi, T. Bauch, E. Il'ichev, F. M. Granozio, D. Stornaiuolo, and U. S. di Uccio. Flavours of intrinsic d-wave induced effects in $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$ grain boundary josephson junctions. *Supercond. Sci. Technol.*, 17(5):S202, 2004.
- [TKO⁺16] K. Takeda, J. Kamioka, T. Otsuka, J. Yoneda, T. Nakajima, M. R. Delbecq, S. Amaha, G. Allison, T. Kodera, S. Oda, and S. Tarucha. A fault-tolerant addressable spin qubit in a natural silicon quantum dot. *Sci. Adv.*, 2(8):e1600694–e1600694, 2016.
- [TLA⁺11] J. D. Teufel, D. Li, M. S. Allman, K. Cicak, A. J. Sirois, J. D. Whittaker, and R. W. Simmonds. Circuit cavity electromechanics in the strong-coupling regime. *Nature*, 471(7337):204–208, 2011.
- [TM96] K. Takeo and I. Masatoshi. Macroscopic quantum tunneling of a fluxon in a long josephson junction. *J. Phys. Soc. Jpn.*, 65(9):2963–2975, 1996.
- [TMS⁺15] G. Tosi, F. A. Mohiyaddin, V. Schmitt, S. Tenberg, R. Rahman, G. Klimeck, and A. Morello. Silicon quantum processor with robust long-distance qubit couplings, 2015, arXiv:1509.08538v2, <https://arxiv.org/abs/1509.08538v2>.
- [TMW16] L. S. Theis, F. Motzoi, and F. K. Wilhelm. Simultaneous gates in frequency-crowded multilevel systems using fast, robust, analytic control shapes. *Phys. Rev. A*, 93(1):012324, 2016.
- [TMWS16] L. S. Theis, F. Motzoi, F. K. Wilhelm, and M. Saffman. High-fidelity rydberg-blockade entangling gate using shaped, analytic pulses. *Phys. Rev. A*, 94:032306, Sep 2016.
- [TPLO⁺10] T. Takakura, M. Pioro-Ladri sre, T. Obata, Y.-S. Shin, R. Brunner, K. Yoshida, T. Taniyama, and S. Tarucha. Triple quantum dot device designed for three spin qubits. *Applied Physics Letters*, 97(21):212104, 2010, <https://doi.org/10.1063/1.3518919>.
- [TvdWOT06] Y. Tokura, W. G. van der Wiel, T. Obata, and S. Tarucha. Coherent single electron spin control in a slanting zeeman field. *Phys. Rev. Lett.*, 96(4):047202, 2006.
- [VAC⁺02] D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve, and M. H. Devoret. Manipulating the quantum state of an electrical circuit. *Science*, 296(5569):886–889, 2002.
- [VAPS⁺15] W. Vinci, T. Albash, G. Paz-Silva, I. Hen, and D. A. Lidar. Quantum annealing correction with minor embedding. *Phys. Rev. A*, 92(4):042310, 2015.
- [VBR08] M. Varnava, D. E. Browne, and T. Rudolph. How good must single photon sources and detectors be for efficient linear optical quantum computation? *Phys. Rev. Lett.*, 100(6):060502, 2008.
- [VC76] R. F. Voss and J. Clarke. $1/f$ noise from systems in thermal equilibrium. *Phys. Rev. Lett.*, 36(1):42–45, 1976.
- [VC05] L. Vandersypen and I. Chuang. NMR techniques for quantum control and computation. *Rev. Mod. Phys.*, 76(4):1037–1069, 2005.
- [VD14] G. Viola and D. P. DiVincenzo. Hall effect gyrators and circulators. *Phys. Rev. X*, 4(2):021019, 2014.
- [vDMV01] W. van Dam, M. Mosca, and U. Vazirani. How Powerful is Adiabatic Quantum Computation? In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, pages 279–287, 2001.
- [VF17] S. Vijay and L. Fu. Quantum error correction for complex and majorana fermion qubits, 2017, arXiv:1703.00459.
- [Vio01] L. Viola. Experimental realization of noiseless subsystems for quantum information processing. *Science*, 293(5537):2059–2063, 2001.
- [VKL99] L. Viola, E. Knill, and S. Lloyd. Dynamical decoupling of open quantum systems. *Phys. Rev. Lett.*, 82(12):2417–2421, 1999.

- [vL10] P. van Loock. A note on quantum error correction with continuous variables. *J. Mod. Opt.*, 57(19):1965–1971, 2010, <http://dx.doi.org/10.1080/09500340.2010.499047>.
- [VPK⁺16] R. Versluis, S. Poletto, N. Khammassi, N. Haider, D. J. Michalak, A. Bruno, K. Bertels, and L. DiCarlo. Scalable quantum circuit and control for a superconducting surface code, 2016, arXiv:1612.08208.
- [VPS⁺14] U. Vool, I. M. Pop, K. Sliwa, B. Abdo, C. Wang, T. Brecht, Y. Y. Gao, S. Shankar, M. Hatridge, G. Catelani, M. Mirrahimi, L. Frunzio, R. J. Schoelkopf, L. I. Glazman, and M. H. Devoret. Non-poissonian quantum jumps of a fluxonium qubit due to quasiparticle excitations. *Phys. Rev. Lett.*, 113(24):247001, 2014.
- [VSB⁺01] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang. Experimental realization of shor’s quantum factoring algorithm using nuclear magnetic resonance. *Nature*, 414(6866):883–887, 2001.
- [vWKG15] D. J. van Woerkom, A. Geresdi, and L. P. Kouwenhoven. One minute parity lifetime of a nbtin cooper-pair transistor. *Nat. Phys.*, 11(7):547–550, 2015.
- [vWPB⁺13] I. van Weperen, S. R. Plissard, E. P. A. M. Bakkers, S. M. Frolov, and L. P. Kouwenhoven. Quantized conductance in an insb nanowire. *Nano Lett.*, 13(2):387–391, 2013.
- [VYH⁺15] M. Veldhorst, C. H. Yang, J. C. C. Hwang, W. Huang, J. P. Dehollain, J. T. Muhonen, S. Simmons, A. Laucht, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak. A two-qubit logic gate in silicon. *Nature*, 526(7573):410–414, 2015.
- [WAB15] F. H. E. Watson, H. Anwar, and D. E. Browne. Fast fault-tolerant decoder for qubit and qudit surface codes. *Phys. Rev. A*, 92(3):032309, 2015.
- [WBE16] J. J. Wallman, M. Barnhill, and J. Emerson. Robust characterization of leakage errors. *New J. Phys.*, 18(4):043021, 2016.
- [Wei88] M. B. Weissman. 1/f noise and other slow, nonexponential kinetics in condensed matter. *Rev. Mod. Phys.*, 60(2):537–571, 1988.
- [WF14a] J. J. Wallman and S. T. Flammia. Randomized benchmarking with confidence. *New J. Phys.*, 16(10):103032, 2014.
- [WF14b] A. C. Whiteside and A. G. Fowler. Upper bound for loss in practical topological-cluster-state quantum computing. *Phys. Rev. A*, 90(5):052316, 2014.
- [WFH11] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg. Surface code quantum computing with error rates over 1, *Phys. Rev. A*, 83(2):020302, 2011.
- [WG17] C. J. Wood and J. M. Gambetta. Quantification and characterization of leakage errors, 2017, arXiv:1704.03081.
- [WHD⁺18] S. Welte, B. Hacker, S. Daiss, S. Ritter, and G. Rempe. Photon-Mediated Quantum Gate between Two Neutral Atoms in an Optical Cavity, *Phys. Rev. X* 8(1):011018, 2018.
- [Wil08] F. K. Wilhelm. Quantum oscillations in the spin-boson model: reduced visibility from non-markovian effects and initial entanglement. *New J. Phys.*, 10(11):115011, 2008.
- [Wil09] F. Wilczek. Majorana returns. *Nat. Phys.*, 5(9):614–618, 2009.
- [Wil13] R. L. Willett. The quantum hall effect at 5/2 filling factor. *Rep. Prog. Phys.*, 76(7):076501, 2013.
- [WKG⁺17] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff. Rapid high-fidelity single-shot dispersive readout of superconducting qubits. *Phys. Rev. Applied*, 7(5):054020, 2017.
- [WKMS11] C. Weitenberg, S. Kuhr, K. Mølmer, and J. F. Sherson. Quantum computation architecture using optical tweezers. *Phys. Rev. A*, 84(3):032322, 2011.

- [WKS⁺16] D. R. Ward, D. Kim, D. E. Savage, M. G. Lagally, R. H. Foote, M. Friesen, S. N. Coppersmith, and M. A. Eriksson. State-conditional coherent charge qubit oscillations in a si/sige quadruple quantum dot. *Nat. Partn. J. Quantum Inf.*, 2(1):16032, 2016.
- [WKWW16] Y. Wang, A. Kumar, T.-Y. Wu, and D. S. Weiss. Single-qubit gates based on targeted phase shifts in a 3d neutral atom array. *Science*, 352(6293):1562–1565, 2016, <http://science.sciencemag.org/content/352/6293/1562.full.pdf>.
- [WL95] P. A. Willems and K. G. Libbrecht. Creating long-lived neutral-atom traps in a cryogenic environment. *Phys. Rev. A*, 51(2):1403–1406, 1995.
- [WL02] L.-A. Wu and D. A. Lidar. Qubits as parafermions. *J. Math. Phys.*, 43(9):4506–4525, 2002, <http://dx.doi.org/10.1063/1.1499208>.
- [WL17] J. R. Wootton and D. Loss. A repetition code of 15 qubits, 2017, arXiv:1709.00990.
- [WLL⁺03] A. Wallraff, A. Lukashenko, J. Lisenfeld, A. Kemp, M. V. Fistul, Y. Koval, and A. V. Ustinov. Quantum dynamics of a single vortex. *Nature*, 425(6954):155–158, 2003.
- [WM11] H. M. Wiseman, G. J. Milburn, *Quantum Measurement and Control*, Cambridge University Press, 2011
- [WMJM17] J.D. Wong-Campos, S.A. Moses, K.G. Johnson, and C. Monroe. Demonstration of Two-Atom Entanglement with ultra-fast Optical Pulses. *Phys. Rev. Lett.*, 119(23):230501, 2017.
- [WPGP⁺12] C. Weedbrook, S. Pirandola, R. García-Patrón, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd. Gaussian quantum information. *Rev. Mod. Phys.*, 84(2):621–669, 2012.
- [WPK⁺17] T. F. Watson, S. G. J. Philips, E. Kawakami, D. R. Ward, P. Scarlino, M. Veldhorst, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen. A programmable two-qubit quantum processor in silicon, 2017, *Nature* 555(2):633–677, 2018.
- [WRR⁺05] P. Walther, K. J. Resch, T. Rudolph, E. Schenck, H. Weinfurter, V. Vedral, M. Aspelmeyer, and A. Zeilinger. Experimental one-way quantum computing. *Nature*, 434(7030):169–176, 2005.
- [WSHG06] F. K. Wilhelm, M. J. Storcz, U. Hartmann, and M. R. Geller. Superconducting qubits ii: Decoherence, 2006, arXiv:cond-mat/0603637.
- [WWZ⁺14] G. Waldherr, Y. Wang, S. Zaiser, M. Jamali, T. Schulte-Herbruggen, H. Abe, T. Ohshima, J. Isoya, J. F. Du, P. Neumann, and J. Wrachtrup. Quantum error correction in a solid-state hybrid spin register. *Nature*, 506(7487):204–207, February 2014.
- [WZP18] L. Wossnig, Z. Zhao, and A. Prakash. *Quantum Linear System Algorithm for Dense Matrices*. *Phys. Rev. Lett.* 120 (5 2018), p. 050502, 2018.
- [XHL⁺17] T. Xin, S. Huang, S. Lu, K. Li, Z. Luo, Z. Yin, J. Li, D. Lu, G. Long, and B. Zeng. Nmrcloudq: A quantum cloud experience on a nuclear magnetic resonance quantum computer, 2017, arXiv:1710.03646.
- [XLM⁺15] T. Xia, M. Lichtman, K. Maller, A. W. Carr, M. J. Piotrowicz, L. Isenhower, and M. Saffman. Randomized benchmarking of single-qubit gates in a 2d array of neutral-atom qubits. *Phys. Rev. Lett.*, 114(10):100503, 2015.
- [XWH⁺18] X. Xue, T. F. Watson, J. Helsen, D. R. Ward, D. E. Savage, M. G. Lagally, S. N. Coppersmith, M. A. Eriksson, S. Wehner, L. M. K. Vandersypen. Benchmarking Gate Fidelities in a Si/SiGe Two-Qubit Device. ArXiv:1811.04002.
- [YC10] P. Y. Yu and M. Cardona. *Fundamentals of Semiconductors*. Graduate Texts in Physics. Springer Berlin Heidelberg, 2010.

- [YFK⁺16] G. Yang, A. Fragner, G. Koolstra, L. Ocola, D. A. Czaplewski, R. J. Schoelkopf, and D. I. Schuster. Coupling an ensemble of electrons on superfluid helium to a superconducting circuit. *Phys. Rev. X*, 6(1):011031, 2016.
- [Y GK⁺16] F. Yan, S. Gustavsson, A. Kamal, J. Birenbaum, A. P. Sears, D. Hover, T. J. Gudmundsen, D. Rosenberg, G. Samach, S. Weber, J. L. Yoder, T. P. Orlando, J. Clarke, A. J. Kerman, and W. D. Oliver. The flux qubit revisited to enhance coherence and reproducibility. *Nat. Commun.*, 7:12964 EP –, 2016.
- [YHO⁺08] T. Yamamoto, K. Hayashi, Ş. K. Özdemir, M. Koashi, and N. Imoto. Robust photonic entanglement distribution by state-independent encoding onto decoherence-free subspace. *Nat. Photon.*, 2(8):488–491, 2008.
- [YJG⁺12] N. Y. Yao, L. Jiang, A. V. Gorshkov, P. C. Maurer, G. Giedke, J. I. Cirac, and M. D. Lukin. Scalable architecture for a room temperature solid-state quantum information processor. *Nat. Commun.*, 3:800 EP –, 2012.
- [YSBK13] K. C. Young, M. Sarovar, and R. Blume-Kohout. Error suppression and error correction in adiabatic quantum computation: Techniques and challenges. *Phys. Rev. X*, 3(4), 2013.
- [YSP14] A. J. Young, B. D. Schultz, and C. J. Palmstrøm. Lattice distortion in single crystal rare-earth arsenide/gaas nanocomposites. *Appl. Phys. Lett.*, 104(7):073114, 2014, <http://dx.doi.org/10.1063/1.4865905>.
- [YTO⁺18] J. Yoneda, K. Takeda, T. Otsuka, T. Nakajima, M.R. Delbecq, G. Allison, T. Honda, T. Kodera, S. Oda, Y. Hoshi, N. Usami, K.M. Itoh, and S. Tarucha. A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%, *Nat. Nano.* 13(2):102–107, 2018.
- [YUvLF08] M. Yukawa, R. Ukai, P. van Loock, and A. Furusawa. Experimental generation of four-mode continuous-variable cluster states. *Phys. Rev. A*, 78(1):012301, 2008.
- [YYK⁺16] J.-i. Yoshikawa, S. Yokoyama, T. Kaji, C. Sornphiphatphong, Y. Shiozawa, K. Makino, and A. Furusawa. Generation of one-million-mode continuous-variable cluster state by unlimited time-domain multiplexing. *APL Photonics*, 1(6):060801, 2016.
- [Zag97] A. M. Zagoskin. The half-periodic josephson effect in an s-wave superconductor - normal-metal - d-wave superconductor junction. *J. Phys.: Condens. Matter*, 9(31):L419, 1997.
- [Zal99] C. Zalka. Grover’s quantum searching algorithm is optimal. *Phys. Rev. A*, pages 2746–2751, 1999.
- [Zal08] C. Zalka. Shor’s algorithm with fewer (pure) qubits. arXiv:quant-ph/0601097, 2008.
- [ZBD14] M. Zwerger, H. J. Briegel, and W. Dür. Hybrid architecture for encoded measurement-based quantum computation. *Sci. Rep.*, 4(1), 2014.
- [ZHL⁺16] C. Zhang, Y.-F. Huang, B.-H. Liu, C.-F. Li, and G.-C. Guo. Experimental generation of a high-fidelity four-photon linear cluster state. *Phys. Rev. A*, 93(6):062329, Jun 2016.
- [ZHM⁺16] D. M. Zajac, T. M. Hazard, X. Mi, E. Nielsen, and J. R. Petta. Scalable gate architecture for a one-dimensional array of semiconductor spin qubits. *Phys. Rev. Applied*, 6:054013, Nov 2016.
- [ZPH⁺17] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z. X. Gong, and C. Monroe. Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator, 2017, arXiv:1708.01044.
- [ZSBS14] J. Zhang, A. M. Souza, F. D. Brandao, and D. Suter. Protected quantum computing: Interleaving gate operations with dynamical decoupling sequences. *Phys. Rev. Lett.*, 112(5):050502, 2014.
- [ZSR⁺17] D. M. Zajac, A. J. Sigillito, M. Russ, F. Borjans, J. M. Taylor, G. Burkard, and J. R. Petta. Quantum cnot gate for spins in silicon, 2017, arXiv:1708.03530.
- [ZVSW03] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley. Geometric theory of nonlocal two-qubit operations. *Phys. Rev. A*, 67(4):042313, 2003.

Keywords and Abbreviations

Abbreviations..... 267

logical qubit..... 62

NMR.....

 Nuclear Magnetic Resonance..... 45, 47f., 50

QFT.....

 Quantum Fourier Transform..... 42, 126

QST.....

 Quantum State Tomography..... 54ff.

RB.....

 Randomized Benchmarking .44, 48, 55ff., 97, 99f., 102, 104, 138, 151, 157, 160, 174, 188, 197, 222, 230f., 236, 258

SPAM.....

 State Preparation And Measurement..... 44, 49, 55