

DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY


Error-corrected quantum processing with neutral atoms

Citation

Bluvstein, Dolev. 2025. Error-corrected quantum processing with neutral atoms. Doctoral Dissertation, Harvard University Graduate School of Arts and Sciences.

Link

https://dash.harvard.edu/handle/1/42725323

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at

https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFlOWI2M2EwYTg

Accessibility

https://accessibility.huit.harvard.edu/digital-accessibility-policy

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story

HARVARD
Kenneth C. Griffin

GRADUATE SCHOOL
OF ARTS AND SCIENCES

THESIS ACCEPTANCE CERTIFICATE

The undersigned, appointed by the

Department of Physics

have examined a dissertation entitled

Error-corrected quantum processing with neutral atoms

presented by Dolev Bluvstein

candidate for the degree of Doctor of Philosophy and hereby certify that it is worthy of acceptance.

Signature _____

Typed name: Prof. Mikhail Lukin

Signature _____

Typed name: Prof. Norman Yao

Signature _____

Typed name: Prof. Vladan Vuletić (MIT)

Date: May 7, 2025

Error-corrected quantum processing with neutral atoms

A DISSERTATION PRESENTED
BY
DOLEV BLUVSTEIN
TO
THE DEPARTMENT OF PHYSICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHYSICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS
MAY 2025

©2025 – Dolev Bluvstein
all rights reserved.

Thesis advisor: Professor Mikhail D. Lukin Dolev Bluvstein

Error-corrected quantum processing with neutral atoms

Abstract

Quantum computers open new scientific avenues, from exploring complex quantum mechanical systems to new computational paradigms, but face the fundamental challenge of decoherence. Remarkably, decoherence can be prevented by creating highly entangled states of physical qubits that encode an error-corrected "logical" qubit. This thesis will describe the development of quantum computing with reconfigurable arrays of neutral atoms and their use for quantum processing with logical qubits. Quantum processing in this approach is based on the coherent transport of atoms shuttled by optical tweezers, enabling any-to-any connectivity, high-fidelity programmable logic, and mid-circuit processing within a zoned architecture. Logical qubit processing is greatly facilitated by parallel control and transversal operations, and is used for experiments ranging from entangling logical qubits to their use for precise simulation of quantum scrambling. Core physical mechanisms for achieving deep-circuit, universal algorithms with logical qubits are identified, and these are leveraged into new techniques that greatly reduce overheads for large-scale computation. Finally, applications of quantum processors will be explored via analog and gate-based quantum simulations. These results, alongside other recent advances, herald a transition to error-corrected quantum processing, establishing foundations that can enable future large-scale quantum computers and their useful applications.

Contents

Title i
Copyright ii
Abstract iii
Table of Contents iv
Dedication vii
Acknowledgments viii
Author Contributions xi

1 Introduction 1
1.1 Quantum error correction: History and background 1
1.2 QEC basics 5
1.3 Estimates of required resources: QEC is challenging! 8
1.4 An inflection point 8
1.5 Overview of thesis 9

2 Quantum computing with neutral atoms 11
2.1 Atoms 11
2.2 Trapping atoms in optical tweezers 13
2.3 Building atomic arrays 15
2.4 Atomic level structure and quantum operations 17
2.5 Programmable quantum circuits via atom reconfiguration 19
2.6 High-fidelity two-qubit gates 24
2.7 Programmable single-qubit gates 34
2.8 Mid-circuit readout and feedforward 41
2.9 Atomic qubit re-use and loss-resolving measurement 45
2.10 Fast cycle rates with atomic qubit re-use 49
2.11 Overview of system and circuit programming approach 49

3 Quantum processing with logical qubits 56
3.1 In classical computing we operate on bits. In quantum computing we operate on *logical* qubits. 57
3.2 Transversal gates 57

3.3 The importance of the transversal CNOT . . . 60
3.4 Zoned architecture and efficient control . . . 61
3.5 Stabilizer measurement and removing entropy . . . 63
3.6 Universal processing. Clifford and non-Clifford operations. . . . 65
3.7 Do we need to flip physical qubits back in hardware? . . . 68
3.8 Considering algorithms as a whole, and correlated decoding . . . 69
3.9 Fast correlated decoding of transversal logical algorithms . . . 71
3.10 Fault-tolerant compiling . . . 74
3.11 Reduction of space-time overheads . . . 75
3.12 Experimental studies of error-corrected processing . . . 78

4 Experiments with logical qubit algorithms **79**
4.1 Logical processor based on atom arrays . . . 82
4.2 Improving entangling gates with code distance . . . 83
4.3 Fault-tolerant logical algorithms . . . 85
4.4 Complex logical circuits using 3D codes . . . 90
4.5 Quantum simulations with logical qubits . . . 94
4.6 Outlook . . . 95

5 Physical mechanisms for a fault-tolerant quantum architecture **97**
5.1 Stabilizer measurement during algorithms . . . 100
5.2 Below-threshold performance . . . 102
5.3 Analog evolution with digital gates . . . 104
5.4 Deep circuits at constant entropy . . . 105
5.5 Outlook . . . 108

6 What comes next? **110**
6.1 Summary . . . 110
6.2 Science of logical qubits . . . 111
6.3 Toward future processors and QEC . . . 113

7 Exploring applications of quantum processors **115**
7.1 Overview . . . 115
7.2 Quantum phases of matter on a 256-atom programmable quantum simulator . . . 117
7.3 Controlling quantum many-body dynamics in driven Rydberg atom arrays . . . 120
7.4 Probing topological spin liquids on a programmable quantum simulator . . . 123
7.5 Quantum optimization of maximum independent set using Rydberg atom arrays . 126
7.6 Quantum coarsening and collective dynamics on a programmable simulator . . . 129
7.7 Hybrid analog-digital quantum simulation . . . 132
7.8 Probing topological matter and fermion dynamics on a neutral-atom quantum computer . . . 135

References 162

To the amazing friends, family, and teammates I went on this adventure with.

Acknowledgments

This thesis is the culmination of an exhilaratingly fun adventure with some of the most amazing people I have ever met. There was no shortage of characters in the cast, and the diverse backgrounds and relationships with all these people have resulted in the most fulfilling period of my life thus far.

I have been very fortunate to see the way that Misha operates - a true master of his craft. I learned many lessons about science from Misha, but I think the most valuable one has been to always ask what *actually* makes sense scientifically and to never hide from obvious truths, which is a deceptively simple tactic that has enormous long-term payoff. Another key lesson from Misha is that if you can do everything yourself, then that really just means that you are bad at picking collaborators, and instead you should be striving to surround yourself with experts in other topics who can do things you cannot. Vladan Vuletic and Markus Greiner have been remarkable examples of this, and have provided exceptional scientific and technical guidance throughout my PhD. I also want to thank Norman Yao for serving on my thesis committee and for being a friend and collaborator since we first met working on NVs.

I have overlapped with many remarkable scientists in the atom array experimental team throughout my PhD: Ahmed Omran, Alex Keesling, Harry Levine, Giulia Semeghini, Sepehr Ebadi, Tout Wang, Tom Manovitz, Simon Evered, Sasha Geim, Sophie Li, and Muqing Xu. Ahmed Omran is possibly one of the kindest humans I have ever met, and the world is better off that he was put here. We worked very closely together on the driven scars project and he taught me many invaluable lessons, including crucially how to manage a larger collaboration with theorists. Alex Keesling was such an effective leader and so level-headed, and taught me many core aspects of AMO physics. Giulia Semeghini brought a significant degree of wisdom and good decision-making to the group, and her scientific insights were always extremely useful in every project she contributed to. Tout Wang never hesitates to help you fix your laser when you're in a bind, and it has been a privilege working together on topics ranging from scientific projects to coordinating efforts in our community.

It is hard to succinctly put how much better off I am for having met Harry Levine. I learned a lot of aspects of the experiment alone in lab during COVID with Harry on zoom in my earphones instructing me like we were playing a bomb defusal game. I vividly remember 4 AM nights in lab when we were first trying to get coherent atom transport working and blasting boiler room sets on the speakers. Not only is Harry an incredible scientist but he has set such an exceptional example of what it means to be good to others and how to be an exceptional teammate. Harry has an incredible talent for simultaneously making lab a fun place while being so productive, and his future students

are luckier than they can imagine.

Sepehr and I worked together for four years and it is amazing how much we accomplished together. We had such incredible complement in skillset and I think that is core to how many different things we were able to achieve. I am still awed by Sepehr's knowledge. I often wondered if I would eventually catch up to his optics knowledge but now I am at the end of my PhD and don't know if I'm even at the level he was at in his second year! Sepehr taught me many core aspects of AMO - not only technical facts, but also key principles of doing things properly.

I am a big believer that productivity is dictated by making good decisions. Bringing Simon Evered into the atom array 1 team is possibly one of the best examples of this from my life. Simon is exceptionally creative and an incredible teammate. Simon and I have quite different personalities - I am a tad hyper, and Simon is calm and peaceful. Simon's calm demeanor coupled with his insatiable desire to make progress makes him an incredible driver in lab while keeping things level. We had quite a few late nights figuring out how to make high-fidelity gates and then would go get bagels in the morning. Simon is really an amazing teammate and friend. When the power goes out, Simon is in lab turning the lasers back on.

When Sasha Geim first joined the lab she neither knew what a lens mount was nor what a BNC cable was. That is hilarious because I think she is now more competent and knowledgeable than I am. Seeing Sasha grow has been one of the most fulfilling parts of my PhD, and working together on developing tons of new tools and using them for quantum error correction has been so much fun.

Sophie Li has grown into an exceptional scientist. The coarsening project was amazing to watch as there were many points where it seemed it wouldn't work out, but Sophie persevered and turned seeming "bugs" into a new physics discovery about Higgs modes. Sophie is a blast to hang out with in lab. Sophie and Sasha are the future of the experiment, and it is in very good hands.

Tom Manovitz brought many new perspectives from his ion trapping background and this deeply enriched our understanding of various key topics. Tom has an incredible knack for being simultaneously kind and confident and I have learned a lot from watching the way he effectively manages projects like the coarsening project or technical projects in the lab.

Muqing Xu and I have only overlapped for a short time so far, but it has been already a lot of fun, and I admire his eagerness to learn and broad knowledge base. I am very excited to see the adventures the next generation gets up to.

It is amazing to reflect on the variety of science I was able to do with Marcin Kalinowski and Nishad Maskara. In our first year I knew nothing about quantum information science and Nishad taught me basic things like CNOT gates and the toric code. We then went onto discover this phenomenon of driven scar stabilization, and the adventures only continued from there. Marcin is a very unique scientist - equipped with extremely sharp technical skills in theory, but with a deep drive to build technology, and was with me in lab late at night getting atoms to move the right way in 2021.

Working with Madelyn Cain has reshaped the way I think about building collaborations. It is amazing how much we accomplished in the experimental and theoretical study of error-corrected algorithms, and opened a new frontier that has blossomed within our group and beyond. Maddie's relentless problem solving abilities coupled with her incredible diligence makes her an indispensable

asset and key driver in every project she touches.

Harry Zhou and I first met when I was an undergrad and we were working on NVs and we clicked from the beginning. Working with Harry taught me the real value in diversity of scientific approaches, where we wouldn't have tried some things if I weren't so keen to disobey rules, and we wouldn't have figured out that they always work in the general case if it wasn't for Harry reading such a broad range of QEC literature. Our dynamic has been extremely fun and productive, and I hope many more adventures are to come.

I would also like to thank the entire Lukin group, in particular the 'atom family' composed of atom array 1, atom array 2, and the atom cavity teams. It has been a great pleasure to work with multiple generations of incredible scientists in these teams, as well as for the broad range of theorists in and out of the group. I would also like to thank the Center for Ultracold Atoms and its extremely collaborative nature.

My undergraduate advisors at UCSB also had a tremendous impact on my career, Tengiz Bibilashvili and Ania Jayich. Tengiz - affectionately known as Dr. B - has an exceptionally broad physics knowledge, and he commits all of this breadth, depth, and creativity into his teaching. The success of the UCSB College of Creative Studies program is no accident. Dr. B taught me all of my foundational physics, and taught me what it means to really understand an idea. Ania taught me essentially everything I know about how to be a scientist - from experimental methods to writing papers. Ania had an incredible talent for throwing me into the deep end - or at least having it seem that way - while providing guidance in the background to help me succeed and grow. I am forever indebted to the undergraduate education I received at UCSB and the College of Creative Studies.

Finally, I'd like to thank my friends, family, and Maine. There are too many stories to tell, too many laughs we've had, and too many lessons learned, to even begin to try to summarize their impact on my life. Sitting here reflecting, my primary feeling is that I couldn't wish it would have been any different at all.

Author Contributions

The work presented in this thesis have been published and co-authored, as outlined below.

Chapter 2 contains content overlapping with the following works:

Sepehr Ebadi, Tout T Wang, Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Dolev Bluvstein, Rhine Samajdar, Hannes Pichler, Wen Wei Ho, Soonwon Choi, Subir Sachdev, Markus Greiner, Vladan Vuletić, Mikhail D Lukin. Quantum phases of matter on a 256-atom programmable quantum simulator. Nature 595, 227-232 (2021).

Harry Levine, Dolev Bluvstein, Alexander Keesling, Tout T Wang, Sepehr Ebadi, Giulia Semeghini, Ahmed Omran, Markus Greiner, Vladan Vuletić, Mikhail D Lukin. Dispersive optical systems for scalable Raman driving of hyperfine qubits. Physical Review A 105, 032618 (2022).

Dolev Bluvstein, Harry Levine, Giulia Semeghini, Tout T Wang, Sepehr Ebadi, Marcin Kalinowski, Alexander Keesling, Nishad Maskara, Hannes Pichler, Markus Greiner, Vladan Vuletić, Mikhail D Lukin. A quantum processor based on coherent transport of entangled atom arrays. Nature 604, 451-456 (2022).

Simon J Evered, Dolev Bluvstein, Marcin Kalinowski, Sepehr Ebadi, Tom Manovitz, Hengyun Zhou, Sophie H Li, Alexandra A Geim, Tout T Wang, Nishad Maskara, Harry Levine, Giulia Semeghini, Markus Greiner, Vladan Vuletić, Mikhail D Lukin. High-fidelity parallel entangling gates on a neutral-atom quantum computer. Nature 622, 268-272 (2024).

Daniel Bochen Tan, Dolev Bluvstein, Mikhail D Lukin, Jason Cong. Compiling quantum circuits for dynamically field-programmable neutral atoms array processors. Quantum 8, 1281 (2024).

Dolev Bluvstein, Simon J Evered, Alexandra A Geim, Sophie H Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J Gullans, Markus Greiner, Vladan Vuletić, Mikhail D Lukin. Logical quantum processor based on reconfigurable atom arrays. Nature 626, 58-65 (2024).

Dolev Bluvstein, Alexandra A Geim, et al., *in preparation*. Physical mechanisms for a fault-tolerant quantum architecture. (2025)

Chapter 3 contains content overlapping with the following works:

Iris Cong, Harry Levine, Alexander Keesling, Dolev Bluvstein, Sheng-Tao Wang, Mikhail D Lukin. Hardware-efficient, fault-tolerant quantum computation with Rydberg atoms. Physical Review X 12, 021049 (2022).

Dolev Bluvstein, Simon J Evered, Alexandra A Geim, Sophie H Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J Gullans, Markus Greiner, Vladan Vuletić, Mikhail D Lukin. Logical quantum processor based on reconfigurable atom arrays. Nature 626, 58-65 (2024).

Qian Xu, J. Pablo Bonilla Ataides, Christopher A Pattison, Nithin Raveendran, Dolev Bluvstein, Jonathan Wurtz, Bane Vasic, Mikhail D Lukin, Liang Jiang, Hengyun Zhou. Constant-overhead fault-tolerant quantum computation with reconfigurable atom arrays. Nature Physics 20, 1084–1090 (2024).

Madelyn Cain, Chen Zhao, Hengyun Zhou, Nadine Meister, J Ataides, Arthur Jaffe, Dolev Bluvstein, Mikhail D Lukin. Correlated decoding of logical algorithms with transversal gates. Phys. Rev. Lett. 133, 240602 (2024).

Dominik Hangleiter, Marcin Kalinowski, Dolev Bluvstein, Madelyn Cain, Nishad Maskara, Xun Gao, Aleksander Kubica, Mikhail D. Lukin, Michael J. Gullans. Fault-tolerant compiling of classically hard IQP circuits on hypercubes. arXiv: 2404.19005 (2024).

Hengyun Zhou, Chen Zhao, Madelyn Cain, Dolev Bluvstein, Casey Duckering, Hong-Ye Hu, Sheng-Tao Wang, Aleksander Kubica, Mikhail D Lukin. Algorithmic fault tolerance for fast quantum computing. arXiv:2406.17653 (2024).

Qian Xu, Hengyun Zhou, Guo Zheng, Dolev Bluvstein, J. Pablo Bonilla Ataides, Mikhail D Lukin, Liang Jiang. Fast and Parallelizable Logical Computation with Homological Product Codes. arXiv:2407.18490 (2024).

Dolev Bluvstein, Alexandra A Geim, et al., *in preparation*. Physical mechanisms for a fault-tolerant quantum architecture. (2025).

Madelyn Cain, Dolev Bluvstein, et al., *in preparation*. Fast correlated decoding of transversal logical algorithms. (2025).

Chapter 4 is derived from:

Dolev Bluvstein, Simon J Evered, Alexandra A Geim, Sophie H Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J Gullans, Markus Greiner, Vladan Vuletić, Mikhail D Lukin. Logical quantum processor based on reconfigurable atom arrays. Nature 626, 58-65 (2024).

Chapter 5 is derived from:

Dolev Bluvstein, Alexandra A Geim, et al., *in preparation*. Physical mechanisms for a fault-tolerant quantum architecture. (2025)

Chapter 6 has overlapping content with:

Josiah Sinclair, Joshua Ramette, Brandon Grinkemeyer, Dolev Bluvstein, Mikhail Lukin, Vladan Vuletić. Fault-tolerant optical interconnects for neutral-atom arrays. arXiv:2408.08955 (2024).

Chapter 7 is derived from:

Sepehr Ebadi, Tout T Wang, Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Dolev Bluvstein, Rhine Samajdar, Hannes Pichler, Wen Wei Ho, Soonwon Choi, Subir Sachdev, Markus Greiner, Vladan Vuletić, Mikhail D Lukin. Quantum phases of matter on a 256-atom programmable quantum simulator. Nature 595, 227-232 (2021).

Dolev Bluvstein, Ahmed Omran, Harry Levine, Alexander Keesling, Giulia Semeghini, Sepehr Ebadi, Tout T Wang, Alexios A Michailidis, Nishad Maskara, Wen Wei Ho, S Choi, Maksym Serbyn, M Greiner, V Vuletić, MD Lukin. Controlling quantum many-body dynamics in driven Rydberg atom arrays. Science 371, 1355-1359 (2021).

Giulia Semeghini, Harry Levine, Alexander Keesling, Sepehr Ebadi, Tout T Wang, Dolev Bluvstein, Ruben Verresen, Hannes Pichler, Marcin Kalinowski, Rhine Samajdar, Ahmed Omran, Subir Sachdev, Ashvin Vishwanath, Markus Greiner, V Vuletić, Mikhail D Lukin. Probing topological spin liquids on a programmable quantum simulator. Science 374, 1242-1247 (2021).

Sepehr Ebadi, Alexander Keesling, Madelyn Cain, Tout T Wang, Harry Levine, Dolev Bluvstein, Giulia Semeghini, Ahmed Omran, J-G Liu, Rhine Samajdar, X-Z Luo, Beatrice Nash, Xun Gao, Boaz Barak, Edward Farhi, Subir Sachdev, Nathan Gemelke, Leo Zhou, Soonwon Choi, Hannes Pichler, S-T Wang, Markus Greiner, V Vuletić, Mikhail D Lukin. Quantum optimization of maximum independent set using Rydberg atom arrays. Science 376, 1209-1215 (2022).

Tom Manovitz, Sophie H Li, Sepehr Ebadi, Rhine Samajdar, Alexandra A Geim, Simon J Evered, Dolev Bluvstein, Hengyun Zhou, Nazli Uğur Köylüoğlu, Johannes Feldmeier, Pavel E Dolgirev, Nishad Maskara, Marcin Kalinowski, Subir Sachdev, David A Huse, Markus Greiner, Vladan Vuletić, Mikhail D Lukin. Quantum coarsening and collective dynamics on a programmable quantum simulator. Nature 638, 86–92 (2025).

Dolev Bluvstein, Harry Levine, Giulia Semeghini, Tout T Wang, Sepehr Ebadi, Marcin Kalinowski, Alexander Keesling, Nishad Maskara, Hannes Pichler, Markus Greiner, Vladan Vuletić, Mikhail D Lukin. A quantum processor based on coherent transport of entangled atom arrays. Nature 604, 451-456 (2022).

Simon J. Evered, Marcin Kalinowski, Alexandra A. Geim, Tom Manovitz, Dolev Bluvstein, So-

phie H. Li, Nishad Maskara, Hengyun Zhou, Sepehr Ebadi, Muqing Xu, Joseph Campo, Madelyn Cain, Stefan Ostermann, Susanne F. Yelin, Subir Sachdev, Markus Greiner, Vladan Vuletić, Mikhail D. Lukin. Probing topological matter and fermion dynamics on a neutral-atom quantum computer. arXiv:2501.18554 (2025).

1

Introduction

1.1 Quantum error correction: History and background

The development of quantum error correction is one of the most important theoretical breakthroughs in quantum information science. In 1994 Peter Shor invented the factoring algorithm and was met with significant excitement [133]. The factoring algorithm was an example that quantum computers could solve an important problem that classical computers are not thought to be able to

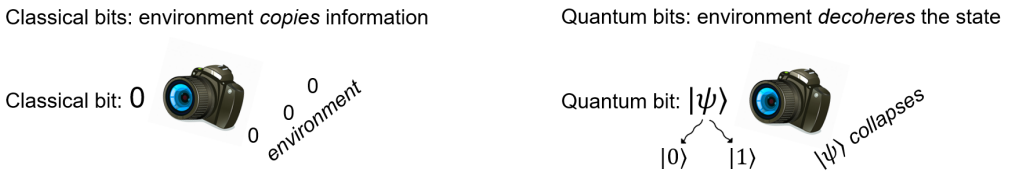


Figure 1.1: Decoherence upon interaction with the environment. Whereas classical information can become redundantly copied and thereby stabilized via its interaction with the environment, interaction with the environment leads to decoherence for a quantum bit.

do. However, there was also widespread skepticism. Around this time, experimenters had started exploring logic operations in the lab and had gate error rates of $\sim 10\%$ [67]. Estimates showed that factoring relevant numbers would require trillions of quantum gates, thereby requiring error rates below $\sim 10^{-12}$ [67,59]. Such a disparity made it seem that there is no conceivable way to use quantum systems for such a complex computation. These error rates should be compared to those of transistors in classical devices, which are well below 10^{-14} [159,82]. Are quantum computer builders just much less competent than classical computer builders? Not quite. Quantum information is fundamentally harder to store and operate on.

Consider for example taking a classical bit and allowing your friends to take pictures of it. This will then copy the bit onto many people's cameras, and *redundantly* stores it. This is what often happens when classical information interacts with its environment - it becomes *copied* and redundantly stored by its environment. This is in stark contrast to what would happen with a quantum bit $|\psi\rangle$. When your friends take pictures of $|\psi\rangle$, it will collapse the wavefunction, and probabilistically jump into either $|0\rangle$ or $|1\rangle$. Whereas classical information is copied when interacting with the environment, quantum information is decohered. So we can already see that quantum information is much more sensitive to interactions with its environment [144,92].

Another issue is that quantum operations are sensitive to small imperfections [92,93]. Consider trying to realize a quantum computation with a chain of spins which are interacting with some strengths J_i, and the computation is encoded in the output state. What happens if $J_i \to J_i + \varepsilon$? Then

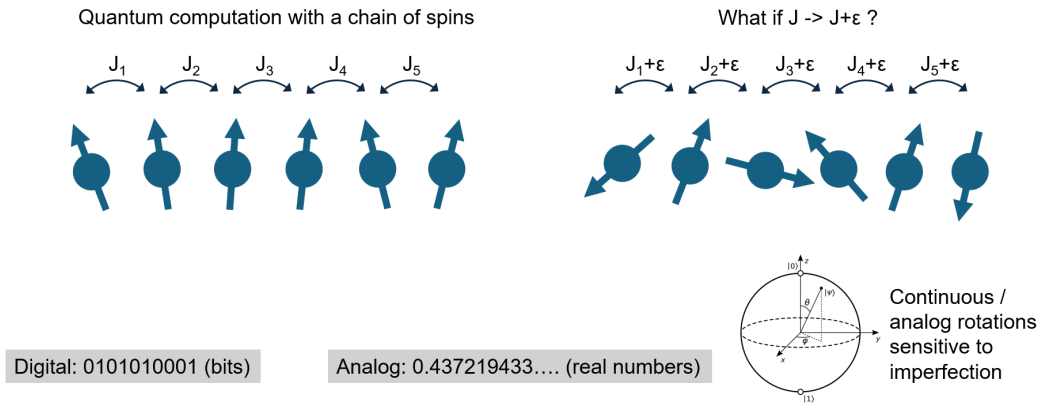


Figure 1.2: Computation with a chain of spins. Small imperfections in tuning parameters can lead to a wrong result as errors can accumulate chaotically. This is why we use digital classical computers, which leverage dissipation to lead to robust digital operation, as opposed to analog classical computers which require perfect evolution to avoid errors. Quantum information lives on a Bloch sphere, seemingly leading to the same challenge with correcting continuous numbers like in analog classical computers.

we know that the output state can be chaotically different[92]. This is further bolstered by deep reasons in complexity. It is, in fact, understood that (idealized) analog *classical* computing is far more powerful than digital classical computing as it can represent real numbers e.g. 0.43121810 81..., but in practice it is not idealized, and is not what we regularly use, because of the inability to correct errors[45,16,93]. Instead in digital classical computers information is represented as *bits* e.g. 00010101011100... The digital nature means that these can be intrinisically error-corrected and protected[91]. So, it somehow seems that the powerful aspect of our quantum computer - that a qubit can be in different superposition states on a Bloch sphere - is also why it can't work robustly due to inability to correct errors.

One then may wonder if it is possible to use error correction techniques to remove decoherence. Classical error correction works in a very simple way - you take a bit and copy it, e.g. $0 \rightarrow 000$. Now, if a single bit flips and the state goes to e.g. 001, then majority voting would reveal that the state was 000, and can be reset to 000, and the bit can be preserved.

Our story of correcting quantum errors hits another roadblock here. You cannot copy quantum

information! It is not possible to take a qubit and copy it, i.e. $|\psi\rangle \rightarrow |\psi\rangle\,|\psi\rangle\,|\psi\rangle$ is not allowed, the infamous *no-cloning* theorem. (The reverse of this process, for example, is clearly non-unitary as it would take $|\psi\rangle \rightarrow |0\rangle$, and so it destroys information). We then hit yet another roadblock - how do we even *check* that a qubit has had an error? If we look at the qubit (i.e. measure it), like in the above, then we will collapse its wavefunction! Between the no cloning theorem, and the fact that we cannot even look at the qubit to check for errors, it seems all but impossible that quantum error correction should be possible.

Remarkably, in 1995 Shor (shortly followed by Steane in 1996, and then various others) came up with a solution to quantum error correction [131,136]. Instead of copying quantum information, quantum error correction works with a very different physical mechanism. *Quantum entanglement* is used as a tool to *delocalize* an error-corrected *logical qubit* of information across many physical qubits. By being very spread out, the logical qubit becomes protected from the environment. More quantitatively, each physical qubit is so entangled with its neighbors that its density matrix is *fully mixed*, and measuring a physical qubit in any basis will return a 50/50 random result which contains no information [131]. So if the environment is now to measure a physical qubit of the block, it will gain identically no information about the underlying state. Moreover, although we cannot directly measure physical qubits to check for errors, we can measure *products* of physical qubits to form so-called *stabilizer checks*, which will allow us to check for errors without corrupting the underlying logical qubit information (more quantitatively, these checks *commute* with the logical operator) [131]. And remarkably, we can realize logic operations on this delocalized wavefunction, leading to unitary and analog-type evolution on the *logical level*, while removing errors via dissipation and error correction on the *physical level* [132]. These will be described more quantitatively and practically in the following sections.

The discovery of quantum error correction has extreme practical and fundamental implications on the development of large-scale quantum processing, on the theory of quantum information and

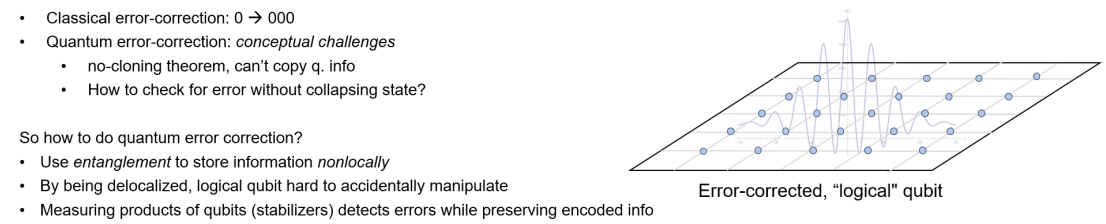


Figure 1.3: Quantum error correction. Quantum information cannot be copied, and cannot be measured without being decohered, making error correction seem impossible. Remarkably, error correction can be achieved by a different mechanism, in which entanglement delocalizes a logical qubit degree of freedom across many independent physical qubits.

quantum mechanics, and on complexity theory and computation. This is well-captured by a famous argument between Peter Shor and famous computer scientist Rolf Landauer, who was one of Shor and quantum computing's biggest skeptics. To paraphrase, the argument went as follows: Landauer publicly questioned Shor, asking "You just reinvented analog computing! We know analog computing doesn't work because we can't correct errors, what's new here?" Shor then replied "The computation is analog, but the errors are *digital*." It is, in fact, the wave-particle duality that makes quantum computers so powerful. The logical qubits can represent an "analog"-type state by being in superposition states, but measuring stabilizer checks causes the system to collapse errors into being *digital*, allowing their correction [131,132]. Analog classical computing does not work in the same way because of the absence of projective measurement. This thesis comes at a time when we are so lucky to be able to experiment with some of these core ideas for the first time.

1.2 QEC basics

We will now introduce some basic concepts in quantum error correction that will be helpful later in this thesis.

Code distance and decoding. $[[n, k, d]]$ notation describes a code with a number of physical qubits n, a number of logical qubits k, and a code distance d. The code distance d sets how many errors a

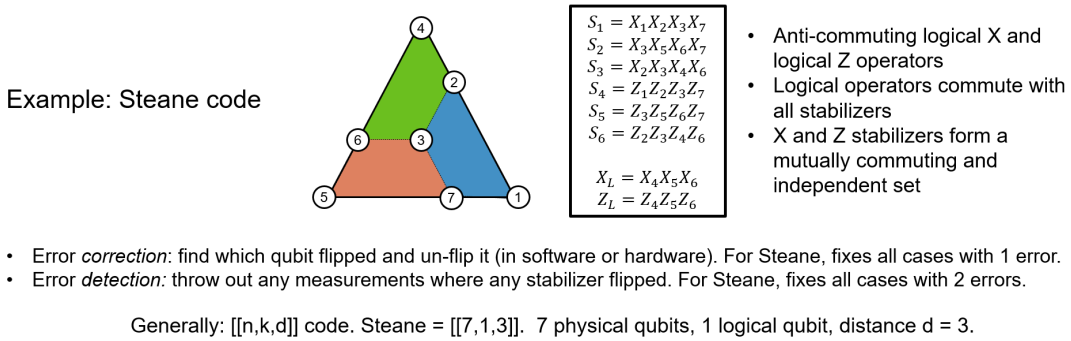


Figure 1.4: Example QEC code: the [[7,1,3] Steane code.

code can detect or correct, and is the minimum Hamming distance between valid codewords (logical states), i.e. the weight of the smallest logical operator[61]. In the case of the two-dimensional surface and color codes studied here, d is equivalent to the linear dimension of the system[57]. A fully explicit example of a quantum code, the Steane [[7,1,3]] code is shown in Fig. 1.4[136].

Following this definition, quantum codes of distance d can detect any arbitrary error of weight up to $d-1$. Such errors cause stabilizer violations, indicating that errors occurred. The correction procedure brings the system back to the closest logical state (codeword); thus, if more than d/2 errors occur, the resulting state may be closer to a codeword different from the initial one, resulting in a logical error[61]. For this reason, codes of distance d can correct any arbitrary error of weight up to $(d-1)/2$ (rounded down if d is even[57]). The process of decoding refers to analyzing the observed pattern of errors and determining what correction to apply to return back to the original code state and undo the physical errors created.

Threshold. As the size of an error correcting code and the corresponding code distance is increased, so are the opportunities for errors to occur as the number of physical qubits increases. This leads to a threshold behavior in quantum error correction: if the density of errors p is above a (possibly circuit-dependent) characteristic error rate p_{th}, then increasing code distance will worsen performance. However, if $p < p_{\text{th}}$, then increasing code distance will improve performance[57]. Theoreti-

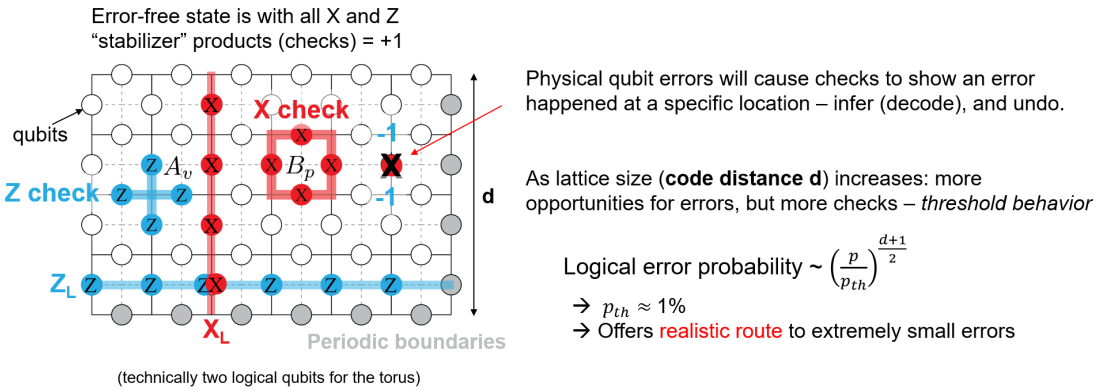


Figure 1.5: Kitaev toric code. The toric code [86] is defined by mutually commuting X and Z stabilizer checks on a 2D lattice of qubits (with periodic boundaries), and X_L and Z_L logical operators propagating vertically and horizontally across the torus. Qubit flips can be inferred via the stabilizer check measurements. As code distance d is increased, the logical error can be exponentially suppressed.

cally, since one requires $(d+1)/2$ errors to create a logical error, the logical error rate will be exponentially suppressed as $\propto (p/p_{\text{th}})^{(d+1)/2}$ at sufficiently low error rates [57]. In the surface code in particular (toric code with open boundary conditions), the logical error rate $P_L \approx 0.03(p/p_{\text{th}})^{(d+1)/2}$ where $p_{\text{th}} \approx 1\%$ [44].

Fault-tolerance. Later in this thesis, we will explore *fault-tolerance* and the role it plays in realizing *computation* with logical qubits. A common definition of fault-tolerance in quantum circuits [61] (which we use here) is that a weight-1 error (i.e. an error affecting one physical qubit), cannot propagate into a weight-2 error (now affecting two physical qubits) within a logical block. This property implies that errors cannot spread within a logical block, and thereby prevents a single error from growing uncontrollably and causing a logical error.

Critically, if an operation is fault-tolerant, then it will allow us to do computation without spreading errors within the block. This will be discussed in detail later, in particular due to the fault-tolerance of *transversal* gates [132].

1.3 Estimates of required resources: QEC is challenging!

If we plug in numbers for the surface code considered above, we find a bleak reality: we need *a lot* of qubits. The typical problems we have in mind require 1000s of logical qubits with logical error rates at the scale of 10^{-10} and below[57,113,59]. With our hardware, we generally expect to be able to get to physical fidelities at the scale of 99.9%, which would correspond to $p/p_{\text{th}} = 0.1$ on the toric code. If we want to get to the scale of 10^{-10} logical error and below, we need code distances of roughly $d = 20$, which would correspond to 400 data qubits just to store one logical qubit! Including ancilla qubits, we are in the ballpark of 1000 physical qubits for a single logical qubit. To tackle large-scale problems with 1000s of logical qubits, we need millions of physical qubits[59]. That is nowhere close to where we are with existing hardware.

The problem of needing many qubits is greatly further compounded by the fact that modern quantum hardware has been built in a qubit-by-qubit approach, where each qubit has its own control lines. Imagine if you built a large-scale processor like this, and you had one million microwave generators all simultaneously sending separate, nearly identical Hadamard pulses to a million different qubits! This so-called "wire problem" poses a significant challenge to large-scale control with existing systems. The other challenge is that once this logical qubit degree of freedom is delocalized, it also becomes very hard to operate on. For these reasons, before the experiments in this thesis, studies in the field had been focused to 1-2 logical qubits and operations. Seemingly, we have a very far way to go! But, things will change faster than we think.

1.4 An inflection point

There has been a lot of "hype" about quantum computers in the past decade. Many people feel like it's "always another 10 years away" or that the progress is very slow and linear. But there are two key reasons why it is likely to change dramatically in the coming years. (1) It is becoming exceptionally

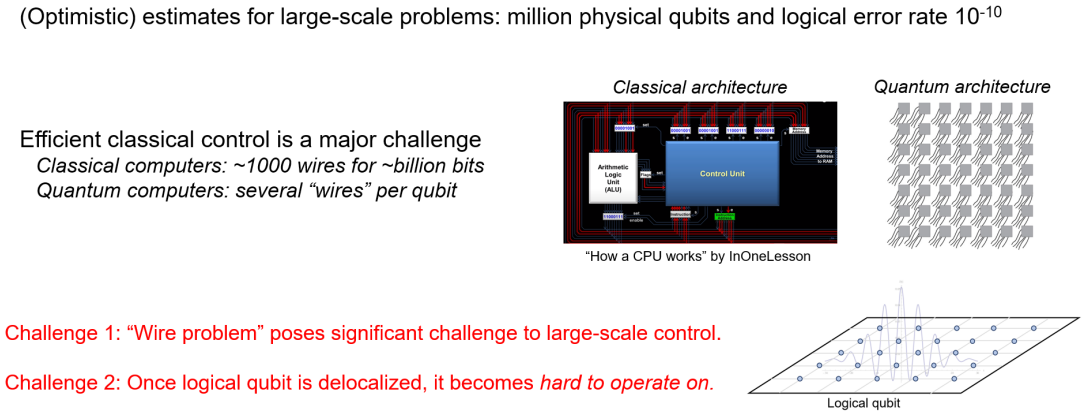


Figure 1.6: Wire problem. The large overheads of quantum error correction are greatly further compounded by the fact that modern quantum hardware has been built qubit-by-qubit, instead of with large-scale parallelism in mind.

clear from experiments that quantum error correction *works*. We see now that we can exponentially reduce errors by increasing the size of an error correction code, and that we can use error correcting codes for computations which are more precise than using physical qubits. We are below characteristic thresholds, and quite literally, the improvement will continue to be exponential from here. (2) Building quantum systems designed for logical-qubit computation can be very different from building systems designed for physical-qubit computation. We do not need to control each individual physical qubit uniquely, and can control all of the physical qubits of a logical qubit *in parallel*. If all the atoms are identical, and they all are supposed to realize the exact same operation with the exact same pulse of light, why does it need to be so hard to control millions of qubits?

1.5 Overview of thesis

In this thesis, we will describe methods we developed for quantum computation with neutral atoms and their use for error-correctd computation. In Chapter 2 we will describe all the key aspects of realizing quantum computation with neutral atoms. This description is meant to be accessible and informal, and additional details can be found in the papers referenced within. In Chapter 3 we will

provide a description of key ideas in quantum error correction, in particular for the realization of error-corrected algorithms. In Chapter 4 we will describe the first implementation of error-corrected algorithms, and in Chapter 5 further studies we have done on the physical mechanisms for a fault-tolerant quantum architecture. Chapter 6 will provide an outlook for the next steps. Chapter 7 will then discuss applications of quantum processors we studied with analog and gate-based methods.

2

Quantum computing with neutral atoms

2.1 Atoms

Neutral atoms are, although a cliche statement, intrinsically very scalable on account of: (1) there are a lot of them and (2) they are all identical. This has been known and commented on many times [120,27,121], and although this observation is not the contribution of this thesis, it is the most useful aspect that we will leverage. In designing our neutral atom quantum computer and the tools

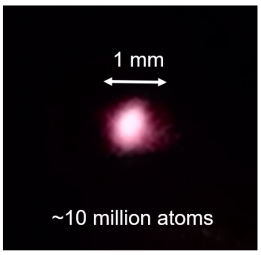


Image of atoms in a magneto optical trap (MOT)

Figure 2.1: Image of atoms in a magneto optical trap (MOT). There are millions of cold, identical neutral atoms in the cloud. Number of atoms is not an issue - the key challenge is efficient classical control.

for controlling the atoms, this is the core principle that we will keep returning back to.

All of our experiments start with large clouds of laser cooled atoms. We use ^{87}Rb atoms in particular - arguably the simplest atom atomic physicists like to use, because it is *hydrogenic* meaning it has only a single valence electron and thus has simple structure, but being sufficiently far down the periodic table such that the optical wavelengths are favorable. Figure 2.1 shows a picture of our magneto optical trap (MOT). This works with 6 counterpropagating laser beams that cool the atoms via *doppler cooling* and *polarization gradient cooling*, and are accumulated into a $\sim$ 1 mm region via the use of a strong magnetic field gradient that makes an energy minimum [115,106]. These pictures are very illuminating of the fact that there are many cold, identical atoms. So the number of atoms is rarely the issue, it's the *efficient classical control.*

Figure 2.2 shows how we achieve quantum circuits with neutral atoms at a high-level, and Figure 2.3 shows a technical diagram of our neutral atom quantum processor with key optical and electronic controls. At a high level, atoms are trapped inside of a vacuum chamber in optical tweezers, they are moved around and zapped by various types of lasers to control their quantum states, and their states are detected by taking pictures with a camera. With the techniques developed in this thesis we will be able to use simple off-the-shelf components to realize advanced quantum computations. In the sections that follow, we will describe how we trap, arrange, image, create superposi-

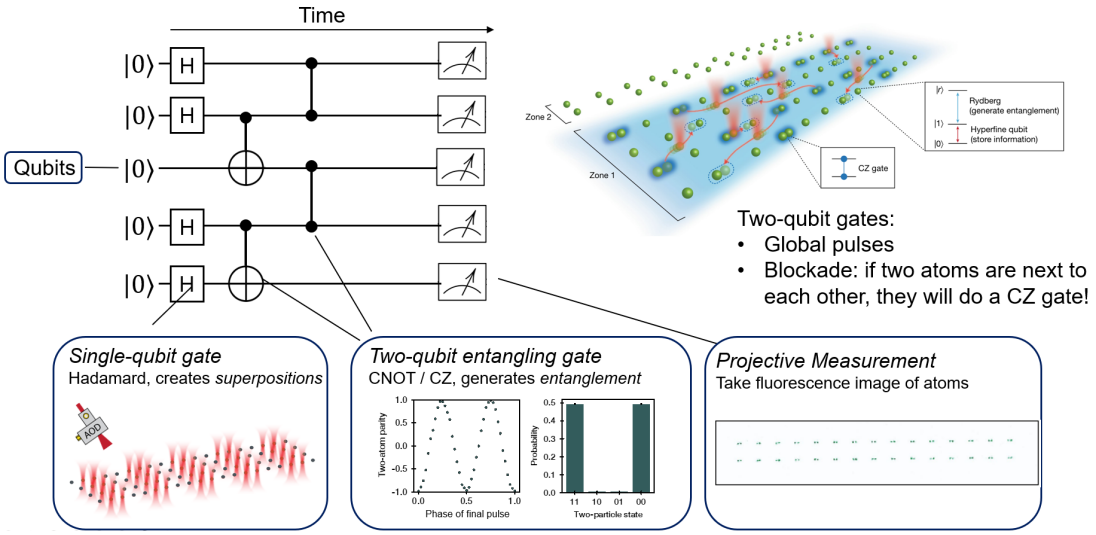


Figure 2.2: Key aspects of computation with neutral atoms. Individual atoms are trapped in individual optical tweezers. Single-qubit gates are realized via locally addressable Raman excitation, two-qubit entangling gates are realized via Rydberg interactions leveraging global pulses and blockade, and projective measurement is realized via fluorescence.

tions, entangle our atoms, as well as non-destructively measure and reinitialize them.

2.2 Trapping atoms in optical tweezers

Neutral atoms are dielectric paticles, meaning that they pick up a dipole $\vec{d}$ in an electric field $\vec{E}$. Functionally, this allows them to be trapped in optical potentials, as the dipole $\vec{d} \propto \alpha\vec{E}$ yields an energy $V \propto -\vec{d} \cdot \vec{E} = -\alpha|\vec{E}|^2$. Optical tweezers are focused beams of light, where the optical intensity $\propto |\vec{E}|^2$ is the highest at the focus, and consequently is the energy minimum for the atom[6]. Procedurally, this is remarkably simple: take a beam of light and focus it to a point, and you can trap atoms there and move them around by moving the light around. My favorite derivation of optical tweezer properties is Wikipedia's.

When we stick our optical tweezer into our MOT, atoms pile into the tweezer because they are cooled to their lowest energy. One of the most remarkable "accidents" of our experiments, is that when the tweezer is so tightly confining to the atoms, there are molecular resonances that the

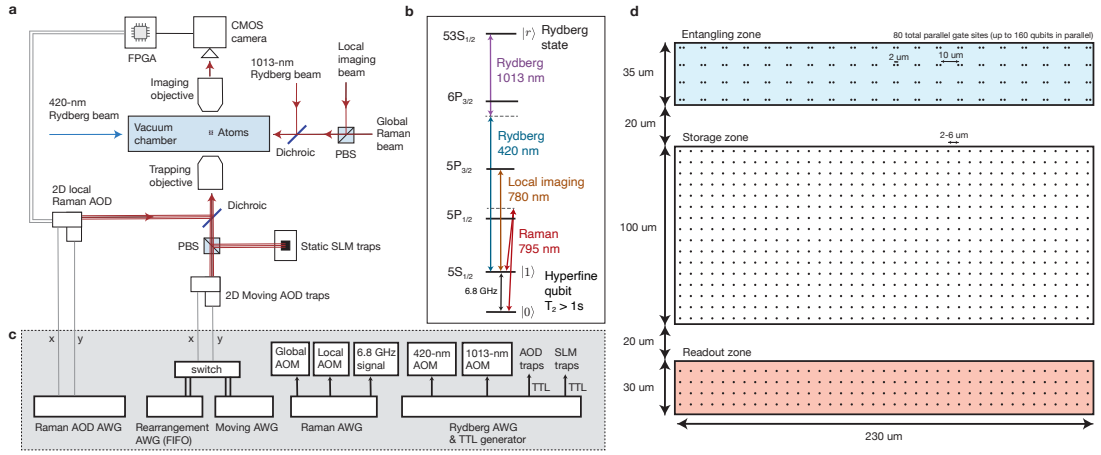


Figure 2.3: Neutral atom quantum computer architecture. **a,** Experimental layout, featuring optical tools including static SLM and 2D moving AOD traps, global and local Raman single-qubit laser beams, 420-nm and 1013-nm Rydberg beams, and imaging system for both global and local imaging. **b,** Level structure for ^{87}Rb atoms, with the relevant atomic transitions employed in this work. **c,** Control infrastructure used for programming quantum circuits, featuring several arbitrary waveform generators (AWGs). In particular, the moving and Raman 2D AODs are each controlled by two waveforms (one for x axis and one for y axis). An additional AWG is used in first-in-first-out (FIFO) mode for rearrangement before the circuit begins, and then the moving AOD control is switched to the Moving AWG. See Ref.[49] for additional SLM and pre-circuit rearrangement details, Ref.[52] for additional Rydberg AWG details and Rydberg excitation details, Refs.[95,19] for additional Raman laser and microwave control infrastructure details, and Ref.[19] for additional moving AWG details. All AWGs (other than rearrangement AWG) are synchronized to $<$ 10 ns jitter. During Rydberg gates the traps are briefly pulsed off by a TTL. The FPGA processes images from the camera real-time and in this work sends control signals to the Raman 2D AOD for local single-qubit control. **d,** Example array layout featuring entangling, storage, and readout zones. Zones can be directly reprogrammed and repositioned for different applications, as well as specific tweezer site locations. Tweezer beams and local Raman control are projected from out-of-plane. The entire objective field-of-view is 400-μm diameter. During two-qubit Rydberg gates, we place atoms $\lesssim 2$ μm apart within a gate site, and gate sites are separated such that atoms in different gate sites are no closer than 10 μm during the gate. At our present $n = 53$ and two-photon Rabi frequency of 4.6 MHz, the blockade radius is roughly 4.3 μm, such that adjacent atoms are well-within blockade and distant atoms are well-outside blockade.

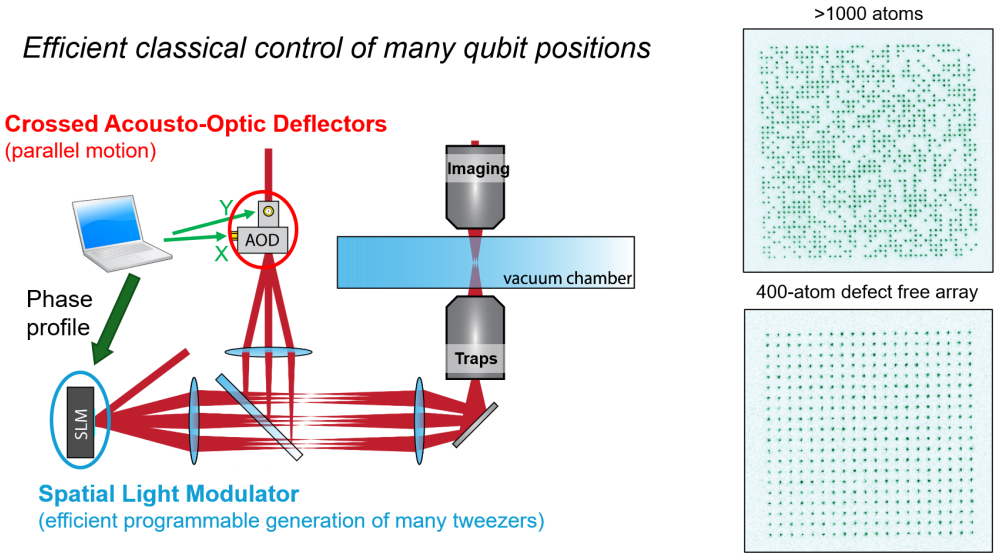


Figure 2.4: Optical control via tweezers. Static tweezer traps are generated via a spatial light modulator (SLM) into programmable geometries. Crossed acousto-optic deflectors (AOD) are used for shuttling atoms both for rearrangement and mid-circuit transport.

atoms get excited to via the laser cooling illumination, which causes *pairs of atoms to eject from the tweezer* [31]. This means that after dispersing the MOT cloud by turning off the gradient coils, the atoms parity-project into having either 0 or 1 atom after several ms of illumination. This process thereby loads *a single atom* into each optical tweezer with roughly 50% probability.

2.3 Building atomic arrays

Additional details for the description in this section can be found in Ref. [49].

Here our key principles of efficient classical control shine. We take a *single* laser beam and impinge it onto a liquid-crystal-on-silicon *spatial light modulator* (SLM) which creates a programmable arrangement of tweezers [10]. We start with a uniform Gaussian beam and the SLM imprints a phase pattern onto it, and then lenses take the Fourier transform of this phase pattern. The phase pattern that leads to the Fourier-transformed-grid of tweezer spots is calculated via the weighted Gerchberg-

Saxton algorithm and allows us to position the tweezers in arbitrary spots[49].

To see where the atoms are and which tweezers have been loaded, we take a camera image of the atoms. We do this by zapping them with the same light as the cooling lasers, and these simultaneously cool the atoms while scattering light off the atoms. We collect this scattered light via a 0.65-NA microscope objective and a CMOS camera from Hamamatsu.

At this point we now know where all the atoms are and can sort the atoms into defect-free grids. We do this using another optical tool known as an *acousto optic deflector* (AOD)[51,10]. This is a crystal with a transducer which uses sound waves to deflect light. Conceptually the dynamically programmable sound wave forms a modulation in the index of refraction and forms a dynamically programmable diffraction grating that deflects the light. In another language, photons of a given momentum impinge on the crystal and are then deflected by high-momentum phonons, leading to deflection of the light by an angle proportional to the frequency of the phonon. By superimposing multiple frequencies into the crystal we can deflect multiple beams with the same device, and the positions of all of these can be modified by modulating the frequency of sound. This converts our input beam into a programmable 1D line of beams.

By adding a second AOD in the orthogonal direction, we can now take each beam in the X direction and multiply it into a pattern in the Y direction. This now allows us to create large rectangular grids of light using two voltage waveforms $V_x(t)$ and $V_y(t)$. Critically, this is done just using two voltage waveforms - one for X and one for Y - and leveraging frequency multiplexing as well as the Cartesian product between the two AODs for controlling large grids.

To sort the atoms into a defect-free grid we just now need to calculate $V_x(t)$ and $V_y(t)$ and stream it to the AOD which we do via an arbitrary waveform generator (AWG) from Spectrum Instrumentation. We have upgraded our rearrangement algorithm in various ways but conceptually it is the same as has been discussed in Ref. [49], and many superior rearrangement algorithms have been developed in parallel in the literature.

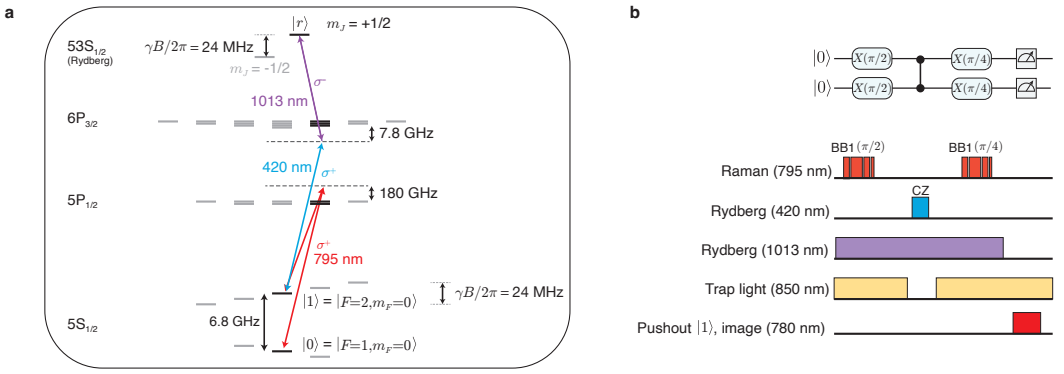


Figure 2.5: Atomic level diagram and pulse sequence. a, Level diagram showing key levels of ^{87}Rb used in our quantum circuits. The clock states, $|0\rangle$, $|1\rangle$ are the qubit states used in this work. Excitation to the Rydberg state between $|1\rangle$ and $|r\rangle$ is carried out by a two-photon transition driven by 420-nm and 1013-nm lasers. Single-qubit rotations are realized with an amplitude-modulated 795-nm laser that drives Raman rotations between the $m_F = 0$ clock states. A DC magnetic field of 8.5 G is applied throughout this work. **b,** Example pulse sequence, here for making a $|\Phi^+\rangle$ Bell state between two qubits. Traps are pulsed off for a few hundred ns during the Rydberg gate to avoid both anti-trapping of the Rydberg state and inhomogeneous light shifts that broaden the transition, and the ground state atoms are then recaptured for roughly $\sim 3\mu$s between consecutive gate applications.

We can rearrange the atoms into various patterns in this way, from square grids to honeycomb arrangements to entirely random arrangements[49]. This convenient way of creating programmable lattices was instrumental in our studies of programmable quantum simulation, discussed toward the end of this thesis.

2.4 Atomic level structure and quantum operations

Placing the atoms into fun arrangements and making atom pictures is certainly amusing but it does get old after a while (if you ever saw an atomica "Happy Birthday" message or an atomica logo of a funding agency - a grad student had to make that!). We would like to use our atoms for processing quantum information. There are many levels within the atom and they can be used for various different purposes. A full level diagram is shown in Figure 2.5.

Conceptually there are three key states that we will use in our quantum information processing[120]. There are two *qubit states* which are $m_F = 0$ clock states (i.e. magnetic-field-insensitive)

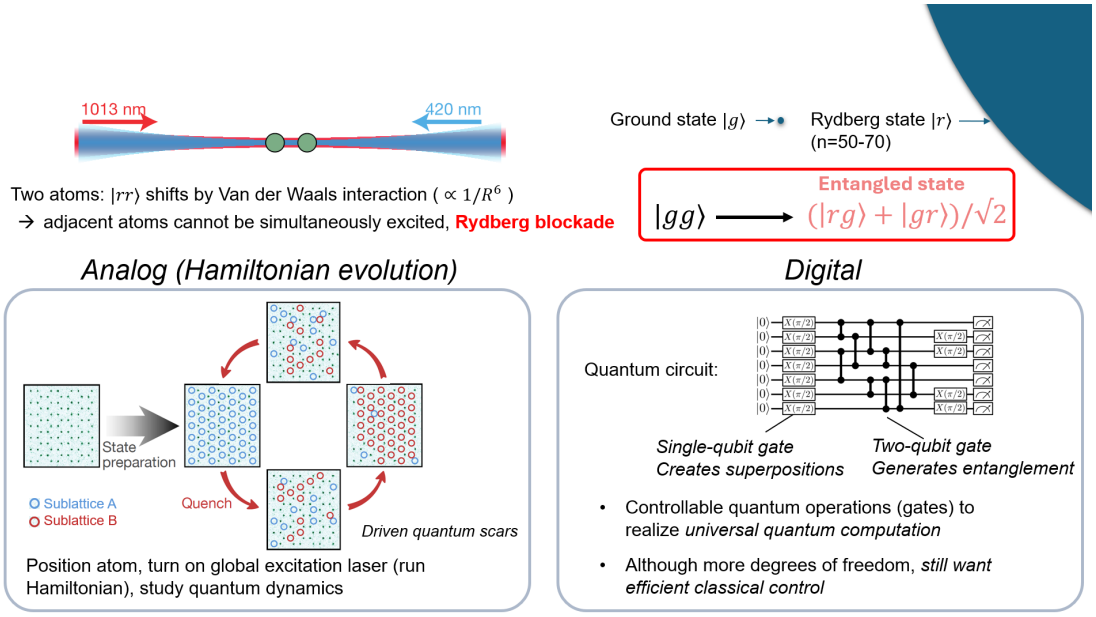


Figure 2.6: Analog and digital processing. Rydberg blockade mechanism generates entanglement in a way which is insensitive to specific interaction strength. Using this for analog processing we studied quantum simulations and quantum many-body phenomena ranging from many-body scars [20], quantum phase transitions [49], quantum coarsening [103], quantum spin liquids [127], and combinatorial optimization [48]. Digital processing has more nuance, but we would still like to develop methods of efficient classical control.

inside the hyperfine ground state manifold which we use for *storing* quantum information and performing *single-qubit gates*. These are magnetically insensitive spin states that are centimeters away from the macroscopic environment, and such can have extremely good properties (several seconds of coherence time are straightforwardly achieved [102,81,107,100]). We also use one *Rydberg state* which we excite to via two-photon excitation which generates *entanglement* between our atoms and is used for *two-qubit gates*.

In this thesis these states are used in two distinct ways. The first is *analog quantum simulation*. In this approach, atoms are placed into programmable lattices and then global Rydberg laser illumination causes the atoms to interact (in this configuration there is only one hyperfine state used, i.e. the ground state $|g\rangle \equiv |1\rangle$) [13,143,122]. This turns on Hamiltonian dynamics between the atoms

and results in rich many-body physics. Using this we studied quantum simulations and quantum many-body phenomena ranging from many-body scars[20], quantum phase transitions[49], quantum coarsening[103], quantum spin liquids[127], and combinatorial optimization[48]. This is conceptually very simple, as it simply relies on a global beam illumination yet resuls in very complex quantum evolution. This approach to studying quantum simulation is remarkably simple and powerful, and examples will be described in the later chapters of this thesis.

The other way these states are used is for *digital processing*, which will be the focus of this thesis. Digital processing tends to be significantly more nuanced than analog simulation - but somehow we still need to find ways to maintain our efficient classical control. Otherwise how could we do these experiments with a few grad students? The subsequent sections describe ways we realize our various programmable quantum logic using efficient classical control.

2.5 Programmable quantum circuits via atom reconfiguration

Additional details for the description in this section can be found in Ref.[19].

Two adjacent atoms in Rydberg states have a very strong Van der Waals interaction falling off as $1/r^6$, and in-practice can be several GHz at a several micron spacing that we can conveniently position atoms optically. The very strong interaction (relative to the excitation Rabi frequency of several MHz) means in practice that excitation of both atoms into the Rydberg state is forbidden: this is the so-called *Rydberg blockade* phenomenon[79,145]. If we place two ground state atoms $|g\rangle$ next to each other, i.e. the state $|gg\rangle$, then when we try to globally excite the atoms into the Rydberg state, they will both try to go into the Rydberg state $|r\rangle$, but they can't both simultaneously go and so will fight and go into the entangled superposition $\frac{1}{\sqrt{2}}(|gr\rangle + |rg\rangle)$. This generates entanglement in a way that uses *global pulses*, and importantly is *insensitive* to the specific interaction strength between the atoms, as it is essentially infinite and it is just the absence of excitation into $|rr\rangle$ that generates

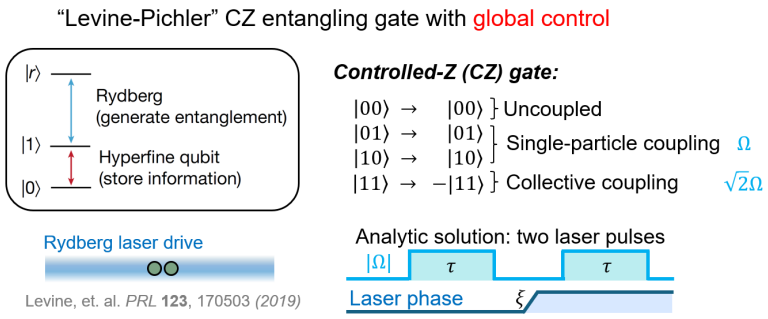


Figure 2.7: Levine-Pichler gate. The Levine-Pichler gate generates entanglement via two fast global Rydberg pulses with a phase jump in-between.

the entanglement. This is not yet a digital gate operation, but is the core underlying mechanism of robust entanglement.

We have Harry Levine to thank for convenient ways of entangling our atoms. He and Hannes Pichler (although Harry always insists it was Hannes, and Hannes always insists it was Harry) came up with the famous Levine-Pichler gate[96]. This is a very important result that allows high-fidelity entangling operations with just two simple global pulses. In 2019 several months before I first started my PhD, Harry and co had just used this to achieve gate fidelities of 97.5%, where the previous record was previously $\lesssim$ 80%. We now use improved versions of this core gate idea that are just one clean global pulse that regularly operates at 99.6% fidelity and will likely see improvement toward 99.9% in coming years. This will be described in the next section. But for now, the key takeaway is that we can generate controlled-Z (CZ) entangling gates between our hyperfine qubits simply by doing one simple, global pulse of the Rydberg entangling laser.

With the use of global pulses and Rydberg blockade, if two atoms are within a blockade radius they will entangle, and if they are outside the blockade radius, they won't. This now gives us a remarkably simple approach to program complex circuits simply via the atomic position. As discussed in the prior section, using optical tweezers and commercial tools of optical multiplexing such as SLMs and AODs, we have efficient methods of controlling the positions of our atoms. Now, simply

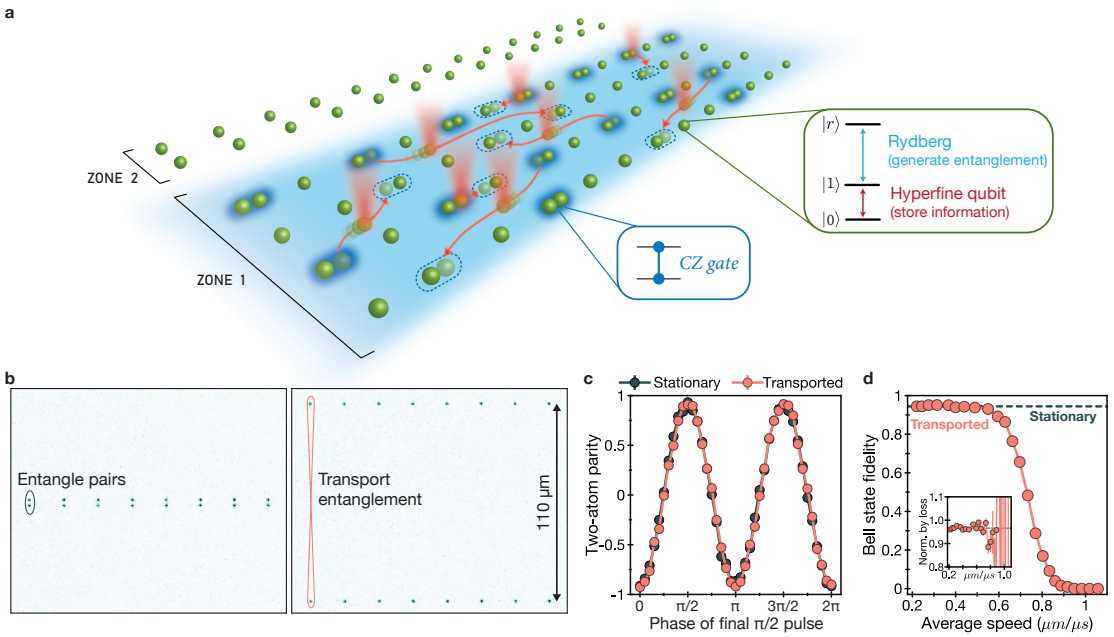


Figure 2.8: Quantum information architecture enabled by coherent transport of neutral atoms. a, In our approach, qubits are transported to perform entangling gates with distant qubits, enabling programmable and nonlocal connectivity. Atom shuttling is performed using optical tweezers, with high parallelism in two dimensions and between multiple zones allowing selective manipulations. Inset shows the atomic levels used: the $|0\rangle$, $|1\rangle$ qubit states refer to the $m_F = 0$ clock states of ^{87}Rb, and $|r\rangle$ is a Rydberg state used for generating entanglement between qubits. **b,** Atom images illustrating coherent transport of entangled qubits. Using a sequence of single-qubit and two-qubit gates, atom pairs are each prepared in the $|\Phi^+\rangle$ Bell state, and are then separated by 110 μm over a span of 300 μs. **c,** Parity oscillations indicate that movement does not observably affect entanglement or coherence. For both the moving and stationary measurements, qubit coherence is preserved using an XY8 dynamical decoupling sequence for 300 μs. **d,** Measured Bell state fidelity as a function of separation speed over the 110 μm, showing that fidelity is unaffected for a move slower than 200 μs (average separation speed of 0.55 μm/μs). Inset: normalizing by atom loss during the move results in constant fidelity, indicating that atom loss is the dominant error mechanism (see Methods for details).

with global laser pulses, this maps into efficient realization of complex quantum circuits. Atoms are positioned, the global entangling laser is pulsed, and then the atoms are reconfigured into a new configuration before a subsequent laser pulse. Remarkably, all the complex reconfiguration is done via a single two-channel AWG card streaming into a two-channel AOD. Broadly, this allows us to program highly complex quantum circuits by programming the two voltage controls of $V_x(t)$ and $V_y(t)$ and pulsing the Rydberg laser at the right time.

Figure 2.8 demonstrates our ability to transport qubits across large distances while preserving entanglement and coherence [14,94,156]. We initialize pairs at an atom-atom distance of 3 μm (Fig. 1b) and then create a Bell state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ in the hyperfine basis (Methods)[96]. To probe the resulting entangled-state fidelity, we apply an additional $\pi/2$ pulse with a variable phase that results in oscillations of the two-atom parity $\langle \sigma_1^z \sigma_2^z \rangle$ (Fig. 1c)[96]. We then repeat this experiment, but now move the atoms apart by 110 μm before applying the final $\pi/2$ pulse. Our transport protocol is optimized to suppress heating and loss by implementing cubic-interpolated atom trajectories (Methods), and is further accompanied by an 8-pulse XY8 robust dynamical decoupling sequence[64] to suppress dephasing. The resulting parity oscillations indicate that two-atom entanglement is unaffected by the transport process[14,46]. Performing this experiment as a function of movement speed shows that fidelity remains unchanged until the total separation speed becomes $> 0.55\,\mu\text{m}/\mu\text{s}$, corresponding to the onset of atom loss (Fig. 1d). We note that the entanglement transport in Figure 1b corresponds to moving quantum information across a region of space that can in principle host $\sim$ 2000 qubits (at an atom separation of 3 μm), on a timescale corresponding to $< 10^{-3}\ T_2$, directly enabling applications in large-scale quantum information systems.

To exemplify the ability to generate nonlocal connectivity between qubit arrays in parallel, we carry out preparation of entangled graph states: a large class of useful quantum information states, with examples ranging from GHZ states and cluster states to quantum error correction codes[69]. As an example, Fig. 2a demonstrates preparation of a 1D cluster state, a graph state defined by a linear

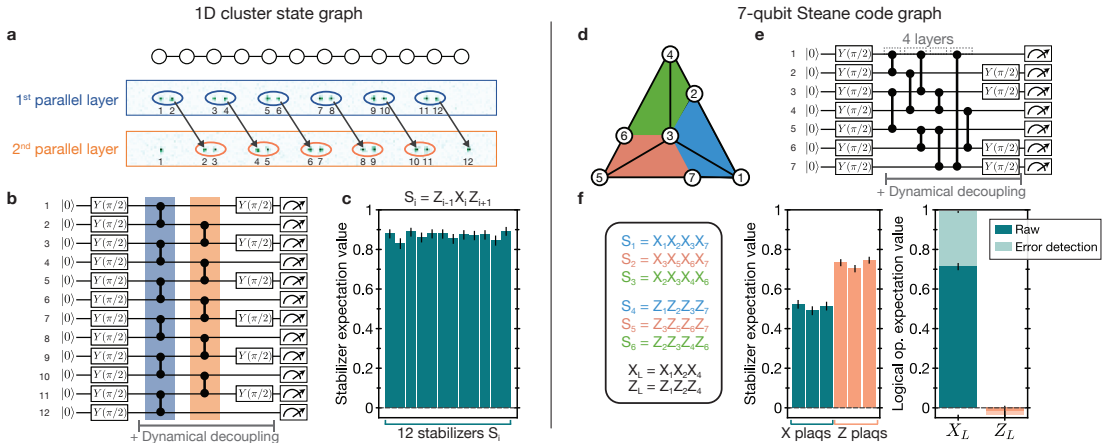


Figure 2.9: 1D and 2D graph states using dynamic entanglement transport. a, Generation of a 12-atom 1D cluster state graph, created by initializing all qubits (vertices) in $|+\rangle$ and applying controlled-Z gates on the links (edges) between qubits. Atom images show the configuration for the first and second gate layers. **b,** Quantum circuit representation of the 1D cluster state preparation and measurement. Dynamical decoupling is applied throughout all quantum circuits. **c,** Raw measured stabilizers of the resulting 1D cluster state, given by $S_i = Z_{i-1}X_iZ_{i+1}$ (X_1Z_2 and $Z_{11}X_{12}$ for the edge qubits). **d,** Graph state representation of the 7-qubit Steane code (colors represent stabilizer plaquettes). **e,** Circuit for preparing the Steane code logical $|+\rangle_L$ state, performed in four parallel gate layers. **f,** Measured stabilizers and logical operators after preparing $|+\rangle_L$. Error detection is done by postselecting on measurements where all stabilizers are +1. For both the 1D cluster state and Steane code, the stabilizers and logical operators are measured with two measurement settings (see text). Error bars represent 68% confidence intervals.

chain of qubits with edges between neighbors [117,124]. To realize this state, we perform one global, parallel layer of CZ gates on adjacent atom pairs, move half the atoms to form new pairs, and then perform another parallel layer of CZ gates (Fig. 2a pictures and Fig. 2b circuit). To probe the resultant twelve-qubit cluster state we measure the stabilizer set $\{S_i\} = \{Z_{i-1}X_iZ_{i+1}\}$ through readout in two measurement settings, given by a local $\pi/2$ rotation on either the odd or even sublattice before projective measurement [141]. We achieve the local rotation by moving one sublattice of qubits to a separate zone and then performing a rotation on the unmoved qubits with a homogeneous beam illuminating the experiment zone (Fig. 2.8a).

However, our approach crucially allows for *nonlocal* connections between qubits. We also used these nonlocal, parallel operations for *quantum error correction* - but this will be discussed in-detail in the next chapter. For now, the takeaway is that this approach allows for *efficient, parallel control* leveraging global laser pulses and tweezer-based reconfiguration of atoms.

The subsequent sections will describe how we realize programmable and high-fidelity digital processing within this reconfigurable atom array approach. The key design principle is parallel, efficient classical control, which as described in the next chapter is extremely well-suited to quantum error correction.

In terms of technical implementation, there are many rules on what moves the AODs can and cannot conveniently do, and what are the best ways of structuring circuits in a way that ensures dynamical decoupling. A description of this is provided at the end of the chapter.

2.6 High-fidelity two-qubit gates

Additional details for the description in this section can be found in Ref. [52].

When we first did the above quantum processing experiments, our circuit fidelity was limited by the two-qubit gate fidelity of 97.5%. Error correction has a threshold of roughly 99% on the two-

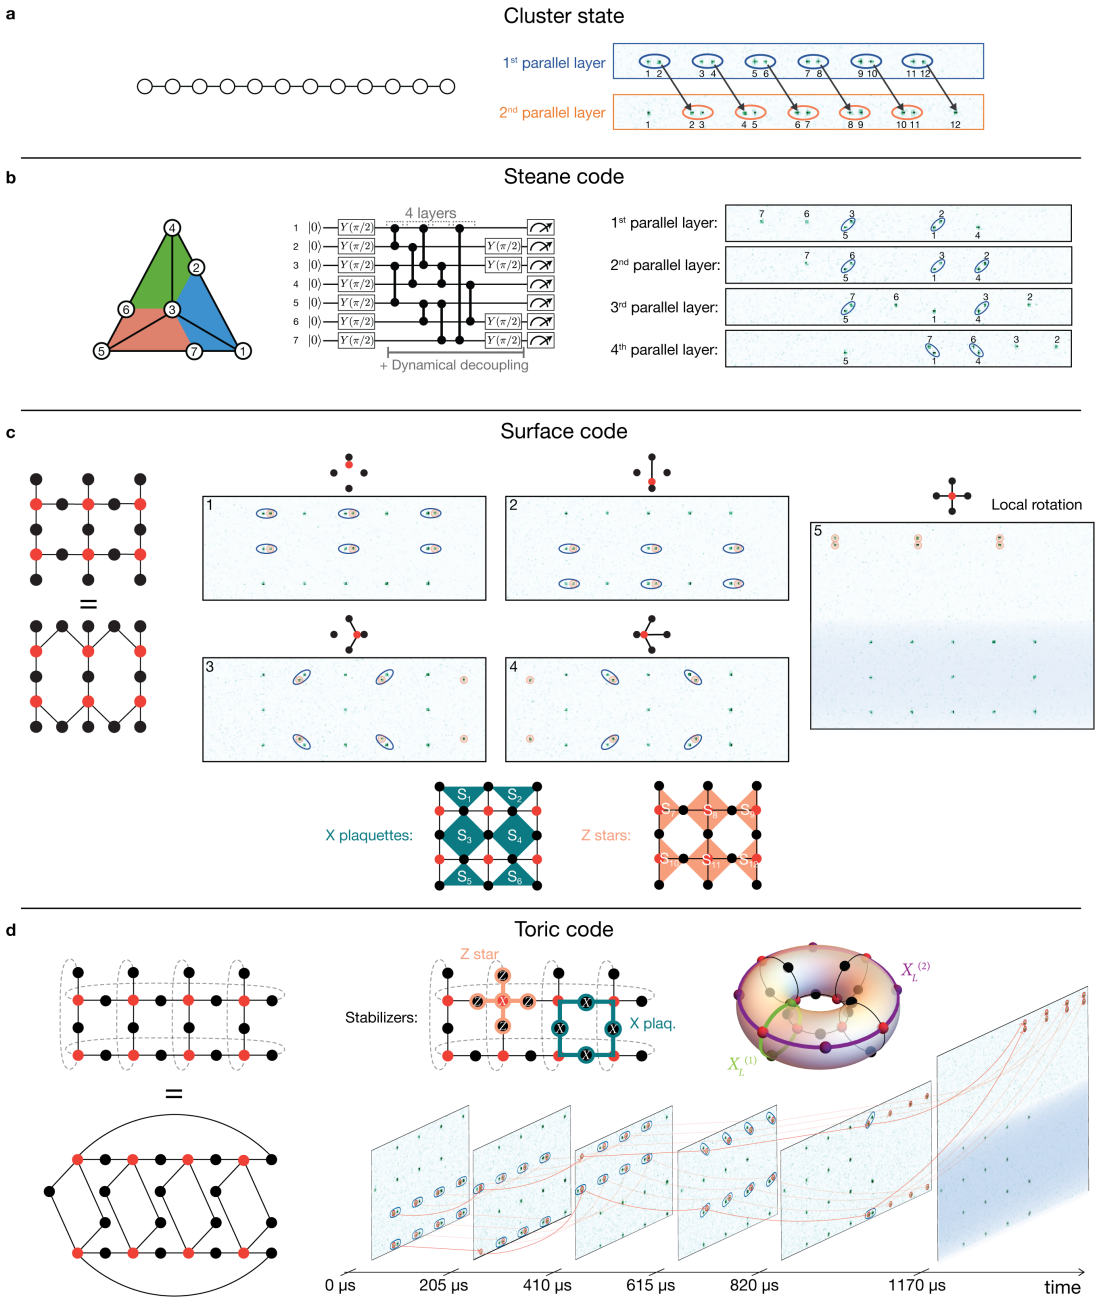


Figure 2.10: Movement schematics. Schematics showing the gate-by-gate creation of **a** the 1D cluster state, **b** the Steane code, **c** the surface code, and **d** the toric code. These various graph states are all generated in the same way, and encoding a desired circuit is a matter of positioning the atoms in different initial positions and applying an appropriate AOD waveform.

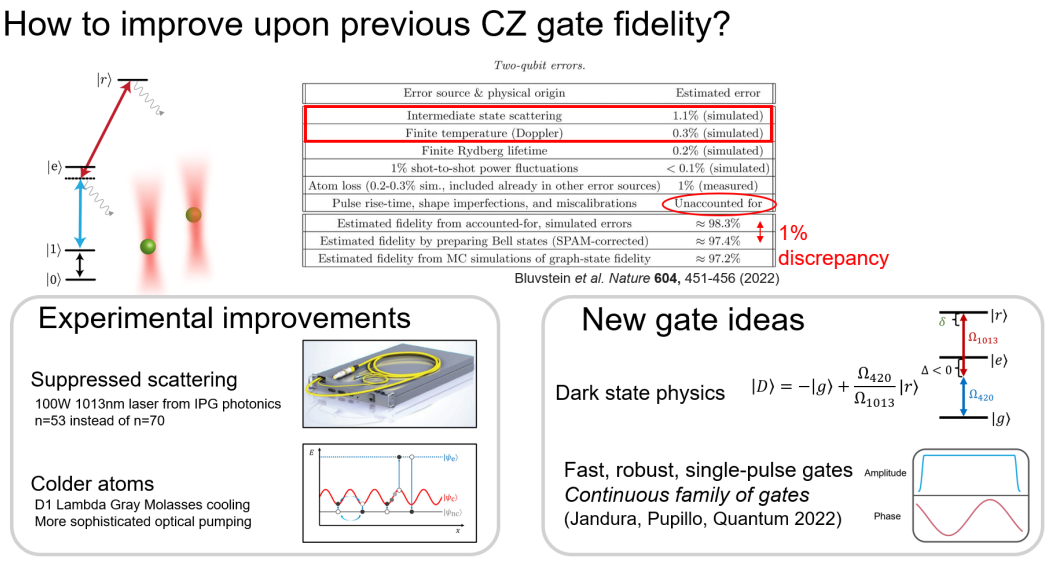


Figure 2.11: Improvements to Rydberg gate to achieve high fidelities. The 97.5% gate fidelity was limited by incoherent error sources but also a discrepancy that was unaccounted for, likely arising from pulse shape imperfections and miscalibrations. To achieve fidelities of 99.5%, the core improvements were: switching to a single-pulse gate which was smoothly parametrizable, optimizing calibration via repeated benchmarking, leveraging dark state physics to suppess scattering, reducing Rydberg n and increasing laser power, and improved cooling of the atoms.

qubit gate[57], meaning our gate fidelity needs to be at least that good. So, it was prudent to try to improve our gate fidelity[52].

At a high level, the Rydberg gate works by pulsing the atoms with a certain optical pulse realized via AOMs. The 1013 pulse is held on at a constant value during the duration of the gate, and the 420 pulse is generated via an AWG card for precise time, frequency, amplitude, and phase control.

The most important improvement to two-qubit gate fidelity in the field was an extremely simple change to the gate protocol of the Levine-Pichler gate. The Levine-Pichler gate works with two pulses of the Rydberg laser with length τ and phase jump φ between the two pulses. This has an analytic solution and in a sense is conceptually very clean. However, there are multiple practical issues likely associated with coherent errors coming from repeated pulse turn-on and turn-off, pulse miscalibration, as well as still-somewhat-mysterious decoherence mechanisms that seem to take place when atoms are left in the Rydberg state without the Rydberg laser continuously driving them.

A seemingly minor, but practically major, improvement to this gate fidelity came from theorists in the field, Jandura et al[80] and Pagano et al[109]. The Jandura-Pupillo / Pagano et al version of this gate is very similar: instead of two laser pulses with a phase jump in-between, it is just a single laser pulse with a smooth variation of the phase. This was found via optimal control and is in fact time-optimal, even in the presence of the ability to have local controls (i.e., having local controls does not make the gate any faster than this simple global-pulse gate). Although only 10% faster than the original Levine-Pichler gate, being smooth and continuous had clear benefit in terms of robustness.

Another conceptual improvement came was realizing that due to the dark-state physics in the system, there is an appropriate sign of two-photon detuning that one should choose, and we were in fact using the wrong one, thereby leading to intermediate state scattering. We also found that this numerically optimized Jandura-Pupillo / Pagano et al gate could in fact be very well-approximated by a simple sinewave modulation of the phase. So the pulse just looks like a square intensity pulse, with a gradual sinewave modulation of the phase with a modulation frequency comparable to the

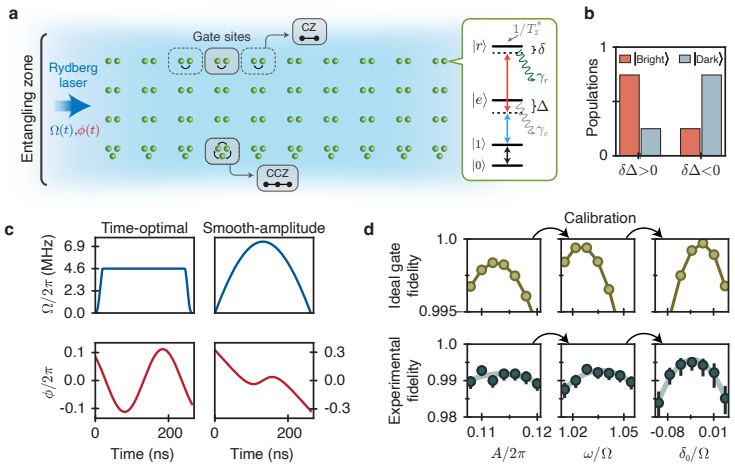


Figure 2.12: Parallel implementation of high-fidelity entangling gates on a neutral atom quantum computer. a, Entangling gates are implemented by arranging atoms into designated gate sites where they interact via Rydberg blockade interactions when pulsing global Rydberg lasers. Two-qubit or three-qubit gates are performed by modulating the Rabi frequency $\Omega(t)$ and phase $\varphi(t)$ profiles of the laser driving the first leg of the two-photon Rydberg excitation. Inset: atoms in the qubit state $|1\rangle$ are excited to the Rydberg state $|r\rangle$ through an intermediate excited state $|e\rangle$, while atoms in $|0\rangle$ are not excited. Main gate error sources include Rydberg state decay γ_r, intermediate-state scattering γ_e, and Rydberg dephasing T_2^*. **b,** Numerical comparison of average bright and dark state populations during the Rydberg gate. Choosing opposite intermediate-state (Δ) and two-photon (δ) detuning signs at the beginning of the gate maximizes population in the dark state, which minimizes the intermediate-state scattering error. **c,** Entangling gates are implemented with a single Rydberg laser pulse with smooth phase modulation $\varphi(t)$, whose slope corresponds to a two-photon detuning $\delta(t)$. A few global parameters characterize the gate, allowing for a family of possible gate implementations, including a parameterized version of the time-optimal CZ gate[?] and a smooth-amplitude CZ gate. **d,** Example gate calibration sequence. Tuning individual parameters of the parameterized time-optimal gate phase profile $\varphi(t) = A\cos(\omega t - \phi) + \delta_0 t$ allows for fast and simple global calibration. Error bars represent 68% confidence intervals.

Rabi frequency - how could it get any simpler than that? [100,35,142]

The other improvements were practical: higher laser power, further detuning from the intermediate state (allowed by the higher laser power), colder atoms, improved Rydberg state choice, more carefully optimized laser polarization, and improved calibration. The Rydberg state choice was in fact the most significant technical improvement - by going to a lower principal quantum number n gate fidelity can be greatly improved. The Rydberg Rabi frequency scales as $n^{-3/2}$, and the optical decay lifetime scales as n^3 and the microwave lifetime (blackbody excitation to adjacent Rydberg states) scales as n^2 (see e.g. Ref[42]). Although naively this seems like it is better to go to higher Rydberg n and that the ratio between Rabi frequency and lifetime scale as $n^{1/2}$ (assuming microwave decay limited), if one wants to entangle N atoms with some fixed amount of laser power, one can simply focus the laser to a much smaller volume at lower n due to the reduced blockade radius so that atoms can be packed closer together (until the interaction strength becomes too weak between adjacent atoms). At a fixed Rabi frequency the blockade radius scales as $n^{11/6}$ (n^{11} comes from the interaction strength and the Blockade radius is set by the $1/r^6$ distance), and so the 2D packing density scales as $n^{-11/3}$. As such, when entangling a fixed N atoms the Rabi frequency effectively gets an additional factor of $n^{-11/6}$, and so the overall ratio between Rabi frequency and lifetime scales as something closer to $n^{-4/3}$ (assuming microwave-decay-limited), meaning it is better to go to lower n until one runs into issues e.g. the interaction being too weak to have robust blockade in the presence of positional disorder.

For improved calibration, we made the conceptually very simple upgrade to just repeat the gates many times in a row in order to amplify the signal and more sensitively calibrate the gate[87,8]. In particular we used the benchmarking sequences shown in Fig. 2.14 where we repeat the CZ gate many times in a row interspersed with random single-qubit rotations, in a global version of randomized benchmarking. We then scan each gate parameter, fit a parabola and take the maximum, scan the next parameter and fit a parabola, and so on.

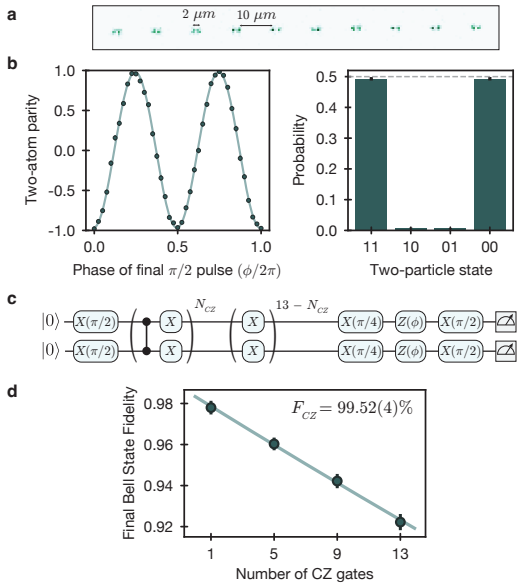


Figure 2.13: High-fidelity CZ gates characterized with Bell states. **a,** Parameterized time-optimal CZ gates are implemented on 20 atoms in parallel. **b,** Raw Bell state measurements upon application of a single CZ gate, with raw Bell state fidelity of 98.0(2)%. We estimate a SPAM-corrected fidelity of 99.4(4)% (not plotted). **c,** Circuit used to benchmark two-qubit gate fidelity by making a Bell state after an odd number of CZ gates interleaved with single-qubit X gates. **d,** Decay of Bell state fidelity after applying a variable number of CZ gates. A fidelity of 99.52(4)% per CZ gate is extracted from an exponential fit. The plot y-axis on (d) is rescaled by a SPAM correction factor which does not affect the extracted gate fidelity. Note that the same number of single-qubit gates are performed for each point, so that no normalization of single-qubit errors is necessary, which thereby results in a reduced y-axis intercept. Error bars represent 68% confidence intervals.

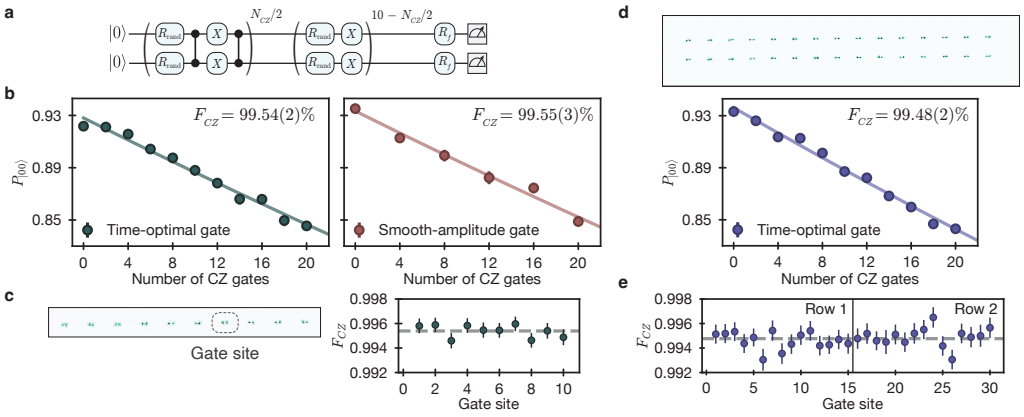


Figure 2.14: Global randomized benchmarking of the CZ gate. a, Circuit used to benchmark two-qubit gate fidelity with global randomized benchmarking, in which random rotations R_{rand} are sampled from a Haar-random distribution (Methods) and a final rotation R_f is precomputed to return population to the initial product state $|00\rangle$ in the absence of gate errors. All 21 single-qubit rotations are applied for each data point, such that the number of single-qubit operations is independent of the number of two-qubit gates. **b,** Benchmarking CZ gates on 20 atoms in parallel. We plot the probability of return to the initial product state as a function of the number of applied CZ gates for both the parameterized time-optimal gate and the smooth-amplitude gate, extracting fidelities of 99.54(2)% and 99.55(3)%, respectively. **c,** The extracted time-optimal gate fidelity is consistent within statistical error across the 10 gate sites. **d,** Scaling to larger systems, we measure a comparable fidelity of 99.48(2)% for the time-optimal gate implemented on 60 qubits in parallel. **e,** Site-by-site analysis reveals homogeneous, high-fidelity gate performance across the two rows.

These various simple improvements improved our two-qubit gate fidelity to 99.2% - 99.3% just by changing to smooth Jandura-Pupillo pulse (implemented as just a simple sinewave), being careful about detuning signs to suppress scattering via dark state physics, changing n to lower values (from $n = 70$ to $n = 53$), and improved calibration via randomized benchmarking and scanning parameters. We were able to improve the gate to 99.55% by cooling the atoms better (via Λ-enhanced gray molasses) and increasing our 1013-nm laser power. Gate fidelities with atoms will soon likely get to 99.9% just by further improvements along these lines. Error-budgeting can be found in the table shown in Figure 2.15. An example approach could be going to 3x higher Rabi frequency (while also suppressing coupling to the other m_J state) and 2x larger detuning, simply requiring 6x higher laser power. We haven't had time to try improving the gate fidelity since our high-fidelity gates project, but routinely reproduce 99.5% - 99.6% gate fidelities.

There are various choices one can make for Rydberg detuning signs and polarizations, and we carefully select them for various optimizations: for suppressing 420-induced differential light shift between $|0\rangle$ and $|1\rangle$, we are red-detuned of $6P_{3/2}$ transition; for utilizing dark-state physics (nominally the phase profile corresponds to a certain sign of two-photon detuning) we thereby choose positive two-photon detuning; which, in turn then suppresses coupling to $m_J = -1/2$ by being primarily on the upper-side of $m_J = +1/2$; and finally, suppressing 1013 light shift similarly prefers being at this single-photon detuning sign since there is a magic wavelength 1 GHz red-detuned of $6P_{3/2}$ for the $|1\rangle \rightarrow 53S_{1/2}$ transition[90]. Two downsides of this detuning choice are that this choice of 420-nm polarization and detuning causes a vector light shift in the hyperfine ground-state manifold that causes the m_F levels to be pushed closer together, as opposed to further apart, which could exacerbate effects arising from 420-induced vector light shifts coupling adjacent m_F states in the ground state manifold (although negligible here), and the other downside is that the $m_J = -1/2$ Rydberg pair states are closer detuned to the two-photon excitation, and so we require a larger interaction strength to suppress their excitation, though the matrix element to these states is smaller.

Error source	Time optimal	Smooth amplitude	Error type X, Y, Z, LG, AL**
Scattering* $\|1\rangle$	0.103% / 0.043%	0.036%	6%, 6%, 25%, 47%, 15%
Scattering $\|0\rangle$	0.019%	0.025%	7%, 7%, 14%, 62%, 10%
Rydberg $T_1 = 88\,\mu$s	0.113%	0.085%	2%, 2%, 6%, 23%, 67%
Rydberg $T_2^* = 3\,\mu$s	0.134%	0.089%	0%, 0%, 75%, 0%, 25%
Position fluct.	0.012%	0.054%	0%, 0%, 96%, 0%, 4%
Rydberg $m_{J=-\frac{1}{2}}$	0.06 - 0.15%	0.01%	
Total fidelity	**99.53 - 99.62%**	**99.70%**	**2%, 2%, 41%, 17%, 38%**

Figure 2.15: Simulated error budget for two-qubit CZ gates. Simulations were performed at an intermediate-state detuning of $\Delta/2\pi = 7.8$ GHz and a two-photon Rabi frequency $\Omega/2\pi = 4.6$ MHz. (*) The scattering error for two detuning signs (bright/dark). The total fidelity values assume the correct (dark) detuning choice. (**) LG — leakage out of the qubit manifold to other m_F states. AL — atom loss from population left in the Rydberg state.

It is important that the Rydberg beam is homogeneous across the atomic array while still maximizing intensity (i.e., not using a massive Gaussian beam). For this reason we shape the Rydberg beams into flat-top potentials via the use of an SLM [49]. So far we are routinely able to homogenize these to the required levels such that inhomogeneity is not limiting our fidelity, but improved methods for homogenization will soon become important.

We briefly overview our Rydberg laser system. The 1013-nm laser seed originates from a Toptica DL Pro external-cavity diode laser which is then locked to and filtered through a Stable Laser Systems ultra-low-expansion cavity (Finesse of 50,000 at 1013 nm) that then injection locks another laser diode, which then seeds our 100W laser. Our 420-nm laser is an 840-nm TiSapph (MSquared) which is locked (but not filtered) to the same cavity (Finesse of 30,000 at 840 nm) and then frequency doubled with an MSquared ECD-X. For realizing the pulses and waveform shaping for quantum gates, we use an arbitrary waveform generator (Spectrum M4i.6631-x8) that allows for arbitrary amplitude, frequency, and phase control which simultaneously drives two 420-nm AOMs (MQ240-B40A0,2-UV, AA Opto-Electronic) in a tandem configuration and maps the RF waveform onto the 420-nm light. The 1013-nm AOM (M1377-aQ80L-1 coated for 1013 nm, Isomet) is pulsed on for several hundred microseconds and intensity stabilized during the duration of the entire circuit.

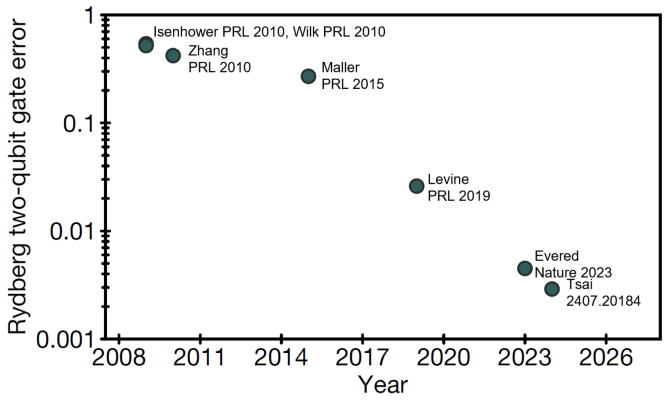


Figure 2.16: Rydberg gate fidelity vs year. Following the current trend, we will reach 99.9% gate fidelity around 2027. 99.9% can be reached from our current results, in theory, by increasing Rydberg laser power by 6x.

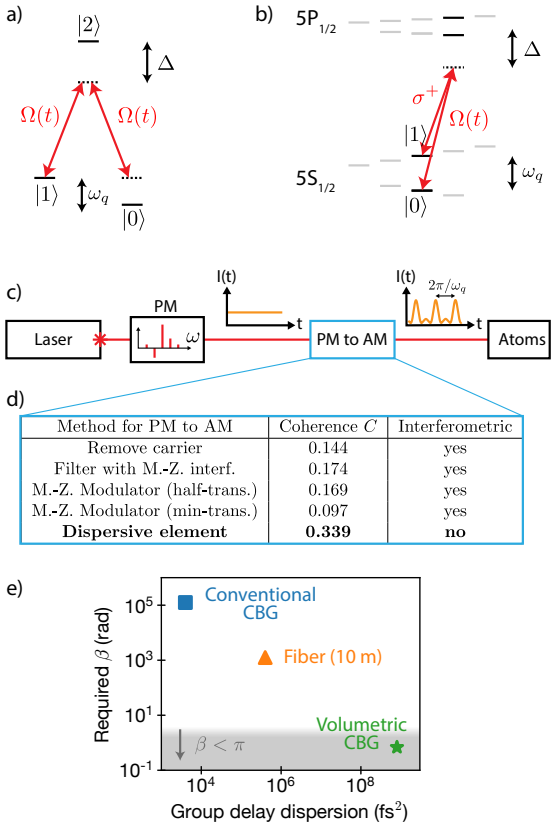
2.7 Programmable single-qubit gates

Additional details for the description in this section about the Raman system can be found in Ref.[95]*, and about coherence and local Raman addressing can be found in Refs.*[19,17].

The $m_F = 0$ hyperfine qubit is an electron-nuclear hybrid within the ground state manifold of the Rb atom with a 6.8 GHz splitting between the two qubit states $|0\rangle \equiv |F = 1, mF = 0\rangle$ and $|1\rangle \equiv |F = 2, mF = 0\rangle$. As these hyperfine states have an electron spin component, they can in principle be driven relatively well with magnetic fields[30,130]. However, even then, as the atoms are so far away from their macroscopic environment, it is quite hard to get much higher than a 50 kHz Rabi frequency by shining microwaves on them. Illuminating microwaves on the atoms also doesn't give us any individual control.

Raman transitions offer an excellent solution to this problem[157,83]. In a typical description, these are two-photon transitions through an intermediate excited state that can drive the qubit transition in an optically-mediated fashion that can have a much higher Rabi frequency. However, the underlying physics is quite intuitive.

Circularly polarized, off-resonant light generates vector light shifts on the atom, which generate



Method for PM to AM	Coherence C	Interferometric
Remove carrier	0.144	yes
Filter with M.-Z. interf.	0.174	yes
M.-Z. Modulator (half-trans.)	0.169	yes
M.-Z. Modulator (min-trans.)	0.097	yes
Dispersive element	**0.339**	**no**

Figure 2.17: Amplitude modulation for driving Raman transitions. (a) Stimulated Raman transitions in a Λ-type 3-level system. Adiabatic elimination of the excited state results in an effective Raman coupling between ground states $|0\rangle$ and $|1\rangle$. (b) Level structure for ^{87}Rb, showing Raman driving of the clock transition from $|0\rangle \;=\; |F=1, m_F=0\rangle$ to $|1\rangle \;=\; |F=2, m_F=0\rangle$. This transition is driven by a time-dependent σ^+ polarized field $\Omega(t)$, which is far-detuned by Δ from the excited state (but not far-detuned relative to the splitting between the $5P_{1/2}$ and $5P_{3/2}$ excited states). (c) Several approaches for Raman driving, including the dispersive approach presented here, operate by converting phase modulation to amplitude modulation at the qubit frequency, which resonantly drives the Raman transition. (d) Comparison of methods for converting phase modulation to amplitude modulation. The dispersive approach benefits both from having the highest coherence metric and from being passively stable since it does not rely on interferometric filtering. (e) Weakly dispersive elements, such as a conventional chirped Bragg mirror or a 10 meter optical fiber, require large modulation depth β to achieve efficient amplitude modulation. The highly dispersive volumetric CBG allows a low β to be used. $\beta \lesssim \pi$ marks the experimentally accessible window of modulation depths.

a fictitious magnetic field in the hyperfine ground state which is indistinguishable from a real magnetic field[95]. A somewhat wrong picture for this fictitious field, that I personally really like, is that the circularly polarized light induces a dipole that rapidly rotates and effectively the orbiting electron generates a static magnetic field along the direction of propagation of the laser. Broadly, the circularly polarized, detuned light leverages the optical transitions to generate a magnetic field on the spin. This gives us a way to create a very strong and localized magnetic field on the atoms. And we can now use this to conveniently drive the qubit transitions.

To realize this, we take our light at a detuning of 180-500 GHz detuning (depending on application) blue-detuned of the D1 transition, pass it through a free space electro-optic modulator (EOM) (leading to phase modulation only) which is driven by a microwave generator controllable via IQ modulator, and then bounce this light off a *chirped Bragg grating*[95]. This chirped Bragg grating results in the different frequency components of the phase-modulated light to travel a different distance into the crystal before reflection, and thereby modifies the phase relationship between the tones to convert the phase modulation into amplitude modulation. In this way, we map the microwave signal directly onto the amplitude modulation of the light, in a way which is passive, has low loss / high efficiency, and requires no active stabilization. Now, we have our amplitude-modulated laser beam which encodes the microwave signal (which is controlled in a usual way via IQ modulation of the carrier by an AWG). Pointing the light at an atom then locally delivers the microwave signal by causing a strong, time-oscillating fictitious magnetic field on the atom which directly drives the qubit transition.

After the amplitude-modulated light is created, it is routed to the atom via two ports. One is a global port that globally illuminates the atoms for dynamical decoupling, global pulse control, and so on. We have done experiments at detunings of up to 500 GHz, and by using composite pulse techniques to be amplitude-error-insensitive (either BB1 which is 6th-order insensitive[148] or SCROFULOUS[38] which is shorter and 4th-order insensitive) our global pulses reach a scattering-

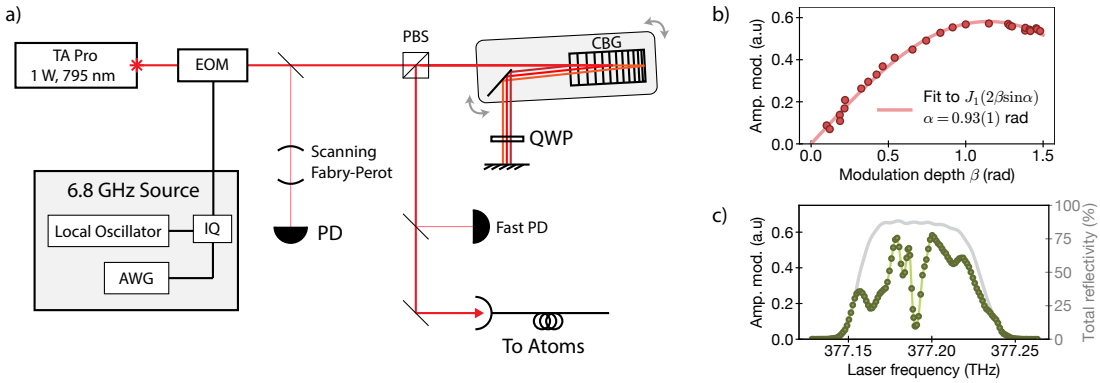


Figure 2.18: Raman laser system using a chirped Bragg grating. (a) Optical setup. The chirped Bragg grating (CBG) and the first mirror afterwards (in the shaded gray region) are mounted on a single rotation mount. Spectral components separate after the first reflection from the CBG, but recombine after the second reflection. A scanning Fabry-Perot cavity measures the sideband spectrum, and a fast photodetector measures the amplitude modulation. (b) The amplitude modulation depends on both the dispersion of the CBG as well as the phase modulation depth (see main text). We observe the expected Bessel function relation, and can extract the dispersion coefficient. (c) As we scan the laser frequency across the CBG bandwidth, we see a high total reflectivity of the CBG system across the $\sim$ 50 GHz bandwidth. The resulting fiber-coupled light should ideally show constant amplitude modulation across the whole bandwidth, but in practice we observe variation with laser frequency due to nonuniform CBG dispersion over its bandwidth. While more uniform CBGs can be used, the current device can be angle tuned to maximize amplitude modulation, and is insensitive to frequency drifts which are $\lesssim$ 1 GHz.

limited single-qubit pulse fidelity of approximately 99.995%.

The other is a local port that is then passed through an additional 2D AOD which is driven by an additional AWG card. This allows us to point the amplitude-modulated light on any given atom and do arbitrary single-qubit gates on each atom. In practice, however, one often just needs to realize large-scale parallel operations. For instance, in realizing a *logical* single-qubit gate (next chapter), one just needs to apply that same single-qubit gate on all the data qubits of the block. This can be conveniently realized by a single, parallel grid illumination via the AODs. As such, the AOD Raman approach allows for arbitrary single-qubit addressing, as well as still benefiting from parallelism where applicable.

There is one subtletly to the story, which is that one needs to be very careful about the direction of the magnetic field that one intends to drive the atoms with. In our experiments, our magnetic

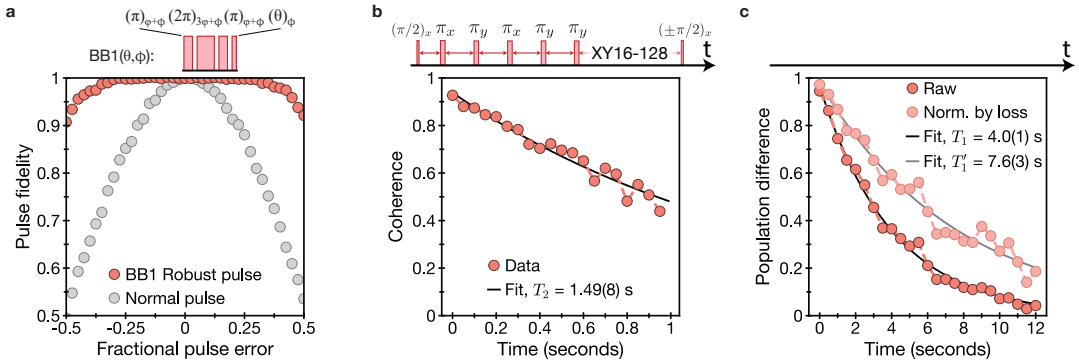


Figure 2.19: Robust single-qubit control and qubit coherence. a, Robust BB1 single-qubit rotation in comparison to a normal single-qubit rotation, as a function of pulse area error. An arbitrary BB1(θ, φ) rotation on the Bloch sphere of angle θ about axis φ is realized with a sequence of four pulses: $(\pi)_{\phi+\varphi}(2\pi)_{3\phi+\varphi}(\pi)_{\phi+\varphi}(\theta)_{\varphi}$, where $\phi = \cos^{-1}(-\theta/4\pi)$[148]. Pulse fidelity is measured here for a π pulse, defined such that the fidelity is the probability of successful transfer from $|0\rangle \rightarrow |1\rangle$, including SPAM correction. **b,** Preserving hyperfine qubit coherence using dynamical decoupling (XY16 with 128 total π pulses). Qubit coherence is observed on a timescale of seconds, with a fitted coherence time $T_2 = 1.49(8)$s. Data is measured with either a $+\pi/2$ or $-\pi/2$ pulse at the end of the sequence, and these curves are then subtracted to yield the coherence y-axis. **c,** Hyperfine qubit T_1, measured by the difference of final $F = 2$ populations between measurements starting in $|F = 2, m_F = 0\rangle$ and $|F = 1, m_F = 0\rangle$. Atom loss without cooling is separately measured (predominantly arising from vacuum loss) and normalized to also measure the intrinsic spin relaxation time T_1' in the absence of atom loss. All data here is measured in 830-nm traps.

field is pointing along the long-axis of the array, where the circularly polarized Rydberg beams, imaging beams, lattice beams, Raman beams etc all naturally propagate along the magnetic field axis. For the Global Raman control, this works well as driving qubit transitions is done via magnetic fields parallel to the external quantization axis. (It is given by the parallel component, not the perpendicular one, because the two $mF = 0$ states are electron-nuclear hybrid states that have different phase i.e. symmetric and antisymmetric wavefunctions, and transitions between them are driven by modifying the wavefunction phase)[95].

This becomes a minor hiccup for the local Raman control, as the light has to be circularly polarized (otherwise it wouldn't generate a fictitous magnetic field i.e. vector light shift) but is propagating from out-of-plane, and so is orthogonal to the convenient axis of the external magnetic field. Somehow we need to generate a parallel magnetic field component. We implemented two com-

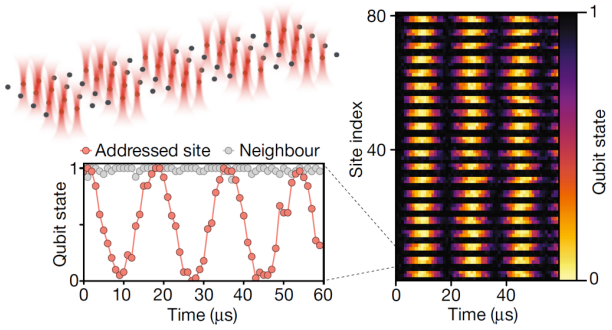


Figure 2.20: Local Raman addressing. Fully programmable single-qubit rotations are implemented using Raman excitation through a 2D AOD; parallel grid illumination delivers the same instruction to multiple atomic qubits.

plementary solutions to this. The first solution is that the detuned Raman beam also contains a DC-field component and thereby rotates the magnetic field axis dynamically (for this reason it is also very important to do Gaussian-pulse shaping of the Raman excitation to remain adiabatic) and leads to having a parallel component of the AC field along the DC quantization axis of the atom. The second solution is that we anyway want the Raman beam to be big on the atoms (see below), which in other words means it is composed of a small angular spread, and so it is in fact quite small in the back aperture of the objective. As such, one has the angular resolution available to displace the beam horizontally within the back aperture and do the local addressing at an *angle* relative to the array axis. We displace the Raman beam by $\sim 1 - 2$ mm within the back aperture of the objective whose size is 5.5 mm. These two combined give us Rabi frequencies of MHz on the atoms and with the use of composite pulse techniques and Gaussian shaping we achieve 99.95% local single-qubit gate fidelities, which is close to the scattering limit on account of the higher intensity required to achieve the same Rabi frequency in this suboptimal field configuration.

We have found it is quite important to make the local Raman beam bigger on the atoms while still making sure to avoid crosstalk on the separated atoms. One reason is that it just makes things easier and more robust and less sensitive to drifts. Another reason is that polarization breakdown

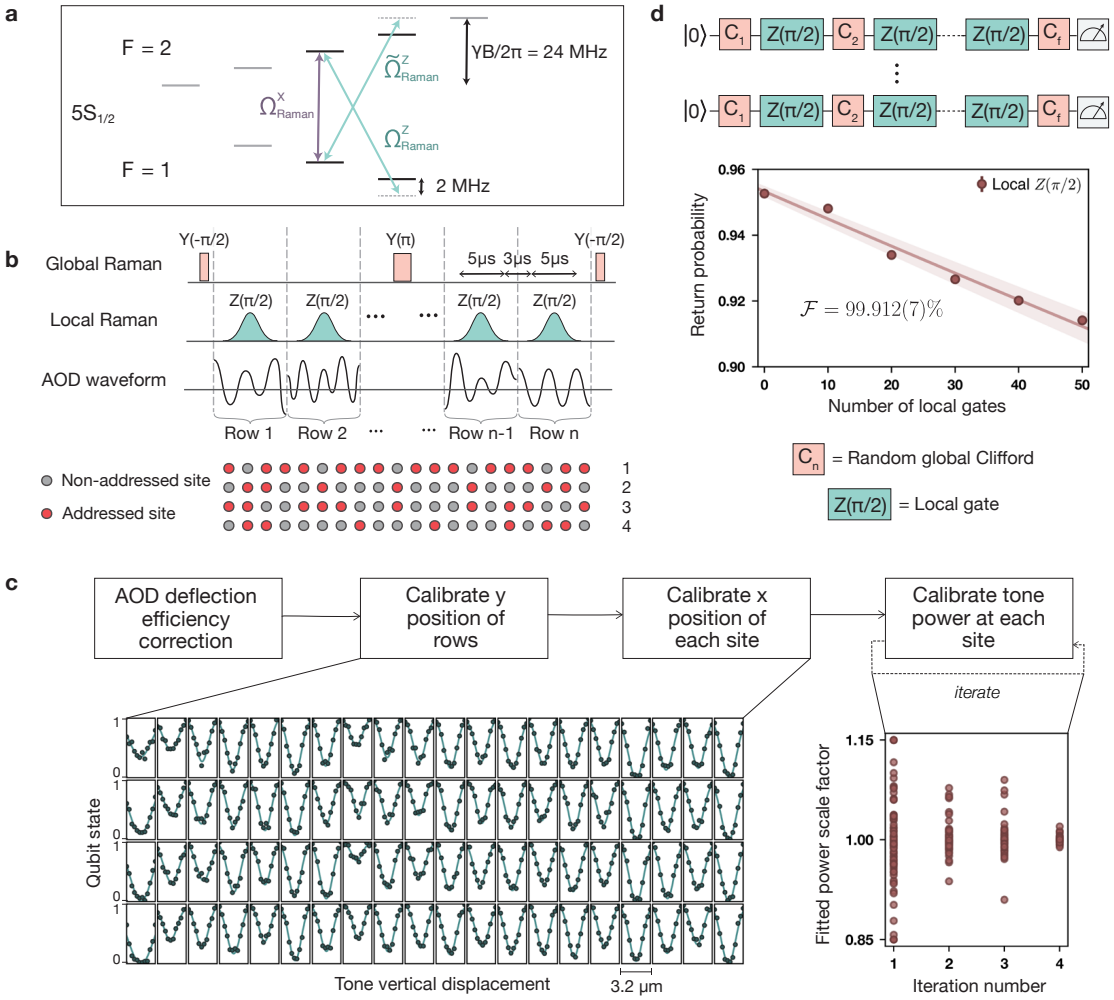


Figure 2.21: Single-qubit Raman addressing. a, $5S_{1/2}$ hyperfine level diagram illustrating the two possible implementations of local single-qubit gates: resonant $X(\theta)$ (purple) and off-resonant $Z(\theta)$ (turquoise) rotations with two-photon Rabi frequencies Ω_{Raman}. In this work, we use the Z rotation scheme and are blue-detuned by 2 MHz from the two-photon resonance. Due to Clebsch-Gordan coefficients, $\widetilde{\Omega}^{Z}_{\mathrm{Raman}} = -\sqrt{3}\Omega^{Z}_{\mathrm{Raman}}$. **b,** Schematic showing the conversion of local $Z(\pi/2)$ into local $X(\pm\pi/2)$ gates, where the pulses before (after) the central $Y(\pi)$ have positive (negative) sign, while leaving non-addressed qubit states unchanged. The Gaussian-smoothed local pulses have duration 2.5 μs for $\pi/4$ pulses and 5 μs for $\pi/2$ pulses, and are performed on single rows at a time with a 3 μs gap between subsequent gates to allow the RF tones in the AODs to be changed (including this, duration is 5-8 μs per row). In this way, arbitrary patterns of qubits, such as the example drawn, can be addressed. **c,** Calibration procedure used to homogenize the Rabi frequency over a 220 μm $\times$ 35 μm array. The position calibration is illustrated for 80 sites: approximate $X(\pi/2)$ gates are locally performed and the horizontal/vertical position of all tones is scanned in parallel such that a Gaussian fit returns the optimal alignment. After this, powers are iteratively calibrated until the fitted scale factors for the individual RF tones converge to unity. **d,** Single-qubit randomized benchmarking of local $Z(\pi/2)$ gates. The local gates are interleaved with random global single-qubit Clifford gates and the final operation C_f is chosen to return to the initial state. Each data point is the average of 100 random sets of Clifford gates, and fitting an exponential decay to the return probability quantifies the fidelity $\mathcal{F}$ per local gate. Note that we apply all 51 global Clifford gates for each data point, such that errors from the global Clifford gates (in addition to SPAM errors) do not contribute to the fitted value.

within the beam can have very complex interactions with the dynamically rotating DC magnetic field and AC field in ways that can lead to destructive interference and dephasing. Making the beam bigger, which we have found empirically and via calculations is essentially identical to just defocusing the beam (to an in-plane waist of $\sim$ 2.5μm) was important for us to achieve the 99.95% single-qubit gate fidelity.

Although the local Raman theoretically allows us to realize large parallel grid illumination, in practice we are currently often doing illumination row-by-row. This allows us to calibrate the amplitude of each atom arbitrarily within each row with the AOD. By keeping the RF power low, we minimize RF beating effects and don't enforce a comb condition on the frequency tones of our local Raman AOD.

2.8 Mid-circuit readout and feedforward

Additional details for the description in this section can be found in Ref.[17].

In quantum error correction, measurement is the mechanism for removing entropy from the system. However, critically, we do *not* want to measure the qubits that still contain quantum information. We take a simple approach to this: we move the qubits we *do* want to measure to a separate "readout" zone where we can scatter photons off of them to read out their qubit state without corrupting the state of the active qubits that are continuing to preserve quantum information.

We have done two different versions of midcircuit readout: destructive and non-destructive. The destructive version works well for quickly getting information out of the system and using it for feedforward operations on our logical (or physical) qubits. To realize such a readout, we quickly convert qubit state 1 to loss, and then quickly image whether an atom is present or not present. The non-destructive version on the other hand is critical to re-using atoms such that many repeated measurements can be done (in lieu of just bringing in new atoms altogether) for continuously removing

entropy from the system. To be non-destructive, it also requires a method for detecting qubit state without relying on converting qubit state 1 to loss. This section will overview the destructive version of mid-circuit readout we first implemented, and then the next section will describe the non-destructive version we used to achieve qubit re-use.

To perform midcircuit readout[43,135,62,101,98,107] of selected qubits without affecting the others, we use a local imaging beam focused on the readout zone which is roughly 100 μm spatially separated from the entangling zone[19,37]. The local imaging beam consists of 780-nm circularly polarized light, with a near-resonant component from $F = 2$ to $F' = 3$ and a small repump component. This beam is sent through the side of our glass vacuum cell, co-propagating with the global Raman and 1013-nm Rydberg beams. We use cylindrical lenses to shape the beam with focused beam waists of 30 μm in the plane of the atom array and 80 μm out of the plane. After moving some of the atoms to this readout zone, we first perform local pushout of population in the $F = 2$ ground state manifold (by turning off the repump laser frequency), followed by local imaging of the remaining $F = 1$ population.

As depicted in Fig. 2.22a, we collect an average of about 50 photons per imaged atom. To avoid losing the atoms too quickly during mid-circuit imaging (which, unlike our global imaging scheme, does not have multi-axis cooling), we use deep (roughly 5-mK) traps (helping retain the atoms), and stroboscopically pulse them on and off out of phase of the local imaging light to avoid deleterious effects of the deep traps such as inhomogeneous light shifts and fluctuating dipole force heating (Fig. 2.22b)[77]. From a double-Gaussian fit to the two distributions in Fig. 3a, we extract an imaging fidelity of over 99.9%. Because this fit can lead to an overestimate of the imaging fidelity (for example due to atom loss during imaging), we compare the total SPAM error (measured by amplitude of Ramsey fringe) with local imaging versus with global imaging for the same state preparation sequence, extracting 0.14(5)% higher error with local imaging; with these considerations we conservatively estimate a local imaging fidelity of around 99.8%.

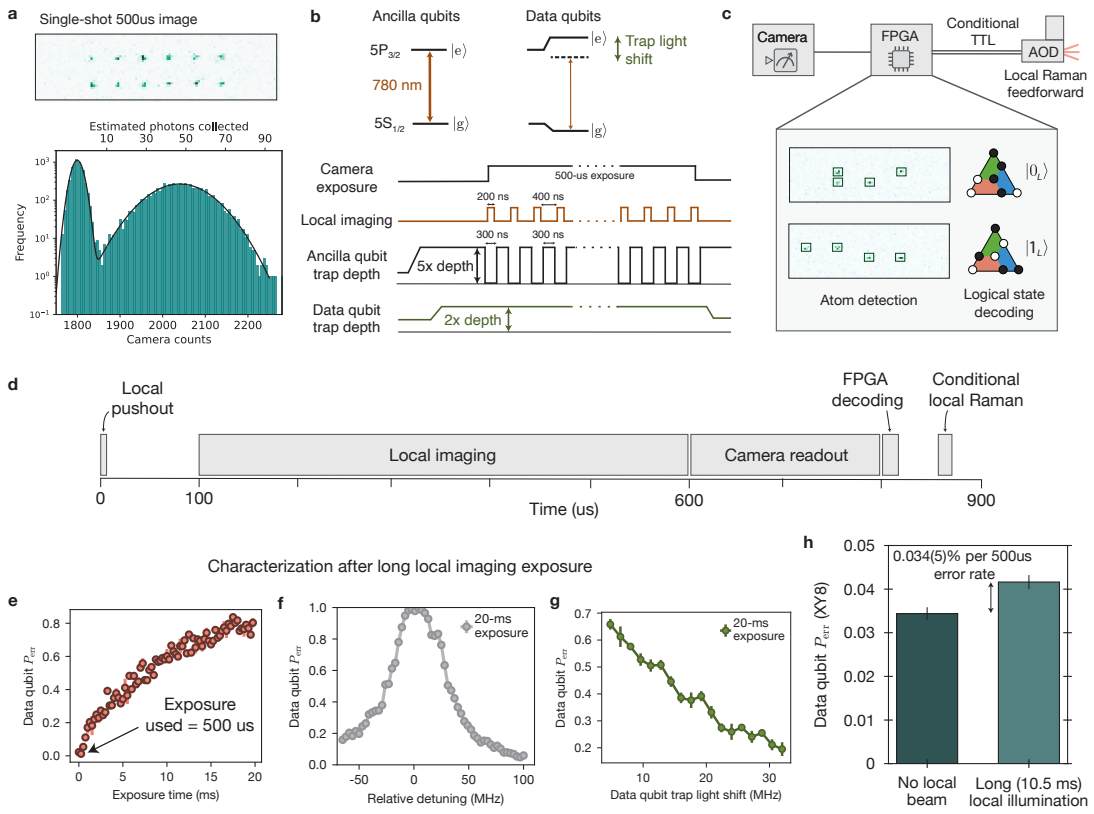


Figure 2.22: Midcircuit readout and feedforward. **a,** Single-shot 500 μs local image in the readout zone, where the peak corresponds to roughly 50 photons collected by the CMOS camera. **b,** Atomic transition and pulse sequence used for local imaging of ancilla qubits. The data qubit trap light shift suppresses data qubit errors, in addition to the large spatial separation between entangling and readout zones. We avoid quickly losing the readout zone atoms during local imaging by using a 5x higher trap depth, and we pulse the ancilla qubit traps and local imaging light to image directly on resonance while avoiding negative effects of large trap light shifts. **c,** Diagram of components involved in midcircuit readout and feedforward steps. Atom detection and logical state decoding occur using the FPGA, which then outputs a conditional TTL to gate local Raman pulses performed on logical qubits in the entangling zone. **d,** Diagram of approximate timings for a midcircuit feedforward cycle. First, $F = 2$ population is pushed out (in 10 μs), and then the remaining $F = 1$ population is imaged locally for 500 μs. The 24 rows of pixels covering the readout zone are read out to the FPGA in 200 μs, after which processing is performed. Finally, a conditional TTL output based on the decoded state gates on or off local Raman pulses. The whole readout and feedforward cycle takes less than 1 ms, and can be sped up in the future by optimizing local imaging and camera readout. **e-g,** Characterization of error probability of data qubits during local imaging. **e,** Data qubit error probability (fraction of population depumped from $F = 2$ to $F = 1$) as a function of local imaging duration out to 20 ms to quantify the effect of the local imaging beam on data qubit coherence for very long illumination. **f,** Data qubit error probability after 20 ms of local imaging, as a function of detuning of the local imaging beam, showing suppression of error both red- or blue-detuned from the data qubit transition. **g,** Equivalently, increasing the trap depth of the data qubits enables suppression of decoherence due to the local imaging beam. Since qubits in the readout zone are imaged while their traps are pulsed off, any light shift of the data qubit transition from the traps contributes directly to the relative detuning. **h,** For a long 10.5 ms local beam illumination with optimal local imaging parameters, we observe a 0.7(1)% increase in data qubit error during an XY8 dynamical decoupling sequence. This suggests a roughly 0.034(5)% error probability for the data qubits during the 500 μs midcircuit readout image employed in this work.

A number of design considerations facilitate local imaging in the readout zone while preserving coherence of the data qubits in the entangling zone (Fig. 2.22e-g)[37]. The main sources of decoherence are rescattering of photons from the locally imaged atoms as well as beam reflections and tails of the local imaging beam hitting the data qubits. For the 500-μs midcircuit imaging used in this work, we are able to achieve unchanged coherence (identical within the errorbars) of the data qubits with the local imaging light on as without it. To understand these effects more quantitatively, we measure the error probability of the data qubits in the entangling zone while the local imaging beam is on in the readout zone for up to 20 ms and with higher intensities than used for local imaging in this work. We suppress decoherence by light shifting the data qubits' 780-nm transition to be different from that of the locally imaged qubits by several 10s of MHz, as studied in Fig. 2.22f-g. Data qubit decoherence is further suppressed by the large spatial separation between the readout zone and the entangling zone, where intensity from the local imaging beam's Gaussian tail should theoretically fall off rapidly. Even at large separations, we find that stray beam reflections (e.g. from the glass cell window and other optical elements) can hit the data qubit region. To mitigate this effect, we displace reflections away from the atom array by angling the local imaging beam as it hits the glass cell window. The estimated effects of re-scattered photons from the imaged atoms, especially with the added relative detuning, is negligible. With all these considerations, we find that we are able to suppress data qubit decoherence rates to $\lesssim 0.1\%$ per 500 μs of local imaging exposure, as illustrated in Fig. 2.22h.

The full mid-circuit readout and feedforward cycle occurs in slightly less than 1 ms, including local pushout, local imaging, readout of the camera pixels, decoding of the logical qubit state on the FPGA, and a local Raman pulse which is gated on or off by a conditional trigger (Fig. 2.22d). In future work, this approach to midcircuit readout and feedforward can be considerably improved to enable mid-circuit readout close to 100-μs scale[129].

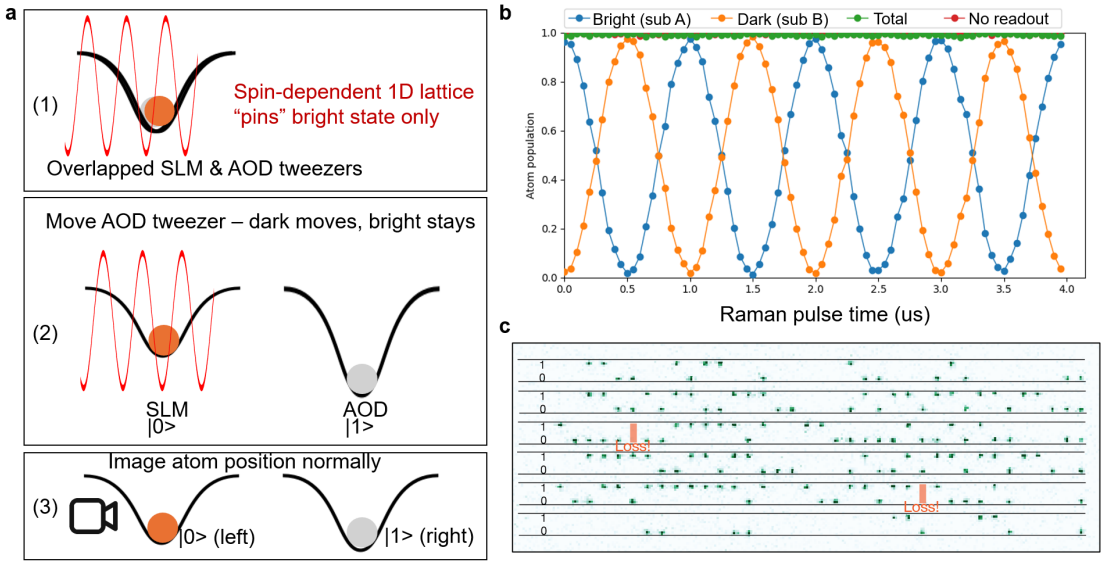


Figure 2.23: State-selective readout with 1D lattice a, Image of processor and illustration of process for removing entropy generated by computation. **b,** Rabi oscillations measured by reuse of atoms with 150 cycles of non-destructive measurement and reinitialization. Each curve is one experimental repetition averaged over 200 atoms in parallel. 3D cooling methods are used in this subplot as coherence does not need to be preserved. **c,** Local cooling with 1D PGC and EIT. The finite magnetic field is compensated simply by applying a relative detuning between the two circularly polarized counterpropagating beams, lending to a rotating frame where the effective field is zero. The loss is constant as a function of time and recovers to steady state within one cycle after applying a perturbation to the system (dashed line). Additional shielding of the data qubits is provided by a 1530 shielding laser (not pictured here, see Methods).

2.9 Atomic qubit re-use and loss-resolving measurement

Additional description for this section can be found in Ref.[18].

Error correction moves entropy from the stabilizer degrees of freedom to ancillary degrees of freedom which can then be measured and reinitialized, such that entropy is removed from the system. As such, it is essential to be able to remove entropy mid-circuit. This section describes the approach used for this, as well as for direct detection of qubit loss for improved entropy removal performance.

Figure 2.23 shows our approach to non-destructively measuring an atom. This has the benefit that entropy is removed from the system but the atom is retained, such that it can be used for extracting more entropy. It also has the benefit that it can now measure an outcome of $\{0, 1, \text{loss}\}$ for

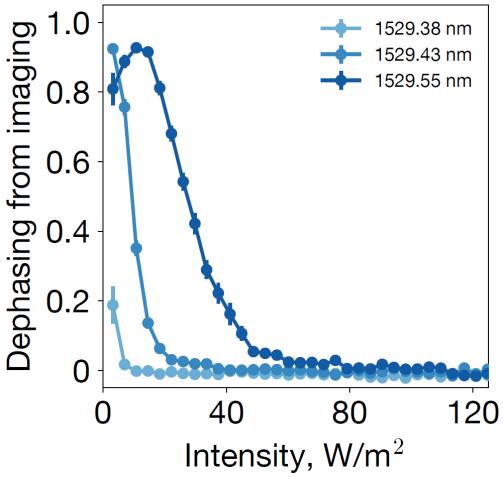


Figure 2.24: Hiding data qubits with 1530 light shift. Applying a 1530 beam to the data qubits shifts their excited states out of resonance by coupling to a 4D excited state, preventing any residual dephasing during imaging while having only minimal impact on the ground state where the qubit is stored.

each atom, which gives significantly more information which can be used effectively in the error correction.

Our non-destructive readout technique uses a 1D optical lattice which is spin-selective and couples to only one of the two spin states, heavily inspired by Ref. [149]. We then perform a normal tweezer motion, and the "bright" spin state is pinned by the lattice but the "dark" spin state moves through the lattice and follows the moving tweezer. Once this spin-to-position conversion has been accomplished, we image the atom position normally, as shown in the figure. This method is particularly elegant as it doesn't involve additional calibration beyond a normal tweezer motion, and functions well with normal / low trap depths. We explored many methods for non-destructive qubit state readout and this one was the simplest, highest fidelity and most reproducible, and worked with the lowest trap depth requirements. One of the elegant aspects of this is that the bright state we use, $|2,-2\rangle$, is pumped into by the lattice - so the bright state remains bright. The dark state, $|2,+2\rangle$, can scatter lattice photons with a scattering probability of the scale of $\sim$ 1%. However, at least two photons need to be scattered to go from $|2,+2\rangle$ to $|2,-2\rangle$, so the probability of transition-

ing all the way to the bright state is suppressed quadratically (inspired by our QEC intuitions of p^2 suppression!). The other elegant aspect is that the spin-selective part, which can lead to transitions between spin states due to necessarily being near-resonant, is on only for a short period to realize spin-to-position conversion, but then the imaging is done normally. These various aspects make this the best non-destructive readout technique we implemented.

This accomplishes non-destructive readout of the qubit state and also loss detection, but more details are involved in mid-circuit re-cooling and re-initialization. Notably, conventional atomic physics methods for imaging and cooling typically use 3D beams - in a zero magnetic field - which illuminate the entire array. We have come up with simple ways to use 1D beams only for imaging and cooling in a finite magnetic field. Simply, all one has to do is simply detune the two counter-propagating circularly polarized beams with respect to each other (Figure 2.25) - this corresponds to a rotating frame, where there is a fictitious field that cancels out the real magnetic field (credit for this extremely simple solution goes to Sasha!)[40,118,36]. So, all one has to do is detune the two beams and one gets the same behavior as in zero magnetic field. Then, it appears empirically to be enough to cool all three axes using just a 1D PGC and EIT scheme, due to mixing of the different axes. We then recover very similar performance to 3D cooling.

These additional cooling lasers generate a lot of near-resonant photons flying around. To protect the data qubits further, we use a hiding laser that shifts the excited state out of resonance by coupling the excited $5P_{3/2}$ state to a $4D_{5/2}$ and $4D_{3/2}$ state[74]. This is elegant as it shifts the data qubit resonance by as much as 10+ GHz, while not affecting the ground state qubit as it is very far off-resonant from any transitions the ground state orbital sees. We are using this to suppress residual decoherence and allows us to achieve our entropy removal architecture in Figure 2.25.

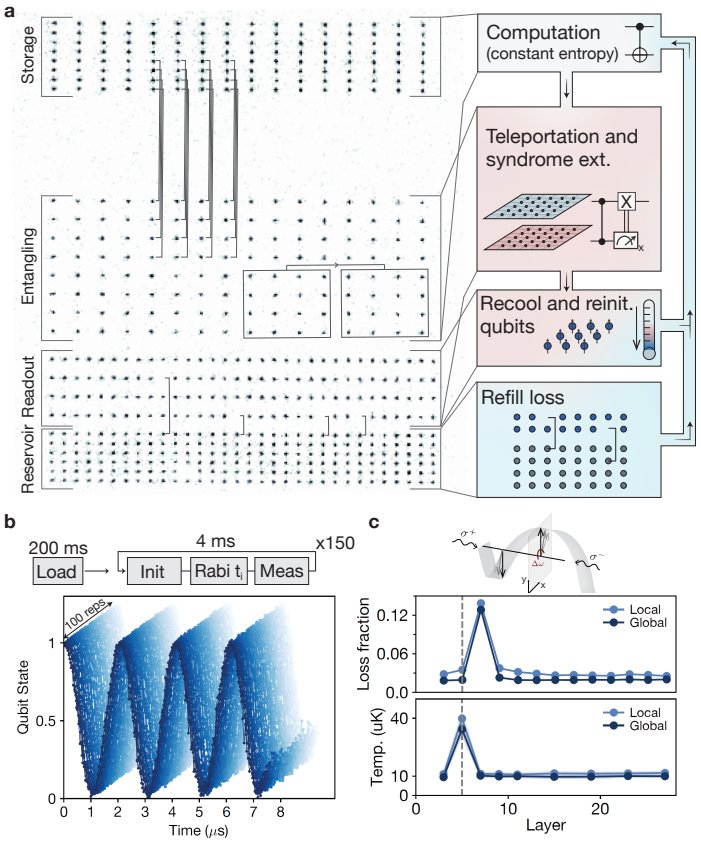


Figure 2.25: Architecture for constant-entropy computation. **a,** Image of processor and illustration of process for removing entropy generated by computation. **b,** Rabi oscillations measured by reuse of atoms with 150 cycles of non-destructive measurement and reinitialization. Each curve is one experimental repetition averaged over 200 atoms in parallel. 3D cooling methods are used in this subplot as coherence does not need to be preserved. **c,** Local cooling with 1D PGC and EIT. The finite magnetic field is compensated simply by applying a relative detuning between the two circularly polarized counterpropagating beams, lending to a rotating frame where the effective field is zero. The loss is constant as a function of time and recovers to steady state within one cycle after applying a perturbation to the system (dashed line). Additional shielding of the data qubits is provided by a 1530 shielding laser (not pictured here).

2.10 Fast cycle rates with atomic qubit re-use

Finally, the ability to re-use the atoms with the 1D lattice approach has significant implications on cycle rates. As an example of this, in Fig. 2.25b we realize fast calibrations of Raman Rabi oscillations. We load the atoms and sort them into a 200-atom ordered array in 200 ms. Then, we do a Rabi oscillation, realize the spin-selective readout with the lattice, image the atoms, and put all the atoms back in their original trap. We then repeat this 150 times, taking the t_i point of the Rabi oscillation curve with the same atoms for each t_i. Then the atoms are re-loaded and the next curve is taken. This reaches 187 Hz cycle rates (250 Hz not including the loading time). This, plus other developments of continuous operation in the field[65,108], are important for deep-circuit QEC but also enable a new generation of high cycle rate AMO experiments. These high cycle rates can help enable getting to 5-10x below QEC thresholds, from the current 2x.

2.11 Overview of system and circuit programming approach

We here overview the system again with everything pieced together, representative of the system during the experiments in Bluvstein et al Nature 2024. See Fig. 2.3 for a technical diagram.

A cloud containing millions of cold ^{87}Rb atoms is loaded in a magneto-optical trap inside of a glass vacuum cell, which are then loaded stochastically into programmable, static arrangements of 852-nm traps generated with a spatial light modulator (SLM), and then rearranged with a set of 850-nm moving traps generated by a pair of crossed acousto-optic deflectors (AODs, DTSX-400, AA Opto-Electronic) to realize defect-free arrays[10,122,49]. Atoms are imaged with a 0.65-NA objective (Special Optics) onto a CMOS camera (Hamamatsu ORCA-Quest C15550-20UP), chosen for fast electronic readout times. The qubit state is encoded in $m_F = 0$ hyperfine clock states in the ^{87}Rb ground-state manifold, with $T_2 > 1$s[19], and fast, high-fidelity single-qubit control is executed by two-photon Raman excitation[19,95]. A global Raman path illuminating the entire array is

used for global rotations (Rabi frequency $\sim$1 MHz, resulting in $\sim$ 1 μs rotations with composite pulse techniques[19]) as well as for dynamical decoupling throughout the entire circuit (typically 1 global π pulse per movement). Fully programmable local single-qubit rotations are realized with the same Raman light but redirected through a local path which is focused onto targeted atoms by an additional set of 2D AODs. Entangling gates (270-ns duration) between clock qubits are performed with fast two-photon excitation using 420-nm and 1013-nm Rydberg beams to n=53 Rydberg states, utilizing a time-optimal two-qubit gate pulse[80] detailed in Ref.[52]. During the computation, atoms are rearranged with the AOD traps to enable arbitrary connectivity[19]. Mid-circuit readout is carried out by illuminating from the side with a locally focused 780-nm imaging beam, with scattered photons collected on the CMOS and processed real-time by a field-programmable gate array, FPGA (Xilinx ZCU102), with feedforward control signal outputs.

The quantum circuits are programmed with a control infrastructure consisting of five arbitrary waveform generators (AWG) (Spectrum Instrumentation), synchronized to $<$ 10-ns jitter. The 2-channel Rearrangement AWG is used for rearranging into defect-free arrangements[49] before the circuit, the 1 channel of the Rydberg AWG is used for entangling gate pulses, the 4 channels of the Raman AWG are used for IQ (in-phase and quadrature) control of a 6.8 GHz source[95,19] (the global phase reference for all qubits) and pulse-shaping of the global and local Raman driving, the 2 channels of the Raman AOD AWG are used for displaying tones that create the programmable light grids for local single-qubit control, and the 2 channels of the Moving AOD AWG are used for controlling the positions of all atoms during the circuit. AODs are central to our methods of efficient control[10], where the two voltage waveforms (one for X-axis and one for Y-axis) control many physical or logical qubits in parallel: each row and column of the grid simply corresponds to a single frequency tone, and these tones are then superimposed in the waveform delivered to the AOD (amplified by Mini-circuits ZHL-5W-1+). The phase relationship between tones is chosen to minimize interference. Most of the system parameters used in our approach do not have hard limits, but instead result from

possible trade-offs. In what follows, as an example we detail some design decisions made for the circuits used in the logical processor work Bluvstein et al Nature 2024.

Zone parameter choices. For simplicity, we keep the entangling zone fixed for all experiments. This conveniently allows to switch between various experiments, without additional calibrations. We choose our entangling zone profile, realized by 420-nm and 1013-nm Rydberg "tophat" beams generated by SLM phase profiles[49], to be homogeneous over a 35-μm tall region. As the Rydberg beams propagate longitudinally, the entangling zone is longer than it is tall. We optimize tophats to be homogeneous over roughly 250-μm horizontal extent. Taller regions are also achievable, with a trade-off with reduced laser intensity and greater challenge in homogenization. The 250-μm width of the zones used here is set by the bandwidth of our AOD deflection efficiency. We position the readout zone on the other side of the storage zone to further minimize decoherence on entangling zone atoms.

Our two-qubit gate parameters are similar to our prior work in Ref.[52]. During two-qubit Rydberg (n=53) gates, we place atoms $\lesssim 2\ \mu$m apart within a "gate site", resulting in $\gtrsim 450$ MHz interaction strength between pairs, significantly larger than the Rabi frequency of 4.6 MHz. Importantly, due to the use of the Rydberg blockade[79,96], the gate is largely independent of the exact distance between atoms. Hence, precise inter-atom positioning is not required. Gate sites are separated such that atoms in different gate sites are no closer than 10 μm during the gate, resulting in negligible long-range interactions. Throughout this work we use 4 gate sites vertically (5 for the surface code experiment) and 20 horizontally, performing gates on as many as 160 qubits simultaneously. Under various conditions, with proper calibration we measure two-qubit gate fidelities in the range $F = 99.3\% - 99.5\%$. We do not observe any error on storage-zone atoms when Rydberg gates are executed in the entangling zone. Even though the tail of the tophat Rydberg excitation beams is only suppressed to $\sim$ 0.1x intensity, the two-photon drive is far off-resonant due to the $\sim$ 20 MHz 1013 light shift detuning which is present for the entangling zone atoms[52]. We natively

realize physical CZ gates; when implementing CNOTs we add physical H gates. We find minimal two-qubit cross-talk between gate sites, as probed with long benchmarking sequences in Ref. [52]. Although Ref. [52] appears to find some small cross-talk seemingly originating from decay into Rydberg P states, this should be considerably suppressed in the practical operation here due to the $\sim$ 200 μs duration between gates, during which time Rydberg atoms should either fly away or decay back to the ground state.

Shuttling and transfers. The SLM tweezers can have arbitrary positions, but are static. The AOD tweezers are mobile, but have several constraints [138,19]. In particular, the AOD array creates rectangular grids (but not all sites need to be filled). During the atom moving operations, they are only used for stretches, compressions and translations of the AOD trap array: i.e., atoms move in rows and columns, and rows and columns never cross [19,138]. Arbitrary qubit motions and permutation is achieved by shuttling atoms around in AOD tweezers, and then transferring atoms between AOD and SLM tweezers as appropriate. We perform gates on pairs of atoms in both AOD-AOD traps and AOD-SLM traps, with no observed difference for gate performance as measured by randomized benchmarking [52].

We find that free-space shuttling of atoms (i.e., no transfers) in AOD tweezers comes essentially with no fidelity cost (other than time overhead), consistent with our prior work [19]. Two additional improvements here are the use of a photodiode to calibrate and homogenize the 2D deflection efficiency of our 2D AODs to percent-level homogeneity across our used region, and engineering atomic trajectories and echo sequences to cancel out residual path-dependent inhomogeneities. For example, we move an atom 100 μm away to realize a distant entangling gate, and then before returning the atom, we perform a Raman π pulse, so that differential light shifts accumulated during the return trip cancel with the first trip. Motion is realized with a cubic profile as in Ref. [19], the characteristic free-space movement time between gates is roughly 200 μs, and acoustic lensing effects from the AOD are estimated to be negligible. We pulse the 1013 laser off during motion to remove loss ef-

fects from the large light shifts. Note that the 1013-induced differential light shift on the hyperfine qubit is only kHz-scale but we still ensure its effects are properly echoed out.

Transferring atoms between tweezers [14] presents additional challenges. We measure the infidelity of each transfer, encompassing both dephasing and loss, to be $\lesssim 0.1\%$. To achieve this performance, in our transfer from SLM to AOD, we ramp up the AOD tones' intensity (with quadratic intensity profile when possible) corresponding to the appropriate sites over a time of 100-200 μs to a trap depth $\sim$2x larger than the SLM trap depth, and then move the AOD trap $1 - 2\,\mu$m away over a time of 50-100 μs. These time scales can likely be shortened considerably while suppressing errors using optimal control techniques. During subsequent motion we leave the AOD trap depth at this 2x value. To transfer an atom AOD to SLM we perform the reversed process. During these transfer processes, the differential light shifts on the transferred atoms are dynamically changing, and can result in large unechoed phase shifts. As such, whenever possible we engineer circuits such that pairs of transfers will echo with appropriately chosen π pulses. When echoing pairs of transfers is not possible, we perform 1 cycle of XY4 or XY8 dynamical decoupling during the transfer. Finally, we note that low-loss transfer is highly sensitive to alignment of the AOD and SLM grid. We fix small optical distortions between the AOD and SLM tweezer grids by fine adjustment of individual SLM grid tweezers, which can be arbitrarily positioned, to overlap with the AOD traps as seen on an image plane reference camera. It is important to adjust the SLM and not the AOD, as small adjustments of individual AOD tones deviating from a frequency comb causes beating and atom loss.

Dynamical decoupling and local gates. In our circuit design, we engineer our echo sequences in order to cancel out as as many deleterious aspects as possible. We ensure that in our dynamical decoupling we have an odd number of π pulses between CZ gates (whenever possible), as this echoes out both systematic and spurious contributions to the single-qubit phase [19,52]. We apply appropriate $X(\pi)$ and $Z(\pi)$ rotations between local addressing with the local Raman to cancel out errors induced by the global $\pi/2$ pulses, as well as between pulses of the 420-nm laser (when used for en-

tangling zone single-qubit rotations [19]) to echo out small crosstalk experienced in the storage zone by the tail of the 420-nm beam. For our global decoupling pulses we use both BB1 pulses [148] and "SCROFULOUS" pulses [38]. To benchmark and optimize coherence during our complex circuits, we perform a Ramsey fringe measurement encompassing the entire movement and single-qubit gate sequence and optimize the observed contrast [19]. When performing properly, our total single-qubit error is consistent with SPAM [52], an effective coherence time of 1-2 s, and the Raman scattering error of all the Raman pulses [95,19]. We note that these measured coherence time include the movement within and between zones; although we use fewer pulses (typically 1 per movement) than the XY16-128 sequence used to benchmark 1.5s coherence in Ref. [19], the coherence times here are naturally longer due to further-detuned tweezers used (852 nm rather than 830 nm.

Local single-qubit gates [63,9] with the Raman AOD are realized in arbitrary positions in space on both AOD and SLM atoms. Targeted logical qubit blocks are addressed by a grid illumination of the logical block. Arbitrary patterns of rotations on the qubit grid (e.g., during color code preparation) are realized with row-by-row serializing, with the targeted x coordinates in each row simultaneously illuminated. The duration of each row is 5-8 μs (i.e., several 10s of μs for an arbitrary pattern of rotations), which can be sped up considerably as discussed in the next section. For simplicity we carefully calibrate rotations on 80-160 specific sites across the array, but also perform rotations in arbitrary spots utilizing the nearest calibrated values.

With the local single-qubit gates and entangling zone two-qubit gates calibrated, the entire circuit is simply defined by the appropriate trapping SLM phase profile, and waveforms for our several AWG channels and TTL pulse generator. These several channels then program complex, varied circuits on hundreds of physical qubits.

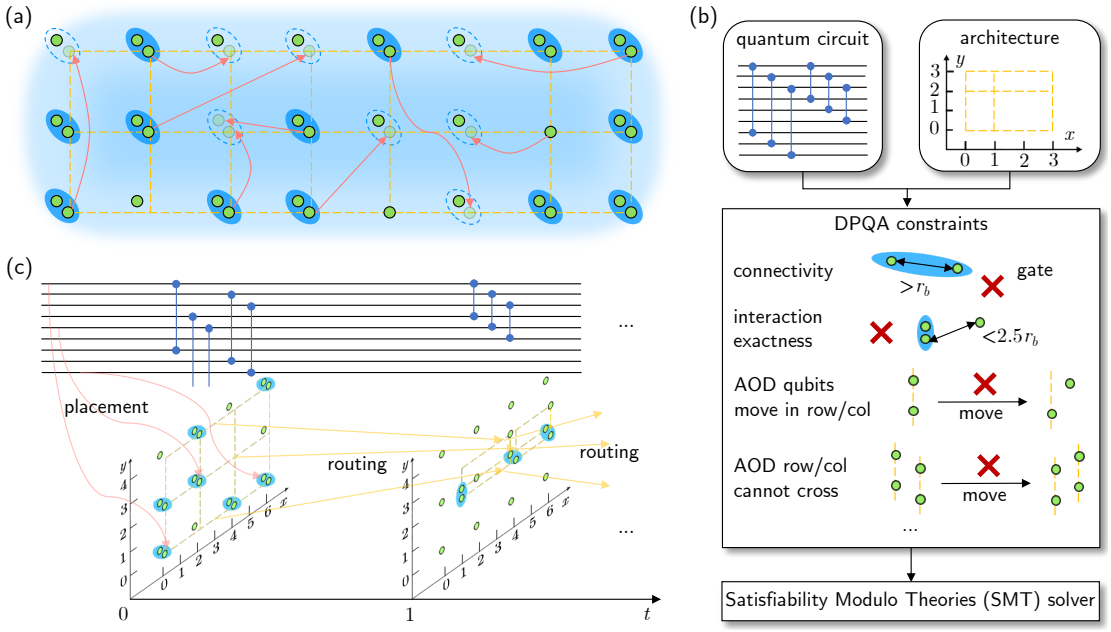


Figure 2.26: Compiling quantum circuits to reconfigurable atom arrays. With our collaborators at UCLA (Jason Cong lab), we developed compilers for atom movements with our currently existing tools[138]. Although not presently used in our experiments, these computational tools are making rapid progress. See Ref.[138] for additional details on this first work as shown in this figure, and Refs.[139,97] for recent more advanced and scalable methods.

3

Quantum processing with logical qubits

In the following sections, we will discuss various key ideas that we used in, and discovered from, our experiments in quantum processing with logical qubits. We will then provide more detailed description of the experimental studies that sometimes came out of, and sometimes inspired, these various key ideas. We've studied different ways of entangling logical qubits, different ways of removing entropy from them, different ways of realizing universal and computationally complex processing, and many different types of encodings.

3.1 In classical computing we operate on bits. In quantum computing we operate on *logical* qubits.

Perhaps the most striking difference between classical and quantum computing architecture design is that classical information is intrinsically digital and is intrinsically error-corrected[91]. In particular, dissipation can be used to correct small imperfections and return all configurations to the same value[91]. For this reason, when we do processing in our classical computers, the errors are so low that we often don't bother explicitly encoding them in error-correction codes (barring notable examples like devices in space). Quantum information needs quantum error correction to enjoy the same digital-type protection, and requires entanglement to create logical qubits. The real analogue to *bits* in classical computing is *logical qubits* in quantum computing. The physical qubits are simply the physical constituents that form the "fabric" of the processor where the logical qubits live via entanglement.

We should then find ways to operate on individual logical qubits as the fundamental units of our processor, instead of individual physical qubits. This turns out to mulitplex remarkably well with optics, and that we can do all of our operations in ways that are entirely parallel.

3.2 Transversal gates

Consider the color code or Steane code in Figure 3.1 (see first chapter for more description). This is a *color code*, which for its distance-3 instantiation is the same as the [[7,1,3]] Steane code. Its logical operators are products of X's and Z's along any of the boundaries, and the stabilizers are products of 4 X's or 4 Z's on the colored plaquettes [136,88].

We don't want to just correct errors like was described in the first chapter, we want to realize *logic operations*. For instance, we would like to realize a Hadamard single-qubit gate, where for physical qubits applying H transforms $X \leftrightarrow Z$, and we would then want that a *logical* Hadamard gate trans-

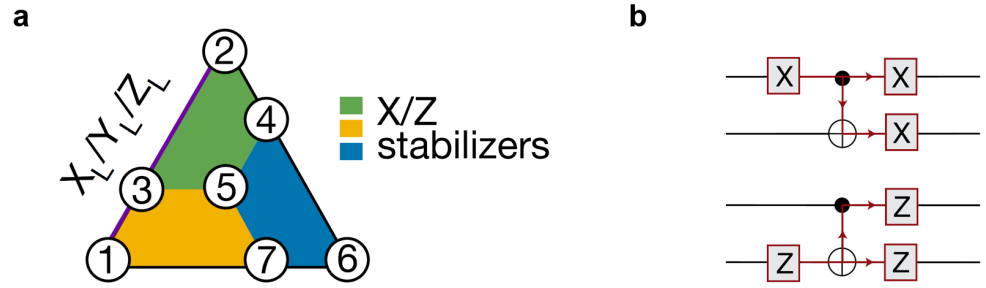


Figure 3.1: Steane code and CNOT propagation. **a,** Steane code logical operator and stabilizers. **b,** Propagation of operators through a CNOT. By applying pairwise CNOTs to the Steane code (and all CSS codes), one realizes a logical CNOT.

forms $X_L \leftrightarrow Z_L$. In the Steane / color code, this logical Hadamard can be realized in a remarkably simple way: just by realizing independent Hadamard on all the physical qubits of the block, e.g. in the Steane code $H^{\otimes 7}$ interconverts $X_L = X_1X_2X_3$ and $Z_L = Z_1Z_2Z_3$ [132]. It's also important that our logical operations preserve the codespace and don't introduce a bunch of new errors. We can see that $H^{\otimes 7}$ preserves the codespace as the stabilizers are also just overlapping products of X and Z, and so the stabilizers interconvert and if all stabilizers = +1 (or are tracked as such) then nothing changes. One can verify that the logical phase S_L gate (i.e. a 90-degree rotation around the Z-axis of the Bloch sphere) can also be realized as $S^{\otimes 7}$, and appropriately interconverts the X_L and Y_L operators while preserving stabilizers (upon multiplying the resulting Y stabilizers with the Z stabilizers to return back to X stabilizers). So we see that we can realize these operations on our logical qubit simply as tensor products of that same identical operation on the physical qubits of the block.

Importantly, this approach naturally multiplexes with optical techniques. For example, to realize a logical single-qubit gate [132], we use the Raman 2D AOD to create a grid of light beams and simultaneously illuminate the physical qubits of the logical block with the same instruction. Such a gate is *transversal* [132], meaning that operations act on physical qubits of the code block independently. This transversal property further implies the gate is inherently fault-tolerant [132], meaning that errors cannot spread within the code block (see first chapter), thereby preventing a physical error from

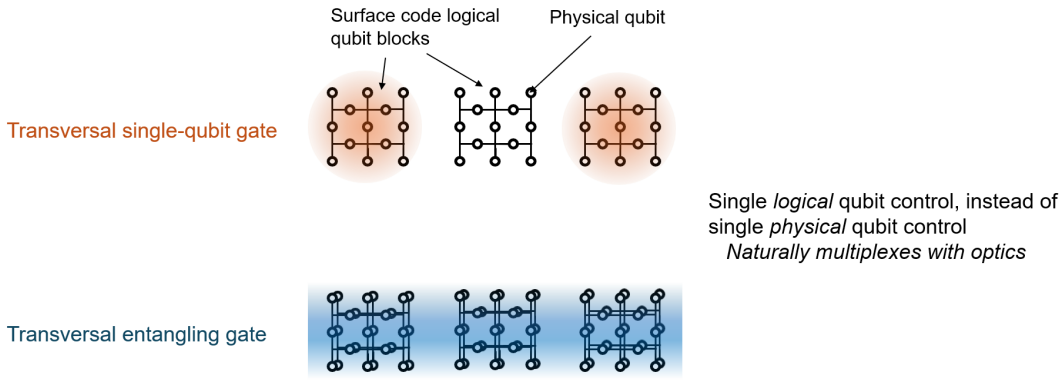


Figure 3.2: Logical operations with parallel control. In performing computation with logical qubits, one can realize operations with parallel controls. For instance, to realize a logical single-qubit gate one does not need to realize a unique operation to each individual physical qubit in the block, but instead can realize that operation in parallel on all the physical qubits of the block. Such an operation is transversal, and is natively fault-tolerant. The same is true for entangling gates: the grids of two logical qubits can be interlaced and a single global pulse of the Rydberg laser can realize a transversal CNOT in a single step.

spreading into a logical fault. Conceptually, this approach is very appealing as it allows us to realize a high-fidelity, fault-tolerant logical operation simply by taking (effectively) one big "Hadamard beam" and illuminating our one "big qubit" with this big beam.

Looking at our Steane code example, we can work out that a logical CNOT gate can also be realized transversally[137]. A CNOT gate takes X on the control and propagates it onto the target, i.e. $XI \to XX$, and takes Z on the target and propagates it onto the control, i.e. $IZ \to ZZ$ (Fig. 3.1), and a logical CNOT should do the same thing. Since again, our logical operators are given just by products of X's or products of Z's, then if we take two codes and apply a transversal CNOT between them, it will map X on each control qubit onto the target (and vice versa for Z), and so will just map $X_L I_L \to X_L X_L$ and $I_L Z_L \to Z_L Z_L$ as intended. Similarly, since the stabilizers directly map between the two codes in the same fashion as they are also products of X and Z, this operation also preserves the codespace. *So one can realize a logical CNOT by just applying pairwise physical CNOTs between each twin qubit of two blocks*. This example here is for a Steane code, but the same is true for a surface code: the X_L and Z_L are cleanly copied between the two code blocks, and the codespace is again pre-

served. In fact, the very large family of codes known as Calderbank-Shor-Steane (CSS) codes all have a transversal CNOT [132,33,136]. This is because CSS codes have stabilizers that are products of exclusively X or exclusively Z, and also logical operators that are products of exclusively X or exclusively Z. This is because CSS codes are constructed from two classical codes - one for correcting X-type errors and one for correcting Z-type errors - and the CNOT is a classical gate (i.e., CNOT just says if control is in state 1, then apply a bit flip to the target).

Crucially, this means a similar approach of parallel optical control can realize logical entangling gates [132,44]. Specifically, we use the grids generated by our moving 2D AOD to pick up two logical qubits, interlace them in the entangling zone, and then pulse our single global Rydberg excitation laser to realize a physical entangling gate on each twin pair of the blocks (Figs. 3.2,3.3). This process realizes a high-fidelity, fault-tolerant transversal CNOT in a single parallel step.

3.3 The importance of the transversal CNOT

The transversal CNOT has both practical and fundamental implications. From a practical perspective, it is an exceptionally convenient way to entangle two logical qubits, and can be realized with a single parallel motion of the atoms and pulse of the Rydberg laser. From a fundamental perspective, there are physical reasons as to why it is a particularly elegant and direct way of entangling logical qubit degrees of freedom.

Conceptually, the way that error correction works is that we use entanglement to delocalize a logical qubit degree of freedom across many physical qubits. But now, if our two delocalized wavefunctions are placed next to each other on a 2D chip, for example, it can be extremely challenging to get these to entangle. In order to do this, we would have to use methods such as "lattice surgery" [72], which is much more complex and has a much higher overhead, as it requires having this delocalized information interact through the boundary between the two codes.

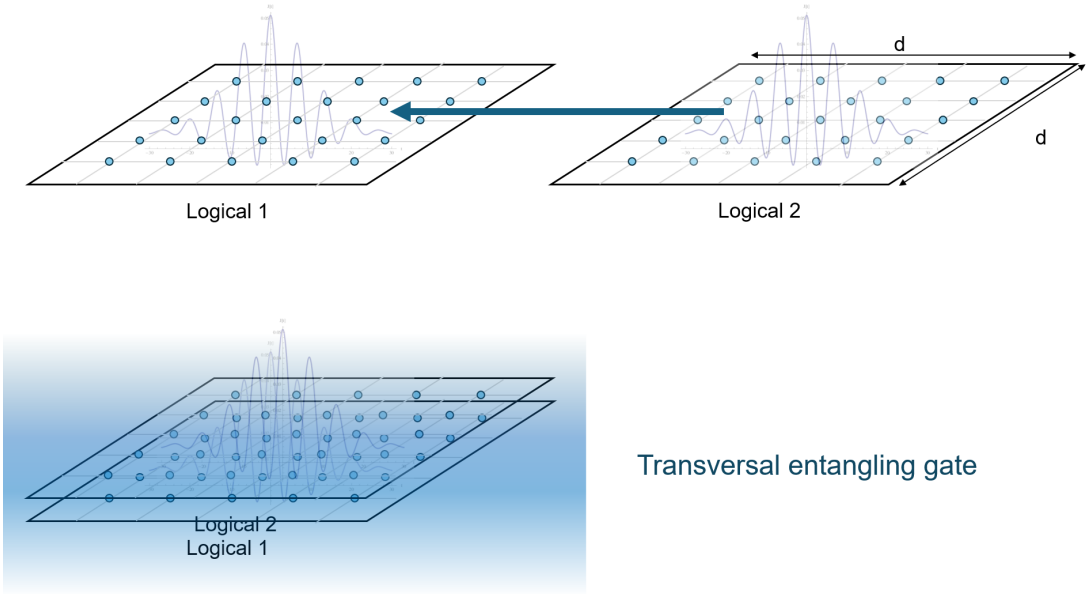


Figure 3.3: Transversal CNOT. The transversal CNOT allows entangling the two delocalized wavefunctions while remaining delocalized and hidden among the many physical qubits of the block.

Conversely, with the transversal CNOT we can just take these two delocalized wavefunctions, put them right on top of each other, and entangle them while they are still delocalized and protected. Importantly, as will be discussed more later, in this approach the logic operations are done *directly* between the data qubits, as opposed to being mediated via ancilla stabilizer measurements. The fact that the transversal CNOT realizes logic direcly between data qubits results in significant simplifications to our error correction.

3.4 Zoned architecture and efficient control

Classical computers leverage abstractions at many different levels in order to enable efficient control. For example, in order to realize processing on billions of bits with only a few wires, one uses different regions of the processor for storage and for logic. Concretely, one takes bits stored in the RAM, which can be addressed with a logarithmic number of wires, then shuttles this information into the

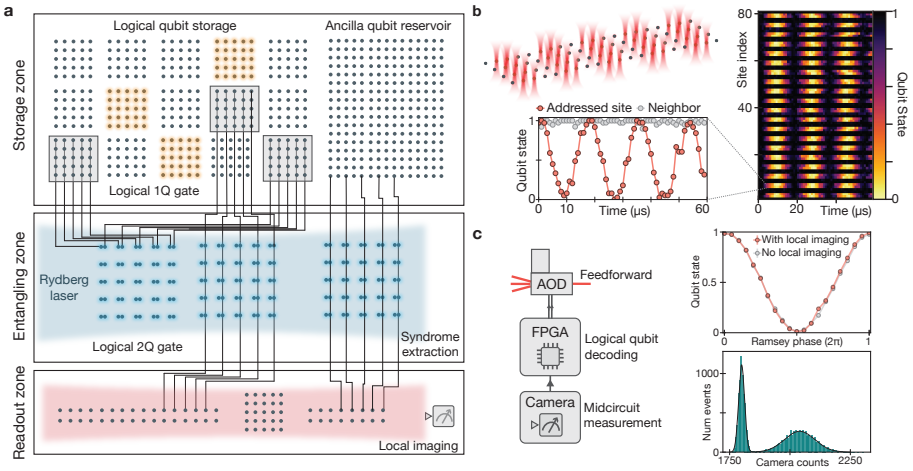


Figure 3.4: A programmable logical processor based on reconfigurable atom arrays. a, Schematic of the logical processor, segmented into three zones: storage, entangling, and readout. Logical single-qubit and two-qubit operations are realized transversally with efficient, parallel operations. Transversal CNOTs are realized by interlacing two logical qubit grids and performing a single global entangling pulse that excites atoms to Rydberg states. Physical qubits are encoded in hyperfine ground states of ^{87}Rb atoms trapped in optical tweezers. **b,** Fully programmable single-qubit rotations are implemented using Raman excitation through a 2D AOD; parallel grid illumination delivers the same instruction to multiple atomic qubits. **c,** Mid-circuit readout and feedforward. The imaging histogram shows high-fidelity state discrimination (500 μs imaging time, readout fidelity $\approx$ 99.8%, Methods), and the Ramsey fringe shows that qubit coherence is unaffected by measuring other qubits in the readout zone (error probability $p \sim 10^{-3}$, Methods). The FPGA performs real-time image processing, state decoding, and feedforward (Fig. 4).

CPU / algorithmic logic unit (ALU) for processing, and then shuttles bits back into the RAM [111].

We were directly inspired by this to develop processing regions or zones in our quantum processor in the aim of developing efficient classical control. Discussions with Michel Devoret were particularly helpful. Our zoned architecture is depicted in Figure 3.4. The storage zone is used for dense qubit storage, free from entangling-gate errors and featuring long coherence times. The entangling zone is used for parallel logical qubit encoding, stabilizer measurements and logical gate operations. Finally, the readout zone enables mid-circuit readout of desired logical or physical qubits, without disturbing the coherence of the computation qubits still in operation.

This now gives very convenient ways for operating on our logical qubits as the fundamental units of the processor. We can take our logical qubit grids and shuttle them around the processor, for

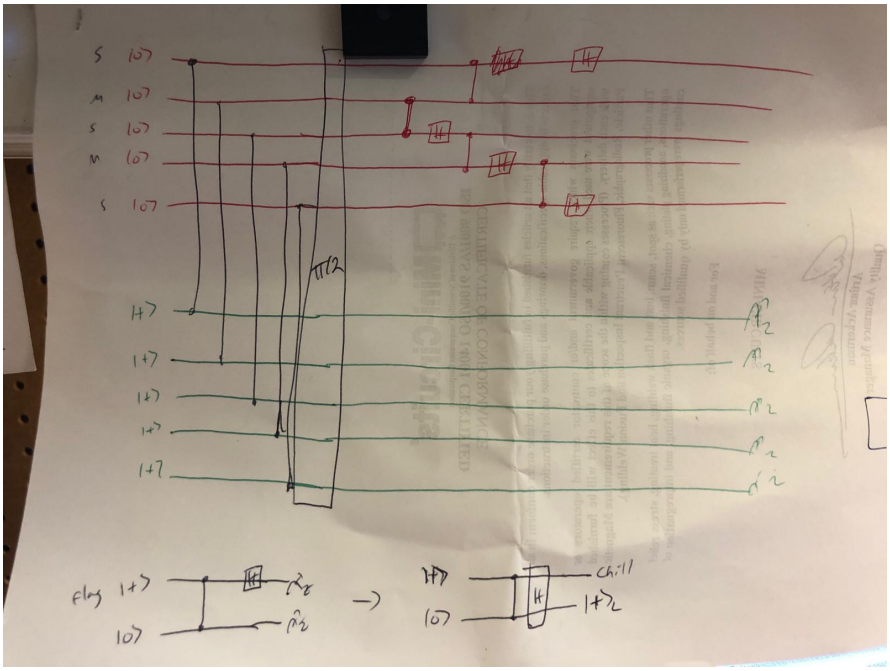


Figure 3.5: Hand-written note for creating logical GHZ state with Steane codes. Parallel control over logical qubit blocks as the fundamental units of the processor greatly facilitates programming logical circuits.

example picking up two logical qubits and moving them into the entangling zone to realize a logic operation, or selectively moving logical blocks to readout. The logical-level control, combined with the zoned abstraction, enabled us to study many different quantum circuits and error-correction codes.

3.5 Stabilizer measurement and removing entropy

We have discussed so far how to do logic with the use of transversal operations. But as we do more and more logic gates, we will create more errors: computation is never perfectly precise, and thereby necessarily generates entropy[91,12]. The error-correction code then allows us to remove errors from the physical qubits without corrupting the state of the logical qubits. From a physics perspective, the computation generates entropy, which we then need to remove in one way or another. Remark-

ably, unitary evolution on the logical qubit degree of freedom proceeds coherently, while dissipative processes on the physical qubit / stabilizer degrees of freedom serve to remove entropy.

There are many different instantiations of how to do this entropy removal, but they all involve some form of coupling 'hot' degrees of freedom used in the logic to 'cold' degrees of freedom for removing errors, and then resetting the qubits. The standard example is that of stabilizer measurement, where ancillary physical qubits are used to do sequential physical gates to data qubits, thereby mapping the value of the stabilizer onto the ancilla [2]. Measuring the ancilla, now, allows us to determine the value of the stabilizer and thereby removes entropy from the system.

While this is the standard example, there are other methods as well for keeping entropy constant, and what works best for a particular setting is dependent on context. A simple extension is to take multiple ancillas and entangle them into a GHZ state that allows extracting the stabilizers from the data qubits in a single step. This is so-called Shor-style syndrome extraction [132], and can have benefits such as use of flags, or practical benefits such as realizing fewer gates on the data qubits which may e.g. get heated from doing too many gates. Another option is to initialize an ancilla *logical* qubit with known stabilizers, and then realize a transversal CNOT between the data logical qubit and the ancilla logical qubit. All the data logical qubit stabilizers (in one basis) now directly map from the ancilla logical qubit, which can then be projectively measured (and then this can be repeated for the other basis as well). In certain contexts, this so-called Steane-style syndrome extraction [137] can have higher thresholds than typical stabilizer syndrome extraction [76] (although it often ends up being very similar / identical if the ancilla logical qubit was just prepared with stabilizer measurement anyway).

Another option to keep entropy constant is not to remove physical entropy from your data logical qubit, but to teleport the logical information from one logical qubit to another. This bares the benefit that the atoms themselves also change, and so all types of errors (e.g., loss, leakage, heating) that may have been present on your "used" qubit get left behind on that old block, while the process

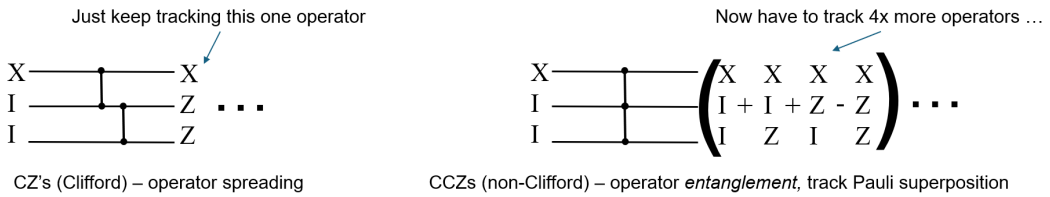


Figure 3.6: Clifford vs non-Clifford. Clifford gates like a CZ gate, by definition, take Pauli operators to other Pauli operators (operator spreading). Non-Cliffords like a CCZ gate, take Pauli operators to superpositions of Pauli operators (operator entanglement). Clifford circuits can be easily simulated by tracking the Pauli operators, whereas adding non-Cliffords exponentially increases the simulation complexity with each non-Clifford gate added (under assumptions).

of decoding and feedforward (which can be done exclusively in the software, see below) propagates the logical qubit information onto the new block. This will come back later as well, and turns out to be a very convenient way of thinking about quantum processing when one has transversal operations, because transversal operations allow us to reduce the number of rounds of error correction to the point that for universal processing we roughly have $\sim O(1)$ teleportation per unit step anyway.

3.6 Universal processing. Clifford and non-Clifford operations.

One of the big things we have swept under the rug so far is that error correction codes only allow a discrete set of gates to be directly implemented[47]. This is not unique to the fact that we realize these circuits transversally instead of with lattice surgery, and is in fact fundamental. As we discussed previously, we need all operations to be discrete and digital in order to be able to correct errors.

In 2D topological error correcting codes such as the surface code or 2D color code, only the *Clifford* gates can be done directly / transversally, i.e. $\{H, S, CNOT\}$, where H is the Hadamard gate (creating superpositions on qubits), S is a 90-degree phase gate around the Z-axis of the Bloch sphere, and $CNOT$ is the entangling two-qubit gate (note that S is not directly transversal on the surface code but is on the 2D color code). By definition, Clifford gates map Pauli operators to other Pauli operators. This has multiple important consequences. First, while these gates can generate

superpositions on *qubits*, they do not generate superpositions on *operators*. This means that the state of the whole computation of N qubits can be tracked by simply tracking the N operators as they propagate throughout the Clifford circuit - while it is not a classical circuit on the level of the qubits, it is a classical circuit on the level of the operators [1]. Another related consequence is that these operations are insufficient for realizing universal processing. For universal processing we need a non-Clifford gate, which can take Pauli operators to superpositions of Pauli operators: generating operator entanglement in addition to operator spreading. The two canonical examples are the T gate, which is a 45-degree rotation around the Z-axis of the Bloch sphere, and a *CCZ* gate, which realizes a *CZ* gate between two qubits controlled on the third qubit being in state $|1\rangle$.

More explicitly, circuits composed solely of stabilizer states and Clifford gates can be simulated in polynomial time due to the Gottesman-Knill theorem [1]. This can be understood as stabilizer tracking: for example, consider a 3-qubit system where a stabilizer of the state is $X \otimes I \otimes I$, where X stabilizes the $|+\rangle$ state and I is the identity. Applying two CZ entangling gates $\mathrm{CZ}_{1,2} \otimes \mathrm{CZ}_{1,3}$ transforms this stabilizer to $X \otimes Z \otimes Z$, since an X flip before the CZ simply changes whether a Z flip will be applied to the other qubits. Even though Clifford circuits create superposition and entanglement between qubits, the N initial stabilizers of the state can simply be tracked as they propagate through the circuit (so-called operator spreading [105]), and thereby simulation of the circuit can be easily accomplished.

The effect of non-Clifford gates, however, is significantly more complex. For example, passing the stabilizer $X \otimes I \otimes I$ through a CCZ maps into a superposition of Pauli strings, i.e. $X \otimes I \otimes I \rightarrow 1/4(X \otimes I \otimes I + X \otimes Z \otimes I + X \otimes I \otimes Z - X \otimes Z \otimes Z)$, as an X flip now changes whether a CZ operator will be applied on the other qubits, resulting in 4x more operators to track after the single CCZ. (The CZ operator matrix is simply equal to $1/4$ $[I \otimes I + Z \otimes I + I \otimes Z - Z \otimes Z]$). This causes not only operator spreading, but also so-called operator entanglement [105]. As we apply additional non-Clifford gates, the number of operators to track will grow exponentially, and eventually will become

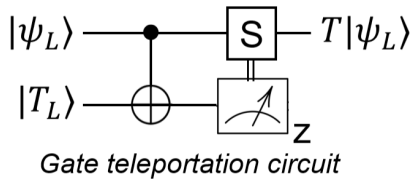


Figure 3.7: T gate teleportation circuit. In 2D codes like the surface and color codes, T gates can be implemented via synthesizing T states and then realizing gate teleportation.

computationally intractable. E.g. state-of-the-art Clifford + T simulators can handle roughly 16 CCZ gates[23]. This is the basis behind our complex sampling circuits with 3D codes in the next chapter.

When we use 2D codes such as the surface code, we cannot realize transversal non-Cliffords. So what we end up doing, then, is synthesizing a resource state, such as a magic state or $|T\rangle$ state, and then teleporting it onto our surface code[57]. This can be realized by the standard circuit shown in Fig. 3.7 with a CNOT gate followed by a feedforward logical S gate. There are fundamental reasons why it is harder to make states that have magic than states that do not, and broadly as we will discuss later, we need much more entanglement within our code blocks for having magic or nonCliffordness. But for now, a key idea is that one can synthesize the $|T\rangle$ state on the side and then teleport it in. Critically however, the teleportation circuit itself is still Clifford, and this will have multiple important implications.

Armed with a universal gate set such as $\{H, T, CNOT\}$, one can decompose universal unitary operations with the discrete digital gate set. Each operation is precise and error-corrected, and these can build analog, continuous angles. This remarkable property is known as the Solovay-Kitaev theorem[85,41], and it allows approximating a unitary transformation such as an arbitrary single-qubit rotation $R(\theta, \varphi)$ with a logarithmic overhead. This then allows us to realize analog-type angles with digital precision and correction - the wave-particle duality mentioned in the first chapter. This will be experimentally explored in Chapter 5.

3.7 Do we need to flip physical qubits back in hardware?

Upon having measured our stabilizers, we know where the physical errors are in the system, and we've also collapsed coherent-type errors into incoherent-type physical Pauli X and Z errors. There is a very natural question: we have already removed entropy from the system, both by projecting coherent errors into incoherent ones (e.g. Zeno effect, see later) as well as finding them and noting their locations, so our system is now in a known state. So, do we need to flip the physical qubits back *in hardware* such that all stabilizers = +1? The answer depends on the function of the stabilizers. In the conventional situation with 2D error correction codes like the surface code and color code, the answer is no [57], because the role of stabilizer measurement here is simply removing entropy (discussed more below). In situations with more 'advanced' QEC codes where transversal non-Cliffords are allowed, then in fact one does have to guarantee that all stabilizers are defined as +1 (either through deterministic state creation or by applying physical feedforward flips), as here the stabilizer constraint and entanglement serves an additional role of constraining the system to allow universal processing (discussed later).

In the conventional case of the surface code and color codes, universal processing is achieved by synthesizing and teleporting in magic states or T states (a 45 degree rotation around the Z-axis of the Bloch sphere). In combination with the native Clifford gates in the surface code, this achieves universal circuits with a discrete, digital gate set of $\{H, T, CNOT\}$. However, although the circuit on the *logical qubit level* is universal and non-Clifford, the circuit on the *physical qubit level* is still just Clifford. This is because the teleportation circuit is still just a Clifford circuit, and it is only because the ancillary logical qubit is in logical state T that one realizes the non-Clifford operation.

As the whole physical circuit is always Clifford, then the physical Pauli-type errors can trivially be tracked in-software through all physical Clifford gates. So one needs to *know* where the errors are in each moment and time and track them (or moreso do a best-effort which is sufficiently below

threshold), but one does not have to physically correct qubits back in hardware. In fact, properly tracking this error propagation is the correct way to decode these transversal circuits and allows us to reduce the number of error correction rounds per logic operation we need by a factor of roughly d.

Although we do not have to physically flip back physical qubits in-hardware, for leakage errors such as qubit loss, these of course need to be corrected in one way or another in-hardware. Otherwise, all the qubits will eventually be lost and disappear from the system! These do not necessarily require feedforward operations either, however. The key is to keep the entropy constant, as we will show later.

Note, this does not mean that one does not need to do feedforward operations, it only means that one does not need feedforward on individual *physical* qubits. But you still need feedforward operations on *logical* qubits in the above scheme, to do the teleportation operation. By the time we do this readout, we need to also know the previous stabilizer measurements we did in order to be able to correctly infer our logical qubit state and realize the feedforward logical operation. But we see here that even for this universal processing, we do not need individual physical qubit control for flipping individual qubits back in-hardware, as we are just directly operating on the logical qubits.

3.8 Considering algorithms as a whole, and correlated decoding

Additional details for the description in this section can be found in Refs.[17,32,158], and Cain et al 2025.

One of the most significant theoretical advances (so far!) that has emerged from our experiments has been that of correlated decoding[17,32,158]. The idea behind correlated decoding is that, if we're doing transversal gates between our logical qubits, it makes a lot more sense to decode them together than separately. This is because transversal operations copy errors between blocks in a deterministic way[15], and the decoding process - where we infer what happened in the system based on the pattern of measured stabilizers - can use much more information if the correlations between the blocks are

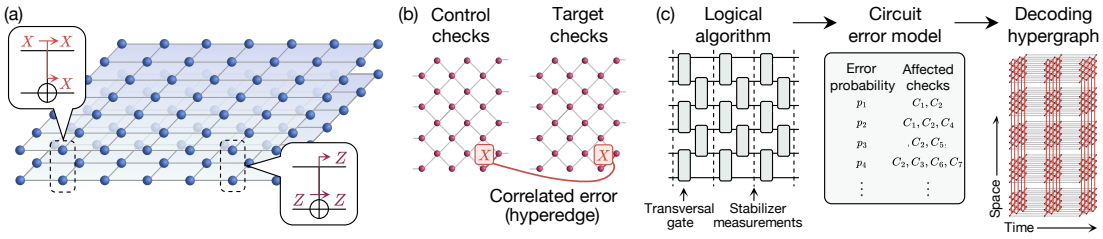


Figure 3.8: Decoding logical algorithms with transversal gates. (a) Transversal CNOT gates copy physical errors between logical qubits. (b) These transferred errors flip the same stabilizers on both logical qubits after the CNOT, generating a hyperedge in the decoding hypergraph connecting the control and target check vertices. (c) Given a physical error model, the decoding hypergraph for a logical algorithm can track these transferred errors, which can appear as hyperedges connecting checks from multiple logical qubits at different points in time.

taken into account. The idea here is very similar to the Steane error correction described above, but here one tracks the stabilizer propagation between the computational logical qubits, as opposed to a separate ancilla logical qubit.

Finding the "ground truth" of the correlated decoding was a journey. The upshot, at the end of the day, is that one can use correlated decoding to realize universal quantum computing with only O(1) stabilizer measurements round per operation, as opposed to the typical d rounds. This comes without additional complexity on the classical decoding side compared to a conventional approach, and in fact corresponds to even less resources due to the reduction of the number of rounds by $O(d)$. The procedure that works for universal computation is precisely the procedure we used in the Bell state experiment presented below, which turns out to be the simplest possible example that accomplishes the key idea.

The key upshot here is that in a transversal Clifford circuit, decoding the logical stabilizer (which propagates throughout the Clifford circuit) as opposed to individual logical qubits, becomes a problem that is very similar to decoding a single logical qubit. The correlated decoding effectively just tracks this propagating operator. Like a single-qubit memory, this can be decoded efficiently with matching in the case of the surface code. It natively does not require d rounds of stabilizer measurements, because any measurement errors on the ancilla just propagate 'alongside' the logical operator.

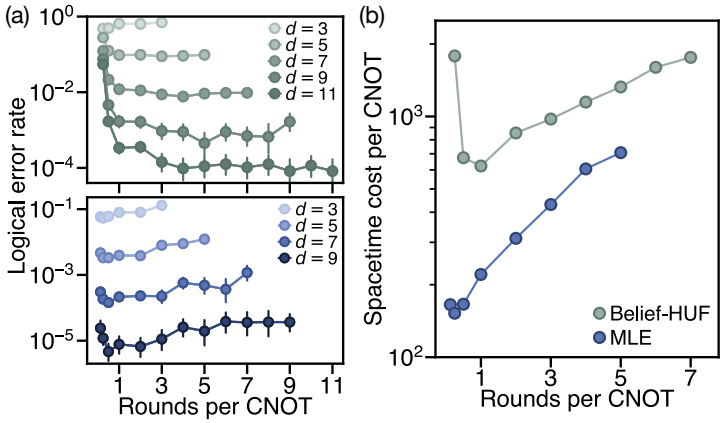


Figure 3.9: Reducing the space-time cost of deep logical circuits. (a) The rounds of syndrome extraction per CNOT can be reduced to $n_r \simeq \frac{1}{2}$ for MLE (bottom) and $n_r \simeq 3$ for BHUF (top), without increasing P_L. (b) The extrapolated space-time cost to reach $P_L = 10^{-6}$ is minimized at $n_r = 1$ for BHUF and $n_r = \frac{1}{4}$ for MLE.

And it is fully compatible with realizing universal computation with only 1 round per logical operation, as this is realized as a Clifford circuit on Pauli and $|T\rangle$ inputs, and the only measurements (or products of measurements) that one is interested in are the ones which are guaranteed to be insensitive to one is interested in products of measurements that are guaranteed to be insensitive to initialization errrors on the logical Pauli's.

3.9 Fast correlated decoding of transversal logical algorithms

Additional details for the description in this section can be found in Cain et al 2025, Fast correlated decoding of transversal logical algorithms.

Jointly decoding logical algorithms with transversal gates can reduce the number of syndrome extraction rounds by a factor of the code distance d. Here, we reformulate the problem of decoding transversal logical circuits by directly tracking relevant logical operator products as they propagate throughout the circuit, yielding a decoding problem closely akin to that of a single-qubit memory propagating through time. This procedure results in significantly simpler correlated decoding, re-

ducing the problem size while maintaining high performance. Specializing to the surface code, this method enables fast decoding using minimum-weight perfect matching.

Universal quantum computation can be performed via an adaptive transversal Clifford circuit acting on logical Pauli states and magic state inputs $|\bar{T}\rangle = \frac{1}{\sqrt{2}}(|\bar{0}\rangle + e^{i\pi/4}|\bar{1}\rangle)$ (Fig. 3.10(a)) (overline indicates logical qubits). Interestingly, such universal circuits can be fault-tolerantly implemented with $O(1)$ SE rounds per transversal logical operation, even in the presence of mid-circuit measurements and conditional gates[158]. The essence of this can be understood by tracking how logical Pauli operators propagate through the circuit. Because the actual circuit executed in any given run is a transversal Clifford circuit (acting on Pauli and non-Pauli inputs), such operators can be tracked deterministically. For each Pauli initial state, there will be an associated logical Pauli stabilizer. Measurements which *anti-commute* with any of the propagated logical Pauli stabilizers are guaranteed to be 50/50 random and thereby provide no information: individually, they do not even need to be decoded. All other measurements, or more generally *products* of measurements, *commute* with all logical Pauli stabilizers: their back-propagation terminates on $|\bar{T}\rangle$ states, which have been fault-tolerantly prepared, or logical Pauli states in the same basis, which have reliable syndrome information in this common basis. These measurement products contain non-trivial information, and must be decoded reliably. By doing so, we can reproduce the logical measurement distribution of the ideal circuit.

Critically, we find here that these logical products of interest can be reliably predicted using only the stabilizer measurements along the back-propagation of the operator through the circuit (Fig. 3.10(b)). This approach vastly simplifies the decoding problem, maintaining high accuracy while achieving speeds comparable to the logical memory setting. With each new logical measurement, we use the following steps to identify appropriate logical products and interpret the measurement result:

1. Find a logical Pauli product including the current measurement, which commutes with all

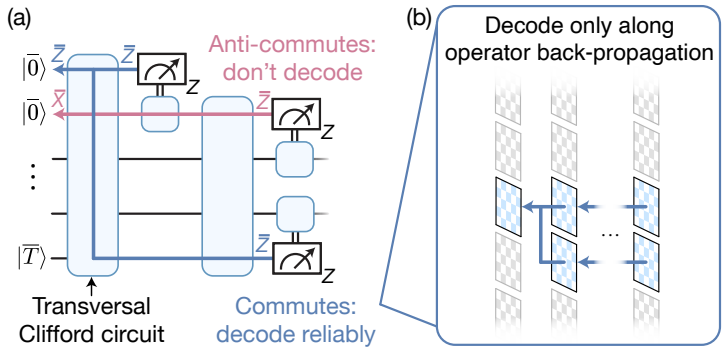


Figure 3.10: Decoding transversal algorithms. (a) Logical measurements in universal computation can be correctly predicted by decoding logical Pauli products which, when back-propagated through the circuit, commute with all logical Pauli initializations. Measurement products which anti-commute with a logical Pauli initialization are 50/50 random and can be assigned randomly. (b) The commuting logical Pauli products can be decoded by tracing back their evolution through the circuit and decoding only the stabilizers along the resulting path. The decoding problem for these operators bears similarity to decoding a single logical qubit, where all time-like edges are simple edges, and in the case of the surface code, MWPM can be applied.

logical Pauli initializations. This can be achieved by solving a system of linear equations. If one does not exist, assign the measurement randomly and skip the remaining steps.

2. Back-propagate the logical Pauli product through the circuit, recording all stabilizer measurements in the same basis along the propagation path.

3. Decode the logical Pauli product using these relevant stabilizer measurements. In the case of the surface code, this results in a decoding problem that can be solved with matching.

4. (Optional) In certain cases, the corrections can be committed in software, reducing the size of future decoding problems.

5. Assign the current logical measurement result using the decoded value of the logical Pauli product and any prior measurements involved in the product.

This method allows us to realize universal computation with surface code and transversal gates, using only $O(1)$ stabilizer measurements per logic gate while maintaining high thresholds, and lead-

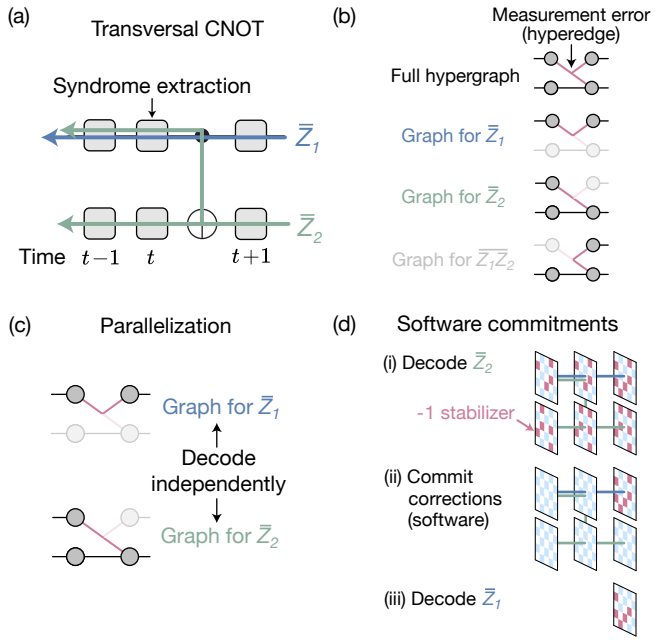


Figure 3.11: Matchable decoding and correction strategies. (a) Logical Pauli products are decoded using stabilizer measurements along their back-propagation through the circuit. (b) During a transversal CNOT, the checks (vertices) are $Z_1^{t-1}Z_1^t$ (top left), $Z_1^tZ_1^{t+1}$ (top right), $Z_2^{t-1}Z_2^t$ (bottom left) and $Z_1^tZ_2^tZ_2^{t+1}$ (bottom right). A measurement error on Z_1^t (pink hyperedge) therefore flips three checks, complicating the decoding problem. By decoding only the checks along the back-propagation of the logical operator, stabilizer measurement errors flip only two checks, enabling efficient decoding via matching. (c) Logical Pauli products can be decoded independently in parallel. (d) For some circuits, the corrections for a logical Pauli product can be committed after decoding (i-ii), reducing the size of future decoding problems (iii).

ing to an overall *reduction* in decoding complexity. Although the surface code allows for fast decoding via matching, we emphasize this approach applies to generic codes where this procedure directly enables promoting a memory decoder to an algorithmic decoder, by simply tracking how operators propagate through transversal gates.

3.10 Fault-tolerant compiling

Another key aspect we learned from our experiments is the importance of considering the structure of an algorithm and how it connects to the structure of the error-correction. Error-correction codes,

as described above, have a very different set of rules compared to physical qubits - for example, continuous rotations are hard while entanglement is easy. It is then prudent to consider creative ways to structure algorithms in ways that are most natively implemented with the native discrete digital gate set in a QEC code.

We first considered such a fault-tolerant compiling for realizing a classically complex quantum simulation with [[8,3,2]] codes, described in the next section. The idea here was to take the native discrete gate set of the code and use it for a quantum scrambling simulation, which works perfectly well with discrete gates.

These ideas of fault-tolerant compiling have since turned out to be very important in the context of high-rate quantum low-density parity-check (qLDPC) codes. The idea behind qLDPC codes is to take many logical qubits and stuff them into a single block, in order to reduce overheads on number of physical qubits required [28]. However, once all the logical qubits are overlapping with eachother, it becomes hard to operate on them! An elegant solution is to note that in many algorithms of interest, *all the logical qubits have to do the same thing anyway* and the logical qubits can be operated on in-parallel. Such a fault-tolerant compiling has led us to develop theoretical constructions with qLDPC codes that reduce overheads for *computation*, not just storage [154].

3.11 Reduction of space-time overheads

Quantum error correction has overheads of both space and time. Space overheads correspond to needing many physical qubits to encode a logical qubit, and time overheads correspond to needing to do many QEC cycles to do a logical operation. Very often, space and time can be traded off with each-other, and so usually the most interesting metric is space-time overhead, meaning how many total time steps on how many total physical qubits [59,15].

There are many examples of such space-time tradeoffs on both the QEC and algorithm level.

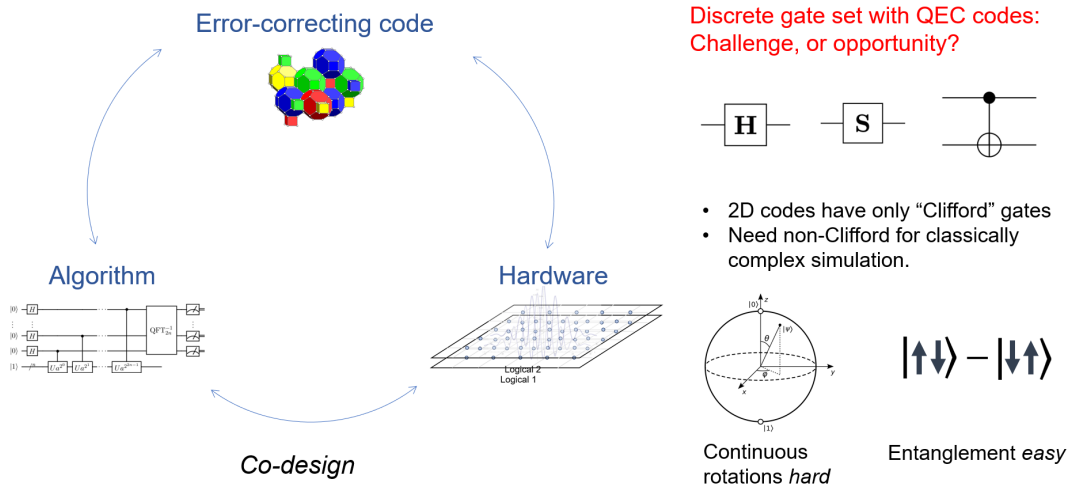


Figure 3.12: Fault-tolerant compiling. Error-correction codes have specific constraints due to only being able to directly implement a discrete digital gate set. Structuring the QEC in a way that directly accomplishes a particular algorithm, as well as being tailored to the hardware, can greatly reduce overheads.

The simplest is that you can often speed up an algorithm by parallelizing [57] across more qubits. On a QEC level, one can reduce the typical d rounds by making ancilla logical qubits fault-tolerantly ahead of time and then doing a transversal CNOT to realize fault-tolerant syndrome extraction in a single shot: this saves time on your data logical qubit but keeps the overall overhead the same. A related example is 3D vs 2D codes: 3D codes allow for single-shot error correction [22], requiring only 1 cycle per logic operation instead of d, but have d^3 space overhead instead of d^2. qLDPC codes, mentioned above, allow storing many logical qubits in a single block with parameters $[[n, k, d]]$, often leading to a saving of k in space overheads, but since they all need to be addressed serially it lends to a corresponding increase in time overhead by a factor of k [152]. Many such examples of space-time tradeoffs exist.

The experimental “discoveries” made below lend to genuine reductions in space-time overheads. The transversal algorithms and correlated decoding, explored in the context of entangling two $d = 7$ surface codes into a Bell state with a single round of stabilizer measurement, allows us to reduce the time overhead by a factor of d for universal computation while reducing the classical decoding com-

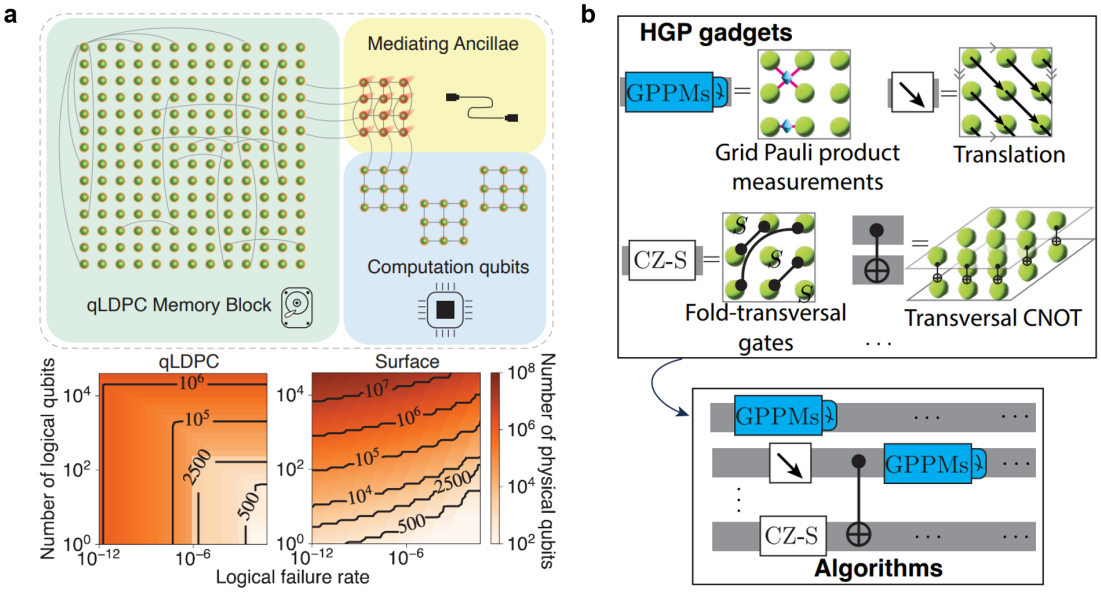


Figure 3.13: Computation with high-rate qLDPC codes. **a,** One method of computation with high-rate qLDPC codes. Many logical qubits are encoded in a memory block, and then telepored into surface codes via ancilla logical qubits for surface-code-based computaiton. This reduces the space overhead of the computation, but with a corresponding increase in time overhead due to the intrinsic serialization. See Ref. [152] for additional details. **b,** Methods for reduced space-time overhead for computation. By using fault-tolerant compiling techniques, like those realized with our [[8,3,2]] experiment, one can find algorithms where highly parallel operations on the logical qubit encoded in qLDPC blocks are well-suited to the algorithm. This results in a genuine reduction to space-time overhead of the computation. See Ref. [154] for additional details.

plexity. Although this sounds "to good to be true", it directly arises from the physics of transversal gates probed experimentally below. Our implementations of fault-tolerant compiling, developed in the context of efficient implementation of classically complex quantum scrambling simulations using [[8,3,2]] codes, has led us to develop qLDPC computations that have true reductions in space-time overheads[154] by leveraging the intrinsic parallelism in an algorithm to avoid serialization. We have also developed clean and simple ways of dealing with leakage like qubit loss in our transversal algorithms[7], leveraging the native role that teleportation plays in quantum error correction, which will be discussed below. We have also found fundamental requirements that cannot be reduced, for example the experiments below observe fundamental reasons why logical states with magic require a certain degree of entanglement within the block. There is clearly still more to learn.

3.12 Experimental studies of error-corrected processing

Having overviewed some of the key ideas, the following sections will describe experiments we performed in studying error-corrected processing with logical qubits. The first set of experiments had the goal of seeing what lessons we can learn from exploring error-corrected algorithms. The second set of experiments set to cement the core requirements and physical mechanism for a fault-tolerant computing architecture.

4

Experiments with logical qubit algorithms

Additional details for the description in this section can be found in Ref. [17].

Suppressing errors is the central challenge for useful quantum computing [113], requiring quantum error correction [132,136,44,119,114] for large-scale processing. However, the overhead in the realization of error-corrected "logical" qubits, where information is encoded across many physical qubits for redundancy [132,136,44], poses significant challenges to large-scale logical quantum computing. Here we report the realization of a programmable quantum processor based on encoded logical qubits

operating with up to 280 physical qubits. Utilizing logical-level control and a zoned architecture in reconfigurable neutral atom arrays [19], our system combines high two-qubit gate fidelities [52], arbitrary connectivity [14,19], as well as fully programmable single-qubit rotations and mid-circuit readout [43,135,62,101,98,107]. Operating this logical processor with various types of encodings, we demonstrate improvement of a two-qubit logic gate by scaling surface code [114] distance from $d = 3$ to $d = 7$, preparation of color code qubits with break-even fidelities [119], fault-tolerant creation of logical GHZ states and feedforward entanglement teleportation, as well as operation of 40 color code qubits. Finally, using three-dimensional [[8,3,2]] code blocks [34,146], we realize computationally complex sampling circuits [5] with up to 48 logical qubits entangled with hypercube connectivity [89] with 228 logical two-qubit gates and 48 logical CCZ gates [26]. We find that this logical encoding substantially improves algorithmic performance with error detection, outperforming physical qubit fidelities at both cross-entropy benchmarking and quantum simulations of fast scrambling [39,75]. These results herald the advent of early error-corrected quantum computation and chart a path toward large-scale logical processors.

Quantum computers have the potential to significantly outperform their classical counterparts for solving certain problems [113]. However, executing large-scale, useful algorithms on quantum processors requires very low gate error rates (generally below $\sim 10^{-10}$) [59], far below those that will likely ever be achievable with any physical device [132]. The landmark development of quantum error correction (QEC) theory provides a conceptual solution to this challenge [132,136,44]. The key idea is to use entanglement to delocalize a logical qubit degree of freedom across many redundant physical qubits, such that if any given physical qubit fails, it does not corrupt the underlying logical information. In principle, with sufficiently low physical error rates and sufficiently many qubits, a logical qubit can be made to operate with extremely high fidelity, providing a path to realizing large-scale algorithms [44]. However, in practice useful QEC poses many challenges, ranging from significant overhead in physical qubit numbers [59] to highly complex gate operations between the

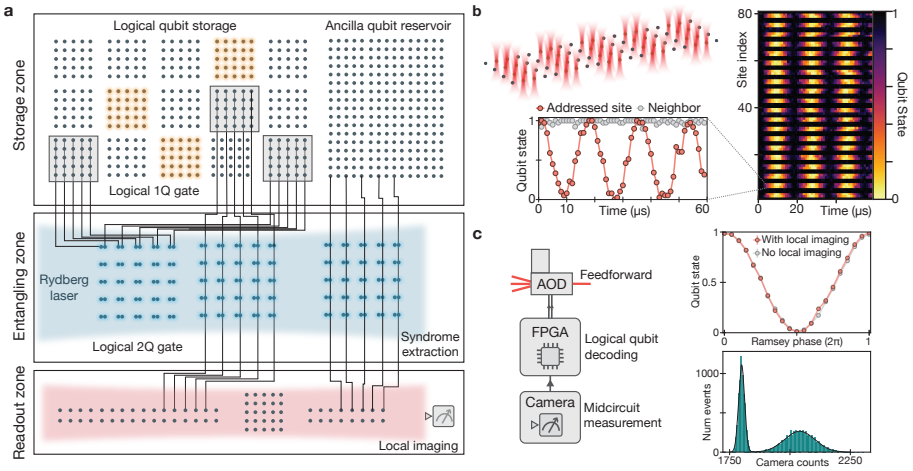


Figure 4.1: A programmable logical processor based on reconfigurable atom arrays. a, Schematic of the logical processor, segmented into three zones: storage, entangling, and readout. Logical single-qubit and two-qubit operations are realized transversally with efficient, parallel operations. Transversal CNOTs are realized by interlacing two logical qubit grids and performing a single global entangling pulse that excites atoms to Rydberg states. Physical qubits are encoded in hyperfine ground states of ^{87}Rb atoms trapped in optical tweezers. **b,** Fully programmable single-qubit rotations are implemented using Raman excitation through a 2D AOD; parallel grid illumination delivers the same instruction to multiple atomic qubits. **c,** Mid-circuit readout and feedforward. The imaging histogram shows high-fidelity state discrimination (500 μs imaging time, readout fidelity $\approx$ 99.8%, Methods), and the Ramsey fringe shows that qubit coherence is unaffected by measuring other qubits in the readout zone (error probability $p \sim 10^{-3}$, Methods). The FPGA performs real-time image processing, state decoding, and feedforward (Fig. 4).

delocalized logical degrees of freedom [57]. Recent experiments have achieved milestone demonstrations of two logical qubits and one entangling gate [119,114], and explorations of novel encodings [126,Honciuc Menendez et al.,Wang et al.,4].

One specific challenge for realizing large-scale logical processors involves efficient control. Unlike modern classical processors that can efficiently access and manipulate many bits of information [111], quantum devices are typically built such that each physical qubit requires multiple classical control lines. While suitable to the implementation of physical qubit processors, this approach poses a significant obstacle to the control of logical qubits redundantly encoded over many physical qubits.

Here, we describe the realization of a programmable quantum processor based on hardware-efficient control over logical qubits in reconfigurable neutral atom arrays [19]. We use this logical pro-

cessor to demonstrate key building blocks of QEC and realize programmable logical algorithms. In particular, we explore important features of logical operations and circuits, including scaling to large codes, fault-tolerance, and complex non-Clifford circuits.

4.1 Logical processor based on atom arrays

Our logical processor architecture, illustrated in Fig. 1a, is segmented into three zones. The storage zone is used for dense qubit storage, free from entangling gate errors and featuring long coherence times. The entangling zone is used for parallel logical qubit encoding, stabilizer measurements, and logical gate operations. Finally, the readout zone enables mid-circuit readout of desired logical or physical qubits, without disturbing the coherence of the computation qubits still in operation. This architecture is implemented using arrays of individual ^{87}Rb atoms trapped in optical tweezers, which can be dynamically reconfigured in the middle of the computation while preserving qubit coherence [14,19].

Our experiments make use of the apparatus described previously in Refs. [49,19,52], with key upgrades enabling universal digital operation. Physical qubits are encoded in clock states within the ground-state hyperfine manifold ($T_2 > 1s$ [19]), and stored in optical tweezer arrays created by a spatial light modulator (SLM) [49,122]. We utilize systems of up to 280 atomic qubits, combining high-fidelity two-qubit gates [52], enabled by fast excitation into atomic Rydberg states interacting through robust Rydberg blockade [79], with arbitrary connectivity enabled by atom transport via 2D acousto-optic deflectors (AODs) [19]. Central to our approach of scalable control, AODs [122,43,135,62,101,98,107,123] use frequency multiplexing to take in just two voltage waveforms (one for each axis) to create large, dynamically programmable grids of light. Fully programmable local single-qubit rotations are realized via qubit-specific, parallel Raman excitation through an additional 2D AOD (Fig. 1b) [63]. Mid-circuit readout is enabled by moving selected qubits $\sim$ 100 μm away to a readout zone and

illuminating with a focused imaging beam [37,19], resulting in high-fidelity imaging as well as negligible decoherence on stored qubits (Fig. 1c). The mid-circuit [43,135,62,101,98,107] image is collected with a CMOS camera and sent to an FPGA for real-time decoding and feedforward.

The central aspect of our logical processor is the control of individual logical qubits as the fundamental units, instead of individual physical qubits. To this end, we observe that during a vast majority of error-corrected operations, the physical qubits of a logical block are supposed to realize the same operation, and this instruction can be delivered in parallel with only a few control lines. This approach naturally multiplexes with optical techniques. For example, to realize a logical single-qubit gate [132], we use the Raman 2D AOD (Fig. 1b) to create a grid of light beams and simultaneously illuminate the physical qubits of the logical block with the same instruction. Such a gate is transversal [132], meaning that operations act on physical qubits of the code block independently. This transversal property further implies the gate is inherently fault-tolerant [132], meaning that errors cannot spread within the code block (see Methods), thereby preventing a physical error from spreading into a logical fault. Crucially, a similar approach can realize logical entangling gates [132,44]. Specifically, we use the grids generated by our moving 2D AOD to pick up two logical qubits, interlace them in the entangling zone, and then pulse our single global Rydberg excitation laser to realize a physical entangling gate on each twin pair of the blocks (Figs. 1a,2a). This process realizes a high-fidelity, fault-tolerant transversal CNOT in a single parallel step.

4.2 Improving entangling gates with code distance

A key property of QEC codes is that, for error rates below some threshold, the performance should improve with system size, associated with a so-called code distance [44,57]. Recently, this property has been experimentally verified by reducing idling errors of a code [114]. Neutral atom qubits can be idly stored for long times with low errors, and the central challenge is to improve entangling operations

with code distance. Thus motivated, we realize a transversal CNOT gate using logical qubits encoded in two surface codes (Fig. 2). Surface codes have stabilizers that are used for detecting and correcting errors without disrupting the logical state[44,57]. The stabilizers form a 2D lattice of 4-body plaquettes of X and Z operators, which commute with the X_L (Z_L) logical operators that run horizontally (vertically) along the lattice (Fig. 2d). By measuring stabilizers one can detect the presence of physical qubit errors, decode (infer what error occured), and correct the error simply by applying a software Z_L/X_L correction[57]. Such a code can detect and correct a certain number of errors determined by the linear dimension of the system (the code distance d).

To test the performance of our logical entangling gate, we first initialize the logical qubits by preparing physical qubits of two blocks in $|+\rangle$ and $|0\rangle$ states, respectively, and performing a single round of stabilizer measurements with parallel operations[19]. While this state preparation is non-fault-tolerant (nFT) beyond $d = 3$, we are still able to probe error suppression of the transversal CNOT (Methods). Specifically, we prepare the two logicals in state $|+_L\rangle$ and $|0_L\rangle$, perform the transversal CNOT, and then projectively measure to evaluate the logical Bell state stabilizers $X_L^1 X_L^2$ and $Z_L^1 Z_L^2$ (Fig. 2c). For decoding and correcting the logical state, we observe there are strong correlations between the stabilizers of the two blocks due to propagation of physical errors between the codes during the transversal CNOT (Fig. 2b)[15]. We utilize these correlations to improve performance by decoding the logical qubits jointly, realized by a joint decoding graph that includes edges and hyperedges connecting the stabilizers of the two logical qubits (Fig. 2b). Using this correlated decoding procedure, we measure ≈ 0.95 populations in the $X_L X_L$ and $Z_L Z_L$ bases (Fig. 2c), showing entanglement between the $d = 7$ logical qubits.

Studying the performance as a function of code size (Fig. 2d) reveals that the logical Bell pair improves with larger code distance, demonstrating improvement of the entangling operation. In contrast, we note that when conventional decoding, i.e. independent minimum-weight perfect matching within both codes[44], is used, the fidelity decreases with code distance. This is in part due to the

nFT state preparation, whose effect is partially mitigated by the correlated decoding (Methods).

We emphasize that while these results demonstrate surpassing an effective threshold for the entire circuit (implying we surpass the threshold of the transversal CNOT), such a threshold is higher due to projective readout after the transversal CNOT. In practice, the transversal CNOT should be used in combination with many repeated rounds of noisy syndrome extraction [114], which is expected to have a somewhat lower threshold and is an important goal for future research.

4.3 Fault-tolerant logical algorithms

All logical algorithms we perform in this work are built from transversal gates which are intrinsically fault-tolerant [132]. We now also use fault-tolerant state preparation to explore programmable logical algorithms. We use two-dimensional $d = 3$ color codes [136,21], which are topological codes akin to the surface code, but with the useful capability of transversal operations of the full Clifford group: Hadamard (H), $\pi/2$ phase (S) gate, and CNOT [21]. This transversal gate set can realize any Clifford circuit fault-tolerantly. As a test case, here we create a logical GHZ state. Figure 3a shows the implementation of a 10-logical-qubit algorithm, in which all ten qubits are first encoded by a nFT encoding circuit (Methods). Then, five of the codes are used as ancilla logicals, performing parallel transversal CNOTs in order to fault-tolerantly detect errors on the computation logicals [60], and are then moved into the storage zone where they are safely kept. Subsequently four computation logicals are used to prepare the GHZ state, and logical Clifford rotations are used at the end of the circuit for direct fidelity estimation [56] and full logical state tomography.

We first benchmark our state initialization (Fig. 3b) [50,112,119]. Averaged over the five computation logicals, we find that by using the fault-tolerant initialization (postselecting on the ancilla logical flag not detecting errors) our $|0_L\rangle$ initialization fidelity is $99.91^{+0.04}_{-0.09}\%$, exceeding both our physical qubit $|0\rangle$ initialization fidelity ($99.32(4)\%$ [52]) and physical two-qubit gate fidelity (99.5% [52]). Then,

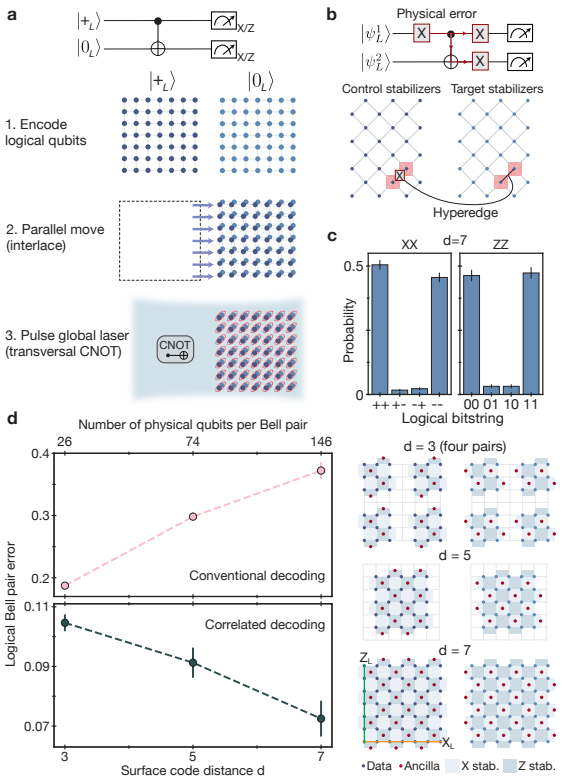


Figure 4.2: Transversal entangling gates between two surface codes. **a,** Illustration of transversal CNOT between two $d = 7$ surface codes based on parallel atom transport. **b,** The concept of correlated decoding. Physical errors propagate between physical qubit pairs during transversal CNOT gates, creating correlations that can be utilized for improved decoding. We account for these correlations, arising from deterministic error propagation (as opposed to correlated error events), by adding edges and hyperedges that connect the decoding graphs of the two logical qubits. **c,** Populations of entangled $d = 7$ surface codes measured in the XX and ZZ basis. **d,** Measured Bell pair error as a function of code distance, for both conventional (top) and correlated (bottom) decoding. We estimate Bell error with the average of the ZZ populations and the XX parities (Methods). To reduce code distance we simply remove selected atoms the grid, as shown on the right, ensuring unchanged experimental conditions (for $d = 3$, four logical Bell pairs are generated in parallel). Error bars represent standard error of the mean.

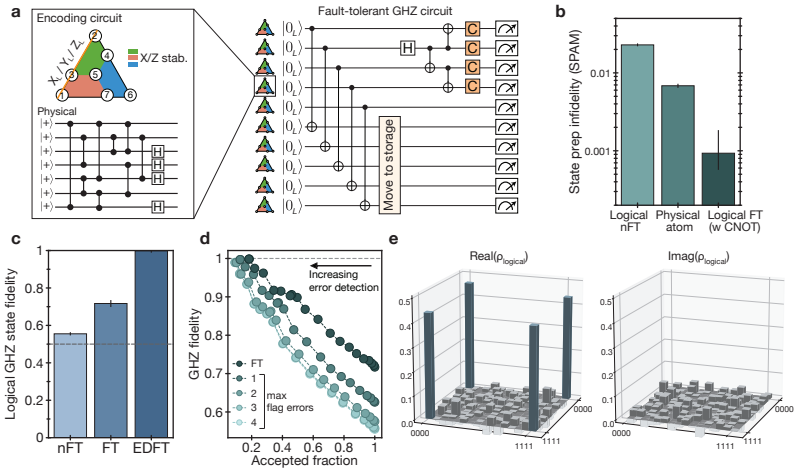


Figure 4.3: Fault-tolerant logical algorithms a, Circuit for preparation of logical GHZ state. Ten color codes are encoded non-fault-tolerantly (nFT), and then parallel transversal CNOTs between computation and ancilla logical qubits perform FT initialization. The ancilla logical qubits are moved to storage, and a 4-logical-qubit GHZ state is created between the computation qubits. Logical Clifford operations are applied before readout to probe the GHZ state. **b,** State-preparation-and-measurement (SPAM) infidelity of the logical qubits without (nFT) and with (FT) the transversal-CNOT-based flagged preparation, compared to physical qubit SPAM. **c,** Logical GHZ fidelity without postselecting on flags (nFT), postselecting on flags (FT), and postselecting on flags and stabilizers of the computation logical qubits, corresponding to error detection (EDFT). **d,** GHZ fidelity as a function of sliding-scale error detection threshold (converted into the probability of accepted repetitions) and of number of successful flags in the circuit. **e,** Density matrix of the 4-logical-qubit GHZ state (with at most 3 flag errors) measured via full state tomography involving all 256 logical Pauli strings.

Fig. 3c shows the resulting GHZ state fidelity obtained using the fault-tolerant algorithm is 72(2)% (again using correlated decoding), demonstrating genuine multipartite entanglement. Additionally, one can postselect on all stabilizers of our computation logicals being correct; using this error detection approach, the GHZ fidelity increases to $99.85^{+0.1}_{-1.0}$% at the cost of postselection overhead.

Since not all nontrivial syndromes are equally likely to cause algorithmic failure, one can perform a partial postselection where syndrome events most likely to have caused algorithmic failure are discarded, given by the weight of the correlated matching in the whole algorithm. Figure 3d shows the measured GHZ fidelity as a function of this sliding threshold converted into a fraction of accepted experimental repetitions, continuously tuning the tradeoff between the success probability of the algorithm and its fidelity; e.g., discarding just 50% of the data improves GHZ fidelity to $\approx$ 90%. (As discussed below, for certain applications purifying samples can be advantageous in improving algorithmic performance.) Finally, fault-tolerantly measuring all 256 logical Pauli strings, we perform full GHZ state tomography (Fig. 3e).

The use of the zoned architecture directly allows scaling circuits to larger numbers, without increasing the number of controls, by encoding and operating on logical qubits, moving them to storage, and then accessing storage as appropriate. This process is illustrated in Figs. 4a,b, where ten color codes are made and operated on with parallel transversal CNOTs, moved to storage, and then more qubits are accessed from storage. Repeating this process four times, we create 40 color codes with 280 physical qubits, at the cost of slow idling errors of $\sim 1\%$ logical decoherence per additional encoding step (Fig. 4c). These storage idling errors primarily originate from global Raman π pulses applied for dynamical decoupling of atoms in the entangling zone, which could be significantly reduced with zone-specific Raman controls.

Since mid-circuit readout[43,135,62,101,98,107] is an important component of logical algorithms, we next demonstrate a fault-tolerant entanglement teleportation circuit. We first create a three-logical-qubit GHZ state $|0_L 0_L 0_L\rangle + |1_L 1_L 1_L\rangle$ (Figs. 4d,e) from fault-tolerantly prepared color codes. Mid-

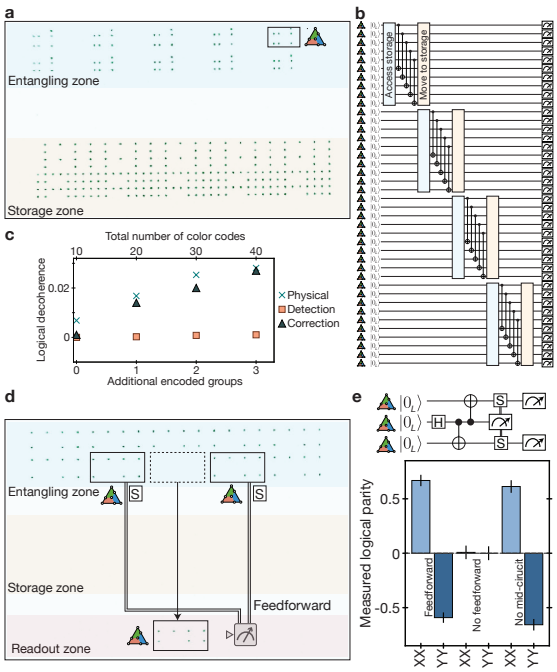


Figure 4.4: Zoned logical processor: scaling and mid-circuit feedforward. a, Atoms in storage and entangling zones, and approach for creating and entangling 40 color codes with 280 physical qubits. **b,c,** 40 color codes are prepared with an nFT circuit, and then 20 transversal CNOTs are used to fault-tolerantly prepare 20 of the 40 codes, whose fidelity is plotted. Logical decoherence is smaller than the physical idling decoherence experienced during the encoding steps. **d,** Mid-circuit measurement and feedforward for logical entanglement teleportation. The middle of three logical qubits is measured in the X-basis, and by applying a mid-circuit conditional, locally pulsed logical S rotation on the other two logical qubits, the state $|0_L 0_L\rangle \quad + \quad |1_L 1_L\rangle$ is prepared. **e,** Measured logical qubit parity with and without feedforward, showing that feedforward recovers the intended state with Bell fidelity of 77(2)% (ZZ parities of 83(4)% not plotted, Methods). No mid-circuit refers to turning off the mid-circuit readout, and postselecting on the middle logical being in state $|+_L\rangle$ in the final readout. By postselecting on perfect stabilizers of only the two computation logicals (error detection in the final measurement) the feedforward Bell fidelity is 92(2)% (not plotted). In **d** three of the additional blocks are flag qubits and the other four are prepared but unused for this circuit.

circuit X-basis measurement of the middle logical creates $|0_L0_L\rangle + |1_L1_L\rangle$ if measured as $|+_L\rangle$, and $|0_L0_L\rangle - |1_L1_L\rangle$ if measured as $|-_L\rangle$. One recovers $|0_L0_L\rangle + |1_L1_L\rangle$ by applying a logical S gate to the first and third logicals conditioned real-time on the state of the middle logical, akin to the magic state teleportation circuit [57]. Measurements in Fig. 4e indicate that while $\langle X_L X_L\rangle$ and $\langle Y_L Y_L\rangle$ indeed vanish without the feedforward step, applying the feedforward correction we recover a Bell state fidelity of 77(2)%, limited by imperfections in the original underlying GHZ state. By repeating this experiment without mid-circuit readout and instead post-selecting on the middle logical being in $|+_L\rangle$, we find a similar Bell fidelity of 75(2)%, indicating high-fidelity performance of the readout and feedforward operations.

4.4 Complex logical circuits using 3D codes

One important challenge in realizing complex algorithms with logical qubits is that universal computation cannot be implemented transversally [47]. For instance, when using 2D codes such as the surface code, non-Clifford operations cannot be easily performed [21], and relatively expensive techniques are required for non-trivial computation [57,29] as Clifford circuits can be easily simulated [1]. In contrast, 3D codes can transversally realize non-Clifford operations, but lose the transversal H [21]. However, these constraints do not imply that classically hard or useful quantum circuits cannot be realized transversally or efficiently. Motivated by these considerations, we explore efficient realization of classically hard algorithms that are co-designed with a particular error-correcting code. Specifically, we implement fast scrambling circuits using small 3D codes, which are used for native non-Clifford operations (CCZ).

We focus on small 3-dimensional [[8,3,2]] codes (Fig. 5a) [34,146,Honciuc Menendez et al.,Wang et al.], which have various appealing features. They encode three logicals per block, feature $d = 2$ ($d = 4$) in the Z basis (X basis), implying error detection (correction) capabilites for Z (X) errors, and can realize a

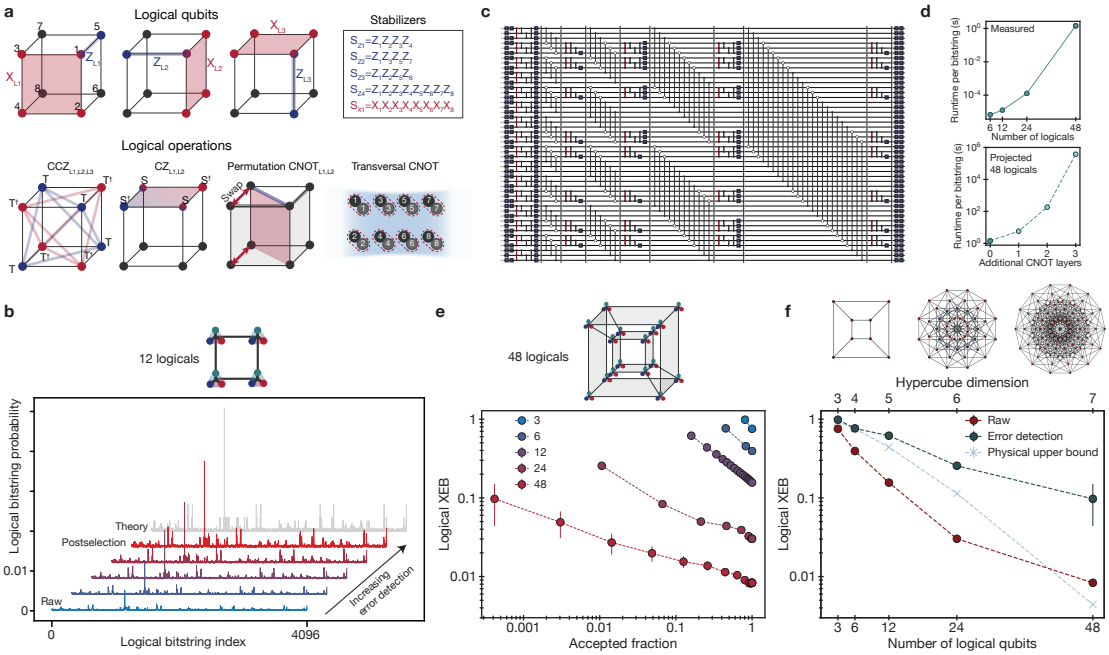


Figure 4.5: Complex logical circuits using 3D codes. a, [[8,3,2]] block codes can transversally realize $\{\text{CCZ}, \text{CZ}, Z, \text{CNOT}\}$ gates within each block, and transversal CNOTs between blocks. By preparing logical qubits in $|+_L\rangle$, performing layers of $\{\text{CCZ}, \text{CZ}, Z\}$ alternated with interblock CNOTs, and measuring in the X-basis, we realize classically hard sampling circuits with logical qubits. **b,** Measured sampling outcomes for a circuit with 12 logical qubits, 8 logical CZs, 12 logical CNOTs, and 8 logical CCZs. By increasing error detection, the measured distribution converges toward the ideal distribution. **c,** Circuit involving 48 logical qubits with 228 logical CZ / CNOT gates and 48 logical CCZs. **d,** Classical simulation runtime for calculating an individual bitstring probability; bottom plot is estimated based on matrix multiplication complexity. **e,** Measured normalized XEB as a function of sliding-scale error-detection for 3, 6, 12, 24, and 48 logical qubits. For all sizes we observe a finite XEB score which improves with increased error detection. Diagram shows 48-logical connectivity, with logical triplets entangled on a 4D hypercube. **f,** Scaling of raw (red) and fully error-detected (black) XEB from **e**. Physical upper-bound fidelity (blue) is calculated using best measured physical gate fidelities. Diagrams show physical connectivity. [[8,3,2]] cubes are entangled on 4D hypercubes, realizing physical connectivity of 7D hypercubes.

transversal CNOT between blocks. Most importantly, by using physical $\{T, S\}$ rotations (T is $\pi/4$ phase gate) one can realize transversal $\{\mathrm{CCZ}, \mathrm{CZ}, Z\}$ gates on the logical qubits encoded within each block, as well as intrablock CNOTs by physical permutation [Honciuc Menendez et al.,Wang et al.] (Methods). This gate set allows us to transversally realize the circuits illustrated in Fig. 5a,c, alternating between layers of $\{\mathrm{CCZ}, \mathrm{CZ}, Z\}$ within blocks and layers of CNOTs between blocks. Although transversal H is forbidden, initialization and measurement in either the X or Z basis effectively allows H at the beginning and end of the circuit.

We use these transversal operations to realize logical algorithms that are difficult to simulate classically [104,110]. More specifically, these circuits can be mapped to Instantaneous Quantum Polynomial (IQP) circuits [26,104,110]. Sampling from the output distribution of such circuits is known to be classically hard in certain instances [26], implying a quantum device can be exponentially faster than a classical computer for this task.

Figure 5b shows an example implementation of a 12-logical-qubit sampling circuit. Here, we prepare all logical blocks in $|+\rangle_L$, implement a scrambling circuit with 28 logical entangling gates, and then measure all logicals in the X-basis. Figure 5b shows the probability of observing each of the $2^{12} = 4096$ possible logical bitstring outcomes, showing that as we progressively apply more error detection (i.e. postselection) in postprocessing, the distribution more closely reproduces the ideal theoretical distribution. To characterize the distribution overlap, we use the cross-entropy benchmark (XEB) [5], which is a weighted sum between the measured probability distribution and the ideal calculated distribution, normalized such that XEB=1 corresponds to perfectly reproducing the ideal distribution, and XEB=0 corresponds to the uniform distribution which occurs when circuits are overwhelmed by noise. Consistent with Fig. 5b, the 12-logical-qubit circuit XEB increases from 0.156(2) to 0.616(7) upon applying error detection (Fig. 5e). We note that XEB should be a good fidelity benchmark for IQP circuits (Methods).

We next explore scaling to larger systems and circuit depths. To ensure high complexity of our

logical circuits, we use nonlocal connections to entangle the logical triplets on up to 4D hypercube graphs, which results in fast scrambling[89]. Exploring entangled systems of 3, 6, 12, 24, and 48 logical qubits, in all cases we find a finite XEB score which improves with increased error detection (Figs. 5e,f). The finite XEB indicates successful sampling, and the improvement with error detection shows the benefit of using logical qubits. While this improvement comes at the cost of measurement time due to error detection, improving the sample quality cannot be replaced by simply generating more samples. Thus, improving the XEB score yields significant practical gains. We obtain an XEB of $\approx$ 0.1 for 48 logical qubits and hundreds of nonlocal logical entangling gates, up to roughly an order of magnitude higher than previous physical qubit implementations of similar complexity[5,150], showing the benefits of a logical encoding for this application.

Assuming our best measured physical fidelities, the estimated upper-bound for an optimized physical qubit implementation in our system is also significantly below the measured logical XEB (blue line in Fig. 5f, Methods). In attempting to run these complex physical circuits, in practice we find that realising non-vanishing XEB is significantly more challenging; we confirm with small physical instances that we measure values well below this upper-bound (Methods). In addition to the error-detecting benefits, it appears the logical circuit is significantly more tolerant to coherent errors, exhibiting operation that is inherently digital, just with imperfect fidelity, consistent with theoretical predictions[25]. We also note that for the logical algorithms we optimize performance by optimizing the stabilizer expectation values (rather than the complex sampling output), providing further advantage for logical implementations.

Our 48-logical circuit, corresponding to a physical qubit connectivity of a 7D hypercube, contains up to 228 logical two-qubit gates and 48 logical CCZ gates. Simulation of such logical circuits is challenging due to the high connectivity (rendering tensor networks inefficient) and large numbers of non-Cliffords[23]. To benchmark our circuits, we structure them such that we can leverage an efficient simulation method (Methods) which takes $\approx$ 2 seconds to calculate the probability of

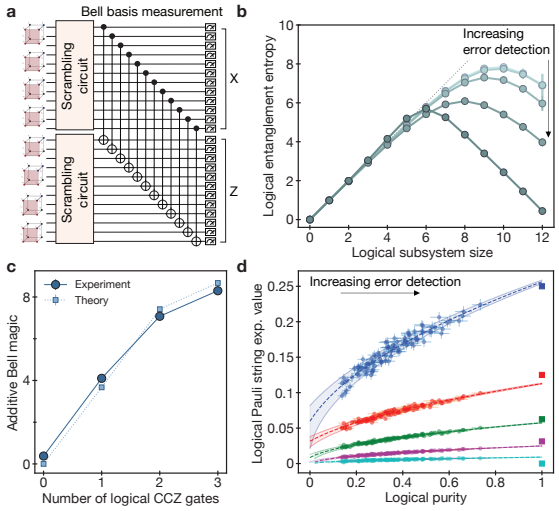


Figure 4.6: Logical two-copy measurement. **a,** Identical scrambling circuits are performed on two copies of 12 logical qubits, and then measured in the Bell basis to extract information about the state. Z basis measurements are corrected with an [[8,3,2]] decoder (when full error detection is not applied). **b,** Measured entanglement entropy as a function of subsystem size, showing expected Page-curve behavior[89] for the highly scrambled state, improving with increased error detection. **c,** Measured and simulated magic (associated with non-Clifford operations) as a function of number of CCZ gates applied, performed on two copies of scrambled 6-logical-qubit systems. **d,** Pauli string measurement and zero-noise extrapolation using logical qubits. Plot shows the absolute values of all 4^{12} Pauli string expectation values, which only have five discrete values for our digital circuit; Pauli strings with the same theory value are grouped. By analyzing with sliding-scale error detection we improve toward the theoretical expectation values (squares) while also improving toward a purity of 1. By extrapolating to perfect purity, we extrapolate the expectation values and better-approximate the ideal values (shaded regions are statistical fit uncertainty).

each bitstring (Fig. 5d). Modeling noise in our logical circuits is even more complicated, as they are composed from 128 physical qubits and 384 T gates, thereby making experimentation with logical algorithms necessary to understand and optimize performance.

4.5 Quantum simulations with logical qubits

Finally, we explore the use of logical qubits as a tool in quantum simulation, probing entanglement properties of our fast scrambling circuits, potentially related to complex systems such as black holes[125,89]. In particular, we utilize a Bell-basis measurement made on two copies of the quantum state (Fig. 6a), which is a powerful tool that can efficiently extract many properties of an unknown

state [39,75,66]. With this two-copy technique, in Fig. 6b we plot the measured entanglement entropy in the scrambled system. We observe a characteristic Page curve [125] associated with a maximally entangled, highly scrambled, but globally pure state. These measurements also reveal a final state purity of 0.74(3), compared to the measured XEB of 0.616(7) in Fig. 5f, consistent with XEB being a good proxy for the final state fidelity. Despite postselection overhead, we find error detection significantly improves signal-to-noise here, as near-zero entropies are exponentially faster to measure.

Two-copy measurements can also be used to simultaneously extract information about all 4^N Pauli strings [75]. Using this property and an analysis technique known as Bell difference sampling [68] we experimentally evaluate and directly verify the amount of additive Bell magic [68] in our circuits as a function of number of applied logical CCZs (Fig. 6c). This measurement of magic, associated with non-Clifford operations, quantifies the number of T gates (assuming decomposition into T) required to realize the quantum state by observing the probability that sampled Pauli strings commute with each other (see Methods). Moreover, combining encoded qubits and two-copy measurement allows for additional error mitigation techniques. As an example, Fig. 6d shows the measured absolute expectation values of all 4^{12} logical Pauli strings with sliding-scale error detection. Since in the two-copy measurements, for each error detection threshold we also measure the overall system purity, we can extrapolate our expectation values to the case of unit-purity (zero-noise) [84]. This procedure evaluates the averaged Pauli expectation values to $\approx$ 10% relative precision of the ideal theoretical values spanning several orders of magnitude (Methods).

4.6 Outlook

These experiments demonstrate key ingredients of scalable error correction and quantum information processing with logical qubits. In addition to implementing the key elements of logical processing, our approach demonstrates practical utility of encoding methods for improving sampling

and quantum simulations of complex scrambling circuits. Future work can explore if these methods can be generalized e.g. to more robust, higher-distance codes and if such highly entangled, non-Clifford states could be utilized in practical algorithms. We note the demonstrated logical circuits are approaching the edge of exact simulation methods (Fig. 5d), and can readily be used for exploring error-corrected quantum advantage. These examples demonstrate that the use of novel encoding schemes, co-designed with efficient implementations, can allow one to implement particular logical algorithms at reduced cost.

Our observations open the door for exploration of large-scale logical qubit devices. A key future milestone would be to perform repetitive error correction [114] during a logical quantum algorithm to greatly extend its accessible depth. This repetitive correction can be directly realized using the tools demonstrated here by repeating the stabilizer measurement (Fig. 2) in combination with mid-circuit readout (Fig. 4). The use of the zoned architecture and logical-level control should allow our techniques to be readily scaled to over 10000 physical qubits by increasing laser power and optimizing control methods, while QEC efficiency can be improved by reducing two-qubit gate errors to 0.1% [52]. Deep computation will further require continuous reloading of atoms from a reservoir source [135,107]. Continued scaling will benefit from improving encoding efficiency, e.g. by using quantum low-density-parity-check (qLDPC) codes [24,Xu et al.], utilizing erasure conversion [151,123,101] or noise bias [37], and optimizing the choice of (possibly multiple) atomic species [135,98], as well as advanced optical controls [63]. Further advances could be enabled by connecting processors together in a modular fashion using photonic links or transport [46,43], or more power-efficient trapping schemes such as optical lattices [140]. Although we do not expect clock speed to limit medium-scale logical systems, approaches to speed up processing in hardware [155] or with nonlocal connectivity [99] should also be explored. We expect that such experiments with early-generation logical devices will enable experimental and theoretical advances that dramatically reduce anticipated costs of large-scale error-corrected systems, accelerating development of practical applications of quantum computers.

5

Physical mechanisms for a fault-tolerant quantum architecture

Additional details for the description in this section can be found in Ref.[18].

Quantum mechanics exhibits features, such as superpositions and a continuous state space, that make it powerful for computation[11,133]. However, these same features make quantum computers extremely fragile[92]. Classical computers, in contrast, are composed of intrinsically robust digital bits

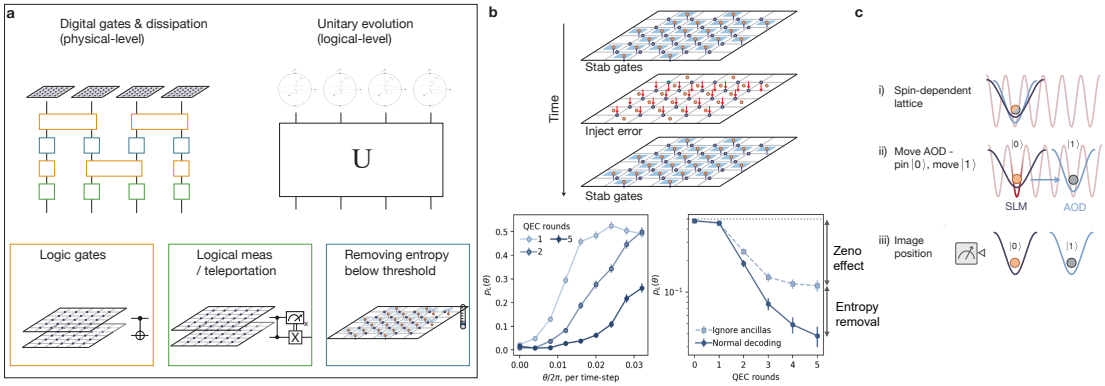


Figure 5.1: Exploring the building block for scalable fault-tolerant quantum computation. **a,** To implement error-corrected computation, the logical information is encoded physical qubits and digital gates are used to apply unitary evolution on the logical level. Multiple key building blocks ranging from logic gates, logical teleportation and below-threshold entropy removal are required for scalable implementations. **b,** Removing errors with repeated error-correction. Stabilizer measurement on d=5 surface code is interspersed with global coherent errors on the data qubits. Repeated correction reduces error build-up through both the Zeno effect and error tracking. **c,** Spin-to-position conversion for non-destructive qubit readout. Readout is accomplished with a state-selective 1D lattice that splits the spin-state.

that are robust to imperfections, due to the use of dissipation[91]. At first glance, quantum computation cannot allow this, since dissipation would destroy the coherent, unitary evolution. Moreover, because the quantum wavefunction evolves in a continuous state space, it is highly sensitive to imperfections — large-scale algorithms, for example, require qubit rotations to have a precision of 10^{-14} [59].

Remarkably, quantum error correction enables unitary evolution at the logical-qubit level while applying dissipation to the physical degrees of freedom to suppress error accumulation. While such stabilizing mechanisms arise naturally in classical systems, they must be carefully engineered in quantum systems. Achieving robust, digital quantum operations involves a range of components[3,57], and various methods exist for realizing them. In this work, we identify and implement the core physical mechanisms required for digital, fault-tolerant quantum computation, along with scalable methods for their use in algorithms. Through this, we uncover simplified approaches for achieving robust operation and effective error correction. In addition to observing long-theorized

behaviors central to quantum information science, these results provide a deeper understanding of these foundational mechanisms and can help guide the design of future algorithms, QEC protocols, and hardware architectures.

Figure 1a depicts the core ingredients required for scalable, fault-tolerant implementation of unitary logical evolution. Scalable fault-tolerant implementations require multiple key features, ultimately resulting in a logarithmic overhead for running arbitrary computations[3]. Logic gates are used for entangling logical qubits. However, this computation generates entropy due to imperfections[91], which must then be removed. Scalable operation involves removing newly created errors, such that the entropy remains sufficiently low and exponential error suppression can be achieved by increasing code distance[44]. Finally, universal processing requires measurement and teleportation to achieve universal gate sets[85,41] and introduce non-unitary evolution on the physical data qubits.

We study these using a logical processor described previously[17], with key upgrades enabling deep, universal error-corrected circuits. Physical qubits are encoded in hyperfine clock states within ^{87}Rb atoms, and high-fidelity entangling operations are realized via fast excitation to Rydberg states[53]. Hundreds of atoms are stored in optical tweezers generated by a spatial light modulator (SLM) in storage, entangling, and readout zones. Quantum circuits are programmed by shuttling qubits in the middle of the computation with a 2D acousto-optic deflector. Crucially this system provides massive control parallelism over logical qubit blocks where for logic gates, repeated error correction, and even universal computation, all the physical qubits within the logical block realize identical operations that can be delivered as parallel instructions with optical control.

Key upgrades enable delayed erasure detection of atoms for improved QEC performance[151,7], as well as techniques for re-initialization and re-use mid-circuit, further enabling an improvement in experimentaly cycle rate by two orders of magnitude. Fig 1c shows a diagram of non-destructive readout with a one-dimensional circularly polarized optical lattice in the readout zone. The optical lattice is spin-selective[149], pinning one qubit state in its starting tweezer location, and a second

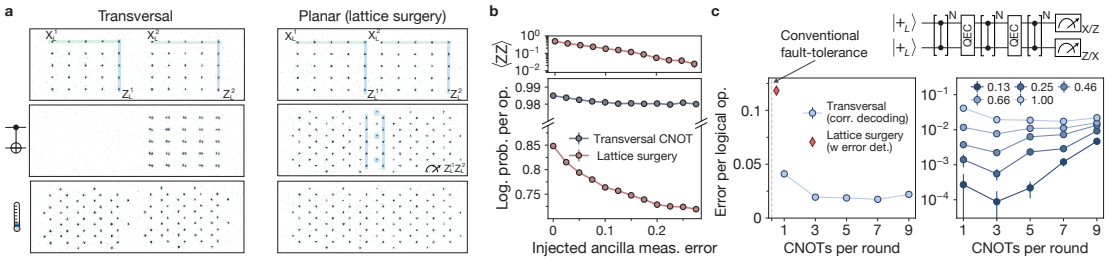


Figure 5.2: Role of stabilizer measurement in logic operations. a, Atom images illustrating two-qubit logic gates and stabilizer measurement. Lattice surgery is realized via measuring logical product and transversal gates are realized via atomic motion. **b,** Dependence of error of the logic operation on the ancilla measurement error. For transversal operations, ancilla measurement errors (injected in post-processing) weakly reduce the error per logical operation, while for lattice surgery the fidelity decays sharply. **c,** N repeated logic operations are interspersed with rounds of stabilizer measurement (QEC). Transversal CNOT has lower error than lattice surgery (left), and has an optimum of roughly 3 CNOTs per round, as seen most clearly with modest postselection (right).

tweezer drags away the unpinned atom into a different position. Conventional camera-based readout of the atomic position then realizes the readout, and subsequently the atoms can be returned to their initial location.

5.1 Stabilizer measurement during algorithms

We first study the role of stabilizer measurement in stabilizing a logical quantum bit. In Fig. 1, we create a distance-5 surface code logical qubit and perform five rounds of stabilizer measurement. Data qubits are held in static potentials in the entangling zone, and an ancilla block is moved over to perform parallel entangling operations. After extracting the stabilizer checks, the ancilla is moved to the storage zone, and a new ancilla block is brought in to repeat the cycle.

To evaluate how repeated stabilizer measurements remove error, we inject global coherent errors at regular intervals. Repeated correction during these intervals significantly suppresses logical error, as coherent errors are projected into incoherent ones (quantum Zeno effect) and tracked. Notably, although the injected errors are global, stabilizer measurements result in uncorrelated physical errors (Fig. 1b)[25]. Furthermore, these coherent physical errors do not translate to the logical level (Fig. 1c),

with logical rotations further suppressed by the randomness of stabilizer projections (Methods).

We next study stabilizer measurements during logic operations using lattice surgery[72] and transversal gates[132,44]. In the transversal case, we apply three rounds of N repeated CNOTs between two surface codes, interleaved with two rounds of stabilizer checks. In the lattice surgery case, we perform a ZZ measurement using two rounds of stabilizer checks[72]. To explore their differing roles, we introduce ancilla measurement errors in postprocessing (Fig. 2b). In the transversal setting, gate performance degrades mildly with increasing error, while in lattice surgery, logical entanglement rapidly degrades. This highlights a key distinction between the two cases. In the transversal gate setting, the logic is realized direcly between the data qubit pairs, and stabilizers play the role of entropy removal, whereas in lattice surgery, the ancilla measurements directly perform the logic operation and have to be made accurately.

Although conventionally d rounds of stabilizer measurements are assumed per gate, in transversal settings — where stabilizer measurements only remove local entropy — only $O(1)$ rounds should be required[17,32,158]. We test this hypothesis by studying performance as a function of N repeated CNOT gates between the two logical block, interleaved with two QEC cycles. We find optimal performance for several CNOTs per QEC round, and using postselection on decoding confidence (Methods), observe that three CNOTs per round is optimal at low error rates. These results are consistent with theoretical predictions for universal computation with these decoding techniques using $\lesssim$ 1 QEC round per transversal operation[32,158]. Importantly, we emphasize that even for universal computation with surface codes, no physical qubit reset is needed - once error locations are known, they can simply be tracked in-software throughout the physical Clifford circuit[57].

5.2 Below-threshold performance

QEC performance is enhanced by extracting as much information as possible about which errors have occurred. As roughly half of our errors are qubit loss, directly detecting the loss provides significant information in addition to inference on error locations inferred by stabilizer measurement results. Such direct detection of an error - a so-called erasure error - has been highlighted by recent theoretical studies to greatly improve QEC performance[151]. Accordingly, we use our lattice-based spin-to-position conversion readout[149] to directly detect qubit loss that occurred during the repeated QEC as seen in Fig 3a. As such erasures are not detected immediately upon happening, we refer to them as a delayed erasure[7].

We find that loss plays a major role in repeated QEC with neutral atoms. As theoretically predicted, we observe that qubit loss results in flickering stabilizer patterns which create time-correlations between detected error events (Fig 3b). Leveraging the delayed erasure information we create superchecks that are robust to the lost atom, and thereby data loss does not increase the detected error fraction (Fig 3c). This significant difference suggests that using the loss information in decoding can greatly improve performance[7], which we realize via most likely error methods as well as machine learning trained on simulated and experimental data (Fig 3f, Methods).

We now characterize repeated QEC performance as a function of the code distance d[2]. We find that after four rounds of repeated error correction, d=5 has a 2.10(12)x lower error than d=3, showing below-threshold behavior for this characterization circuit. We note that while clearly below threshold for the repeated QEC here, theoretically the threshold will get worse by roughly $\sim$ 1.1x when also incorporating $\sim$ 1 transversal gate per QEC round, and worse by $\sim$ 1.2x in the limit of many repeated rounds. The average error during the four rounds of error correction is 0.6% per round for the d=5, and reduces toward 0.1% per round on shots with no qubit loss (Fig 3e), again highlighting the role of loss and delayed erasure detection. Although these results indicate promis-

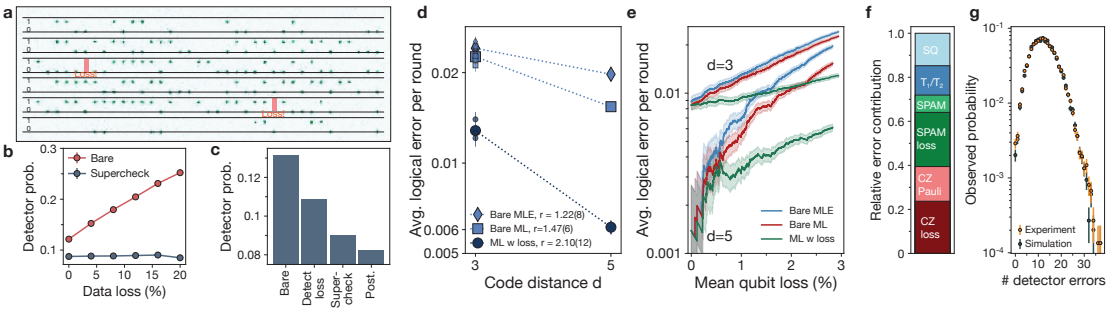


Figure 5.3: Repeated quantum error correction below-threshold leveraging loss-detection. a, Picture of repeated rounds of d=5 surface code using loss detection. Two losses occurred in this shot. **b,** Detector error probability as a function of data qubit loss. Detector probability is constant with use of a supercheck. **c,** Detector error probability for different metrics. **d,** Average logical error per round for surface codes after 4 rounds. Using machine learning (ML) with loss detection renders the error per round for d=5 as 2.1x better than d=3. **e,** Average logical error per round as a function of the mean qubit loss in the shot. **f,** Relative physical error contribution to detector error and threshold. **g,** Distribution of detector errors per shot. Experiment agrees with simulation, showing the absence of large-scale correlated errors.

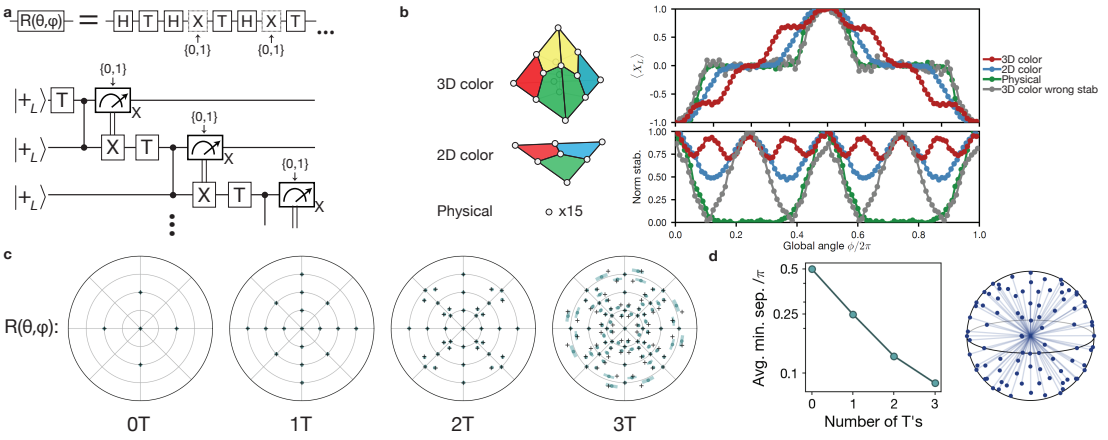


Figure 5.4: Synthesizing analog angles with fault-tolerant digital gates. a, Programmable angles $R(\theta, \varphi)$ are realised by an alternating sequence of H and T gates. Circuit shows an implementation of such a sequence with T implemented transversally and H implemented via teleportation. **b,** Subjecting codes of various dimensionality to global rotations. Plateaus are seen at robust angles where the stabilizer also revives. 3D codes, only with all positive stabilizer signs, have an additional plateau at T angles.**c,** Generated angles using entangled Reed-Muller codes, measured by single qubit tomography, for different total number of T gates. Experiment is consistent with theory (crosses) within statistical error. **d,** Minimum separation between generated angles decreases exponentially with length of sequence, as predicted by the Solovay-Kitaev theorem. Bloch sphere shows all measured angles.

ing performance, large-scale computation will require reducing physical errors by another factor of ~ 5; we provide a tabulation of estimated error contribution in Fig 3f and discuss in the Methods how these physical improvements can be made.

5.3 Analog evolution with digital gates

In the above we characterize the role that repeated error correction plays in removing errors created during logic operations, and find that such performance improves with increasing system size. Critically however, these correctable logical operations are necessarily discretized, digital gates. This is necessary in order to be correctable and robust to imperfections in, for example, miscalibrations in a global angle θ.

The Solovay-Kitaev theorem guarantees that arbitrary single-qubit rotations can be approximated to exponential precision using only digital gates like H and T[85,3,41]. Alternating layers of fault-tolerant H, T, and X gates enable access to a dense set of angles (Fig. 4a). Since H and T cannot simultaneously be implemented unitarily in a fault-tolerant way[47], we use teleportation to realize the H gate[116]. This approach bypasses the Eastin-Knill theorem[47] and naturally removes entropy via measurement and feedforward[91,137].

Figure 4b explores fault-tolerant T gate generation using 3D color codes. We compare unencoded qubits, 2D color codes (Steane)[136], and 3D color codes (Reed-Muller) under global θ rotations[116]. As in Fig. 1, errors are suppressed by the Zeno effect and error correction near allowed angles. Crucially, 3D codes exhibit robustness at T angles due to their transversal T gate[22]. Combined with teleportation-based H gates, this enables universal processing. While we study the core physics here with 3D color codes, the same principles underpin many techniques including magic state distillation (Methods)[57].

Figure 4c shows the set of angles generated by the circuit in Fig. 4a using up to three T gates. The

resulting logical states span a range of points on the Bloch sphere and match the expected angles with high precision (within statistical uncertainty; see Methods). As predicted by the Solovay-Kitaev theorem [85,1], the angular spacing between accessible states shrinks exponentially with the number of T gates, enabling precise synthesis of arbitrary rotations with only a logarithmic number of steps [3].

These results emphasize two key points. First, teleportation is a core enabling mechanism for universal processing, allowing analog rotations to be built from digital gates while breaking the constraints of the Eastin-Knill theorem. Due to our ability to realize transversal gates and teleportations, the state synthesis and state measurement natively accomplishes constant entropy operation [137] without the need for additional repeated stabilizer measurement. Second, the role of stabilizers differs fundamentally when generating logical magic. For Clifford circuits, stabilizer measurements primarily remove entropy. But for transversal T gates, correct stabilizer signs (e.g., $+1$ eigenvalues) are required to realize the operation (Fig. 4b). This reveals a direct connection between logical magic and the underlying physical entanglement structure (Methods) [73].

5.4 Deep circuits at constant entropy

To support deep-circuit computation on the logical level, it is important that the processor is kept at a constant entropy (Figure 5a) [3]. Computation inevitably introduces errors in the physical qubit state, as well as increasing entropy in the degrees of freedom of the physical atoms. Transversal teleportation is a direct approach to guarantee the computation remains at constant entropy, as the logical information propagates to a new block while the physical errors, and used physical atoms, get left behind on the previous block. Measuring this block then accomplishes reset, recooling, and reinitialization of the physical atoms. And when the block is prepared again, it is again at a low entropy, and this ensures that the overall computation stays at a constant entropy.

To realize such constant entropy behavior we ensure the atom internal state, atom temperature,

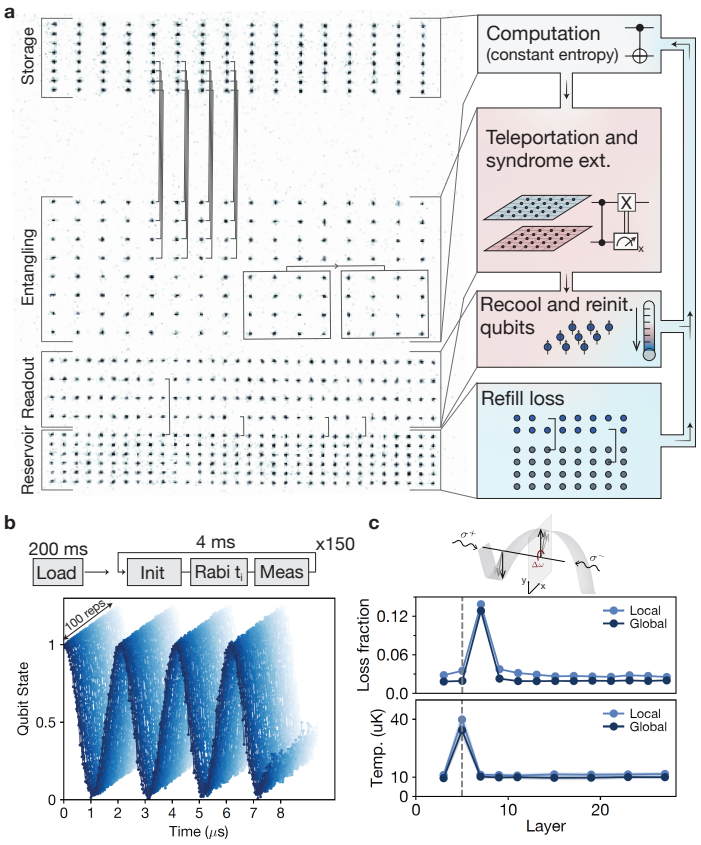


Figure 5.5: Architecture for constant-entropy computation. **a,** Image of processor and illustration of process for removing entropy generated by computation. **b,** Rabi oscillations measured by reuse of atoms with 150 cycles of non-destructive measurement and reinitialization. Each curve is one experimental repetition averaged over 200 atoms in parallel. 3D cooling methods are used in this subplot as coherence does not need to be preserved. **c,** Local cooling with 1D PGC and EIT. The finite magnetic field is compensated simply by applying a relative detuning between the two circularly polarized counterpropagating beams, lending to a rotating frame where the effective field is zero. The loss is constant as a function of time and recovers to steady state within one cycle after applying a perturbation to the system (dashed line). Additional shielding of the data qubits is provided by a 1530 shielding laser (not pictured here, see Methods).

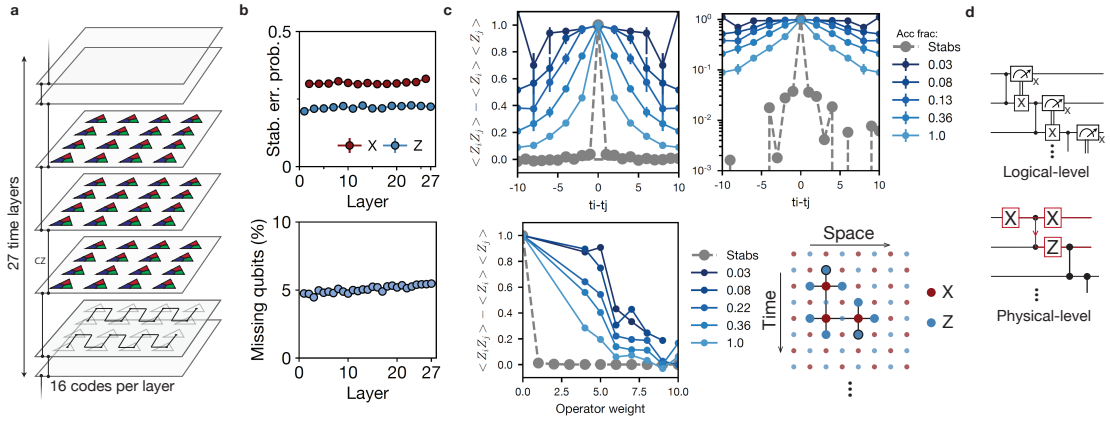


Figure 5.6: Deep circuits at constant entropy. a, Schematic of a 2D cluster state of Steane codes in space and time. The same atomic qubits are reused every other time layer for up to 27 layers. **b,** Repeated logical state preparation. The stabilizer error is constant as a function of cycle, and the loss grows slightly as the qubit reservoir is slowly depleted. **c,** Logical and physical correlations in 1D (top) and 2D (bottom) cluster states. Logical correlations persist while stabilizer error correlations rapidly decay. **d,** Logical-level evolution is unitary and allows the logical operator to propagate throughout the algorithm whereas physical-level evolution is dissipative and does not let physical errors propagate.

and atom filling are all reinitialized. We use the lattice to measure the internal state non-destructively (Fig. 1c). For reinitializing the atom temperature, whereas typical atomic physics techniques operate using 3D beams in zero magnetic field, we develop methods that enable high-fidelity imaging and cooling operating with 1D beams and in a finite field (Fig. 5c)[40,118,36] while also keeping the data qubits coherent (Methods)[74]. Finally, holes in the array are refilled by atoms from a reservoir. With all these methods combined, in Fig. 5c we measure the performance when subjecting the atoms to all the operations (except entanglement) involved for a deep circuit. Applying a perturbation on the 5th cycle, we find that the atomic filling and temperature quickly recover to a steady-state. The 1D methods nearly reproduce the conventional 3D performance, limited by tweezer-depth inhomogeneity (which we found at the end of our measurements). Such continuous operation experiments also allow for fast experimental cycle rates. As an example of this, we study fast Rabi calibrations with a 4-ms cycle time in Figure 5b.

Having established repeated reinitialization of the physical atom, we implement repeated reini-

tialization of the logical block [137] as a constant-entropy resource. We create 32 blocks of [[7,1,3]] codes (16 blocks per alternating group) for up to 27 layers. Importantly, we find that the stabilizer expectation value is constant as a function of the cycle number, indicating a steady-state internal entropy. We note that while the entropy is constant, our fidelity here is limited by multiple sub-optimalities in our circuit design structure which we chose for creating generic hypercube codes, and can be improved significantly for future experiments (Methods).

As an example algorithm, we entangle the [[7,1,3]] codes in both the time and space direction to realize 1D and 2D cluster states and probe the resulting correlations. In Figure 6c we plot the correlations on both the algorithm level as well as the physical error level by calculating the correlator $\langle Z_i Z_j \rangle - \langle Z_i \rangle \langle Z_j \rangle$ between events in both time and space. On the logical-qubit level we observe the expected correlations between logical outcomes corresponding to successful algorithm evolution, with a decay with distance corresponding to an implicit algorithmic error rate, which is dependent on the internal entropy as seen by a varying acceptance fraction. Conversely, for the internal stabilizers and physical errors, we observe the correlations are rapidly suppressed. This is because, whereas the logical decoding and logical feedforward (here done in-software) maintains unitary evolution on the logical level, the physical-level circuit results in errors traveling maximum 1 layer, and thereby guarantees that the algorithm proceeds at constant entropy.

5.5 Outlook

These experiments isolate and demonstrate the core mechanisms required for scalable, fault-tolerant implementations of arbitrary algorithms as well as their implementations with neutral atoms. Removing entropy from logic gates, below-threshold performance, synthesis of arbitrary angles, and deep circuits at constant entropy are all essential ingredients of large-scale processing with scalable overheads. Although these key elements are all demonstrated separately here, future work can com-

bine them to enable advanced fault-tolerant processing. Although here the reservoir slowly depletes, in a companion paper on a separate apparatus, we demonstrate coherent continuous operation on over 3000 atomic qubits by continuous replenishment, with techniques fully compatible with the methods presented here. Taken together, these techniques can enable small-scale fault-tolerant implementations of arbitrary algorithms. While currently approximately a factor of 2 below the threshold, this needs to be maintained in the deep-circuit context as well as improved by a factor of 3-5 (see Methods).

Our experiments further elucidate core physical mechanisms and highlight questions in QEC algorithm design. It is found that the transversal gate setting requires less entanglement within the logical block to realize logic operations compared to the lattice surgery setting. Conversely, entanglement is found to play a central role in realizing logical magic. Such an observation hints to a potential deeper understanding of the essential elements of protecting algorithmic outputs as opposed to individual qubits. The central role of transversal teleportation, and its ability to keep errors from propagating, also highlights that consideration of entropy on both the qubit states as well as the physical atoms is important in the design of fault-tolerant architectures. Establishing the core foundations of fault-tolerant computation, these results help illuminate the path forward toward large-scale computation.

6

What comes next?

6.1 Summary

In the above chapters, we described methods we developed for realizing quantum computation with neutral atoms as well as exploring error-corrected algorithms. We learned a lot about what is involved in running error-corrected algorithms, although there is clearly still much to learn. We found that understanding the core physics of the logical qubits is illuminating to understanding

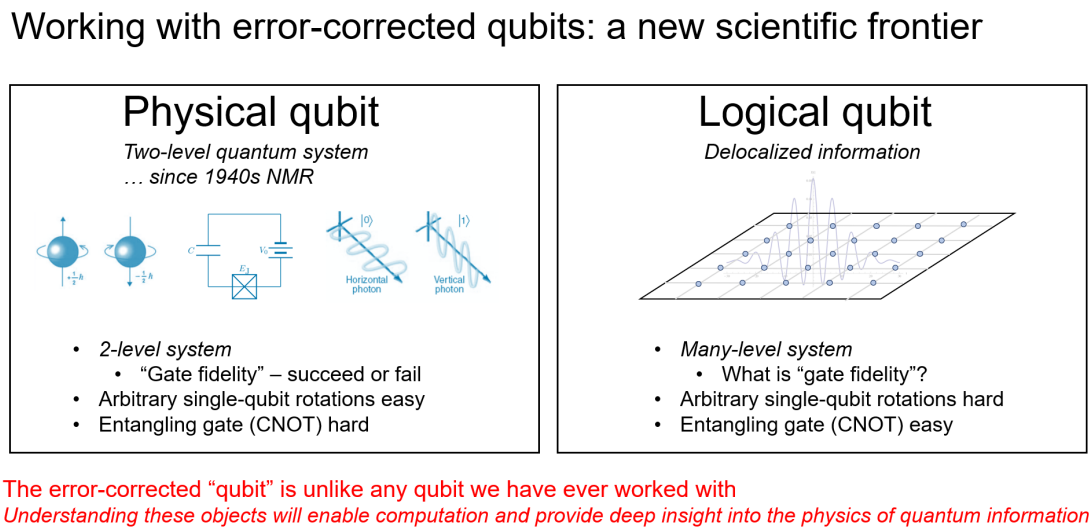


Figure 6.1: Physical and logical qubits are very different objects. Whereas we have experimentally studied physical qubits - i.e. two-level quantum systems - for nearly a century, logical qubits are very new to us. They behave in very different and strange ways. If we are using quantum computers in the future, which we expect is very likely, then it is extremely likely that they will be made out of logical qubits. How exciting and fun it is that we get to start exploring them and see how they behave.

ways to improve the quantum error correction.

The prudent question: what comes next from here? For one, it is becoming extremely clear that quantum error correction actually does work, and that you can suppress errors by increasing code distance. The progress will - quite literally - continue to be exponential from here.

6.2 Science of logical qubits

What (hopefully) was also clear from the above is that logical qubits are extremely interesting and understanding them is a rich science and physics frontier. We have been working with physical qubits since the 40's (spins in NMR), and remarkably, all these qubits behave the same - whether they are photons, spins, superconductors, mechanical oscillators, or so on. The reason these qubits behave the same is that they are two-level systems. We have developed extremely good intuition for these - Rabi oscillations, notions of gate fidelities, and so on. Logical qubits however, are no-

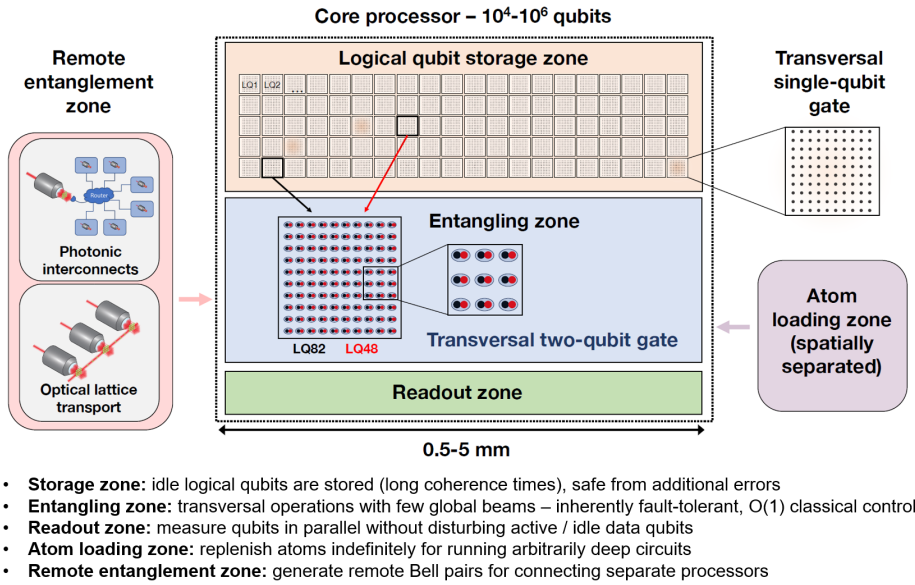


Figure 6.2: Expanding current tools to a large-scale processor. Atom re-loading zones facilitate running arbitrarily deep circuits, and remote entanglement zones facilitate using networking to attain larger sizes as well as explore other applications of quantum systems.

tably *not* two-level systems - they are many-level systems. As such, they can behave in complex and surprising ways! For example, with a physical qubit the notion of "gate fidelity" is very clear and well-defined: you either do the operation you asked for, or you do something exactly orthogonal with some probability p. What is gate fidelity with logical qubits? As discussed above, it is clearly entropy-dependent, so even just having a single number cannot be a good enough description anymore. As we also discussed above, it often doesn't make too much sense to think about the logical qubit block, it often makes more sense to think about the logical operator which spreads throughout the circuit. This is not such a surprise, since a "logical qubit" is anyway just an abstracted logical unit of quantum information, which can spread throughout the circuit in various complex ways. The things that physical qubits and logical qubits can natively do are also strikingly different - the "native evolution" of physical qubits is often analog evolution under Hamiltonian dynamics, whereas the "native evolution" of logical qubits is often digital evolution under Clifford circuits.

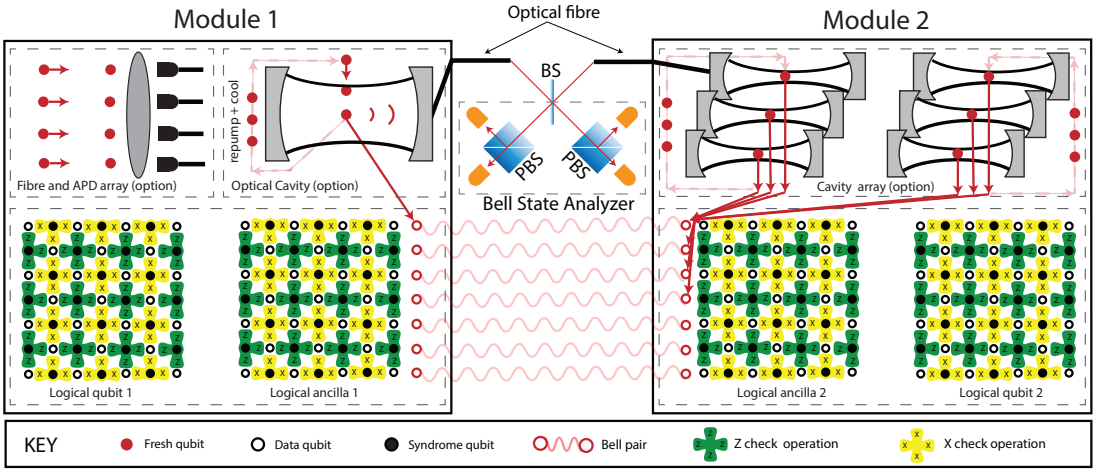


Figure 6.3: Connecting modules via remote Bell pairs. Logical qubits in each module are realized using an array of atoms and connected using teleported gates. The Bell pairs are generated using either a lens, an optical cavity, or an array of optical cavities. See Ref.[134] for additional details. Importantly, the fidelity threshold of required Bell pairs in such a scheme is only 10%[58,134], as the surface codes are joined on a boundary, effectively given repetition-code-like thresholds due to the reduced dimensionality.

If we are using quantum computers in the future - which we expect is very likely - then it is extremely likely that they will be running algorithms via logical qubits. How exciting and fun it is that we get to start experimentally exploring the physics of these objects and finding that they behave so differently than what we're used to.

6.3 Toward future processors and QEC

That said, we clearly still have a very far way to go on a technical level as well. There are many ideas about how these things may continue. A prominent perspective is that at some point we will need to treat each individual processor as a "module" and begin to link them together via remote Bell pairs (see Figs. 6.2, 6.3). Significant progress is underway with such methods[43,46], and is receiving more and more attention. Such capabilities will be useful not only for making larger quantum computers, but also for exploring a rich variety of possibilities when quantum computers are connected

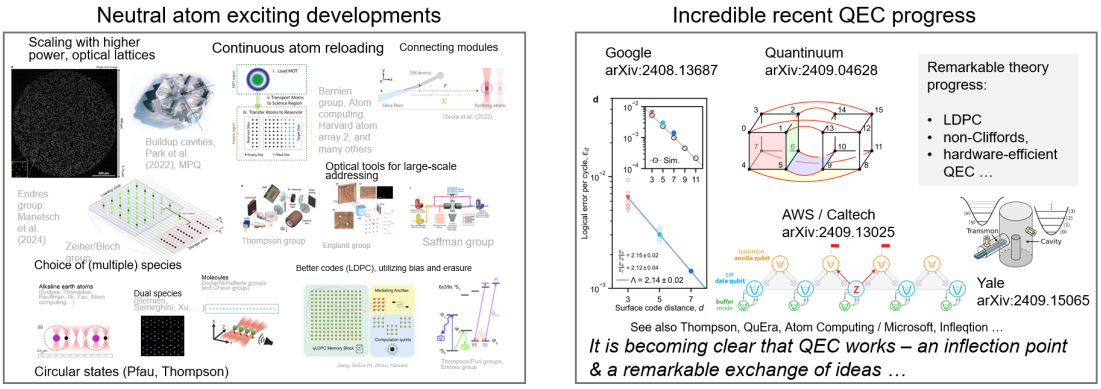


Figure 6.4: Exciting recent developments in neutral atoms and quantum error correction..

on a quantum network. This is an exciting prospect especially given that we do not yet know exactly what we will use such quantum systems for.

One may wonder if there are ways to improve things significantly beyond our current approach. It will be an exciting next few years where creative ideas can be explored. It seems likely that there will be significant progress to come in both neutral atoms and QEC broadly, and that the field is at a true inflection point (see Fig. 6.4).

7

Exploring applications of quantum processors

7.1 Overview

In the above, we spent a lot of time discussing the physics of how to build quantum processors - on both an atomic physics level as well as on the logical qubit level. But let's say we had one tomorrow,

and it worked exactly the way we wanted it to. What would we even *do* with it? This is also a broad question in the field [113].

A natural application of quantum computers is to simulate quantum mechanics [55]. In the below I overview various quantum simulations and computations we studied, covering quantum phases of matter, non-equilibrium and non-thermalizing dynamics, topological quantum spin liquids, combinatorial optimization, quantum coarsening dynamics, hybrid analog-digital quantum simulations, and digital quantum simulations of non-abelian spin liquids and fermionic dynamics. Only a brief overview and key figures will be presented here, and references to additional description are included in each section.

7.2 Quantum phases of matter on a 256-atom programmable quantum simulator

Additional details for this section can be found in Ref.[49].

Motivated by far-reaching applications ranging from quantum simulations of complex processes in physics and chemistry to quantum information processing, a broad effort is currently underway to build large-scale programmable quantum systems. Such systems provide unique insights into strongly correlated quantum matter, while at the same time enabling new methods for computation and metrology. Here, we demonstrate a programmable quantum simulator based on deterministically prepared two-dimensional arrays of neutral atoms, featuring strong interactions controlled via coherent atomic excitation into Rydberg states. Using this approach, we realize a quantum spin model with tunable interactions for system sizes ranging from 64 to 256 qubits. We benchmark the system by characterizing high-fidelity antiferromagnetically ordered states and demonstrating quantum critical dynamics consistent with an Ising quantum phase transition in (2+1) dimensions. We then create and study several new quantum phases that arise from the interplay between interactions and coherent laser excitation, experimentally map the phase diagram, and investigate the role of quantum fluctuations. Offering a new lens into the study of complex quantum matter, these observations pave the way for investigations of exotic quantum phases, non-equilibrium entanglement dynamics, and hardware-efficient realization of quantum algorithms.

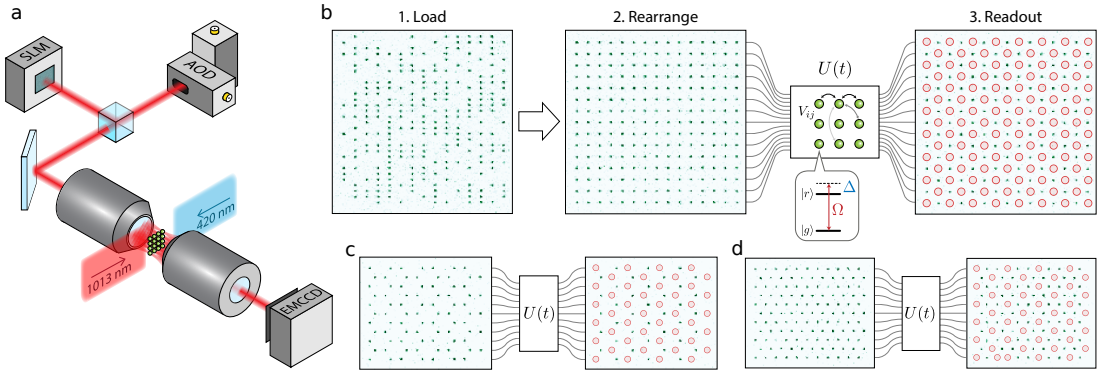


Figure 7.1: Programmable two-dimensional arrays of strongly-interacting Rydberg atoms. a. Atoms are loaded into a 2D array of optical tweezer traps and rearranged into defect-free patterns by a second set of moving tweezers. Lasers at 420 nm and 1013 nm drive a coherent two-photon transition in each atom between ground state $|g\rangle = |5S_{1/2}, F = 2, m_F = -2\rangle$ and Rydberg state $|r\rangle = |70S_{1/2}, m_j = -1/2, m_I = -3/2\rangle$. **b.** Fluorescence image of initial random loading of atoms, followed by rearrangement to a defect-free 15×15 (225 atoms) square array. After this initialization, the atoms evolve coherently under laser excitation with Rabi frequency $\Omega(t)$, detuning $\Delta(t)$, and long-range interactions V_{ij}, realizing a many-body unitary $U(t)$. Finally, the state of each atom is read out, with atoms excited to $|r\rangle$ detected as loss and marked with red circles. Shown on the far right is an example measurement following quasi-adiabatic evolution into the checkerboard phase. **c, d.** Similar evolution on honeycomb and triangular lattices result in analogous ordered phases of Rydberg excitations with filling $1/2$ and $1/3$, respectively.

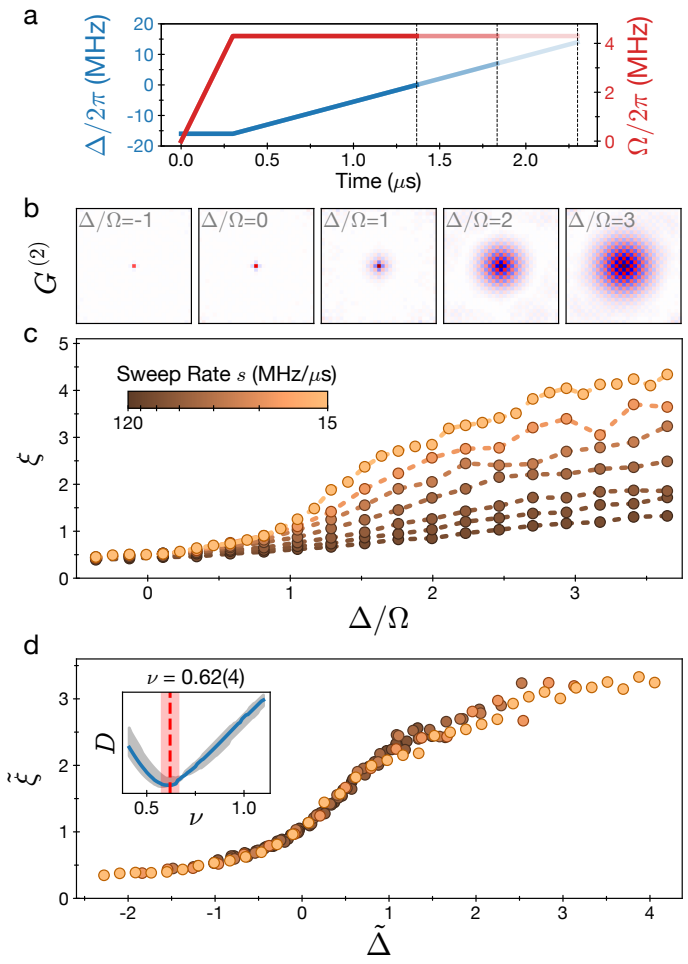


Figure 7.2: Observation of the (2+1)D Ising quantum phase transition on a 16×16 array. a. The transition into the checkerboard phase is explored using a linear detuning sweep $\Delta(t)$ at constant Ω. The resulting checkerboard ordering is measured at various endpoints. **b.** Example of growing correlations $G^{(2)}$ with increasing Δ/Ω along a linear sweep with sweep rate $s = 15$ MHz/μs. **c.** Growth of correlation length ξ for s spanning an order of magnitude (15, 21, 30, 42, 60, 85, and 120 MHz/μs). ξ used here measures correlations between the coarse-grained local staggered magnetization (see Methods). **d.** For an optimized value of the critical exponent ν, all curves collapse onto a single universal curve when rescaled relative to the quantum critical point Δ_c. Inset: distance D between all pairs of rescaled curves as a function of ν (see Methods). The minimum at $\nu = 0.62(4)$ (red dashed line) yields the experimental value for the critical exponent (red and gray shaded regions indicate uncertainties) and is consistent with the theoretical prediction $\nu = 0.629$.

7.3 Controlling quantum many-body dynamics in driven Rydberg atom arrays

Additional details for this section can be found in Ref.[20].

Controlling non-equilibrium quantum dynamics in many-body systems is an outstanding challenge as interactions typically lead to thermalization and a chaotic spreading throughout Hilbert space. We experimentally investigate non-equilibrium dynamics following rapid quenches in a many-body system composed of 3 to 200 strongly interacting qubits in one and two spatial dimensions. Using a programmable quantum simulator based on Rydberg atom arrays, we probe coherent revivals corresponding to quantum many-body scars. Remarkably, we discover that scar revivals can be stabilized by periodic driving, which generates a robust subharmonic response akin to discrete time-crystalline order. We map Hilbert space dynamics, geometry dependence, phase diagrams, and system-size dependence of this emergent phenomenon, demonstrating novel ways to steer entanglement dynamics in many-body systems and enabling potential applications in quantum information science.

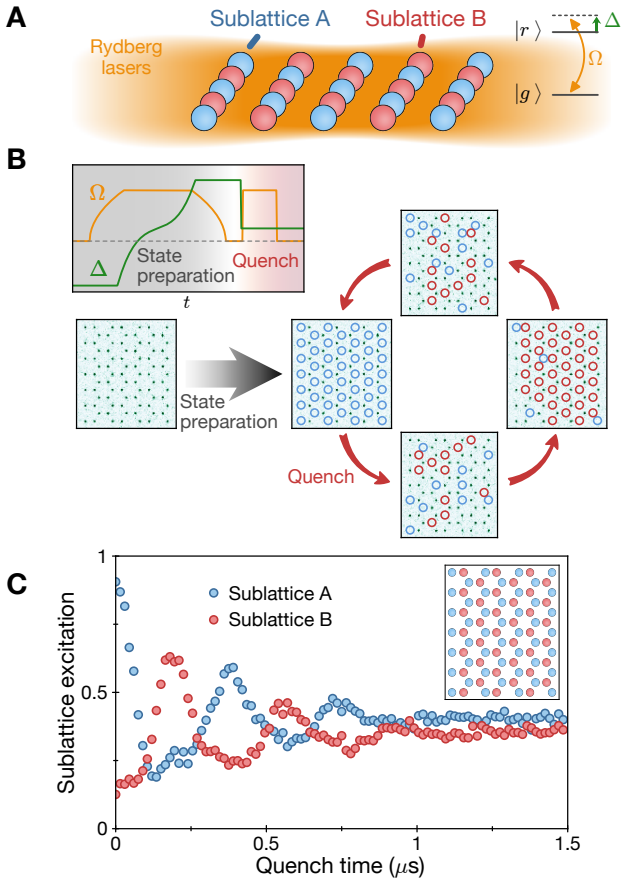


Figure 7.3: Experimental investigations of quantum many-body scars. (A) Two-dimensional atom array subject to global Rydberg lasers with Rabi frequency Ω and detuning Δ. **(B)** A quasi-adiabatic ramp of Δ and Ω prepares an antiferromagnetic state $|\mathrm{AF}_1\rangle$ with sublattice A excited, and a detuning quench launches non-equilibrium dynamics. Atoms in $|g\rangle$ are imaged in optical tweezers via fluorescence while atoms in $|r\rangle$ (empty circles) are expelled and detected as atom loss. **(C)** The Rydberg population on sublattices A and B undergo periodic oscillations (Inset: geometry used here).

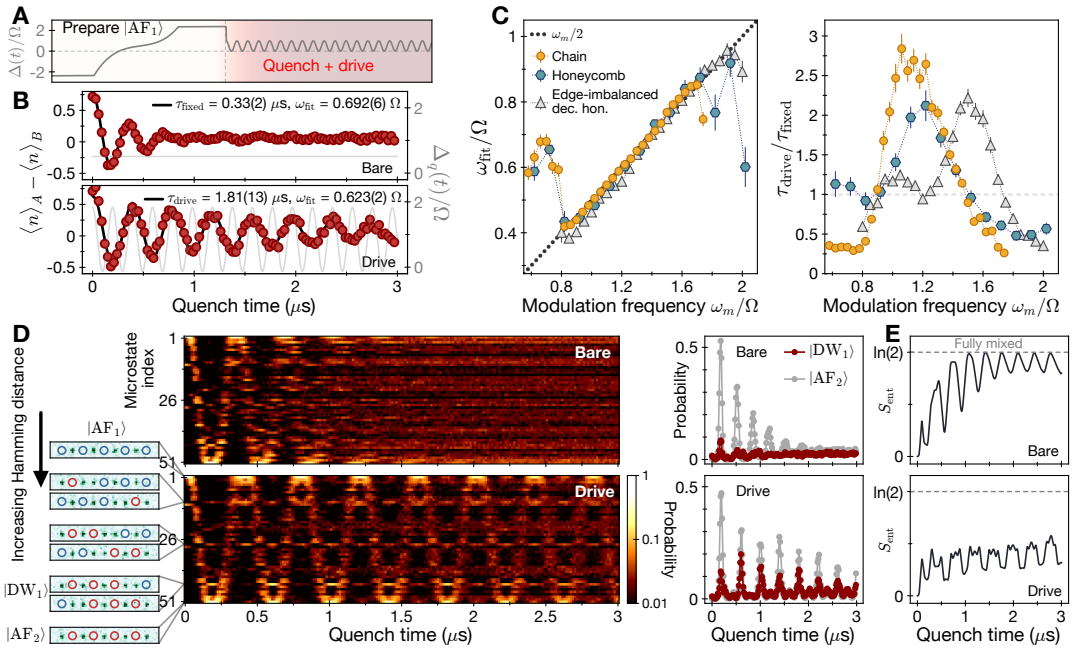


Figure 7.4: Emergent subharmonic locking and stabilization. (A) Pulse sequence showing state preparation and quench with $\Delta_q(t)$. **(B)** Scar dynamics on a chain during quench to fixed optimal detuning (bare) with lifetime τ_{fixed}, and time-dependent detuning (drive) with modulation frequency ω_m = 1.24 Ω and lifetime τ_{drive}. The drive increases the scar lifetime and changes its frequency to $\omega_m/2$. **(C)** Scar lifetime and response frequency as a function of ω_m, showing a lifetime increase and subharmonic locking. **(D)** Dynamics of the entire Hilbert space measured with experimental snapshots (0.5 million total bit strings). The microstates of the constrained Hilbert space are ordered by n_A — n_B, or equivalently by Hamming distance from $|\mathrm{AF}_1\rangle$. Right subplots highlight $|\mathrm{AF}_2\rangle$ and a state with a domain wall $|\mathrm{DW}_1\rangle$. **(E)** Reduced density matrix of a single atom in a chain (numerics) shows that driving reduces the growth of entanglement entropy S_{ent}.

7.4 Probing topological spin liquids on a programmable quantum simulator

Additional details for this section can be found in Ref.[128].

Quantum spin liquids, exotic phases of matter with topological order, have been a major focus of explorations in physical science for the past several decades. Such phases feature long-range quantum entanglement that can potentially be exploited to realize robust quantum computation. We use a 219-atom programmable quantum simulator to probe quantum spin liquid states. In our approach, arrays of atoms are placed on the links of a kagome lattice and evolution under Rydberg blockade creates frustrated quantum states with no local order. The onset of a quantum spin liquid phase of the paradigmatic toric code type is detected by evaluating topological string operators that provide direct signatures of topological order and quantum correlations. Its properties are further revealed by using an atom array with nontrivial topology, representing a first step towards topological encoding. Our observations enable the controlled experimental exploration of topological quantum matter and protected quantum information processing.

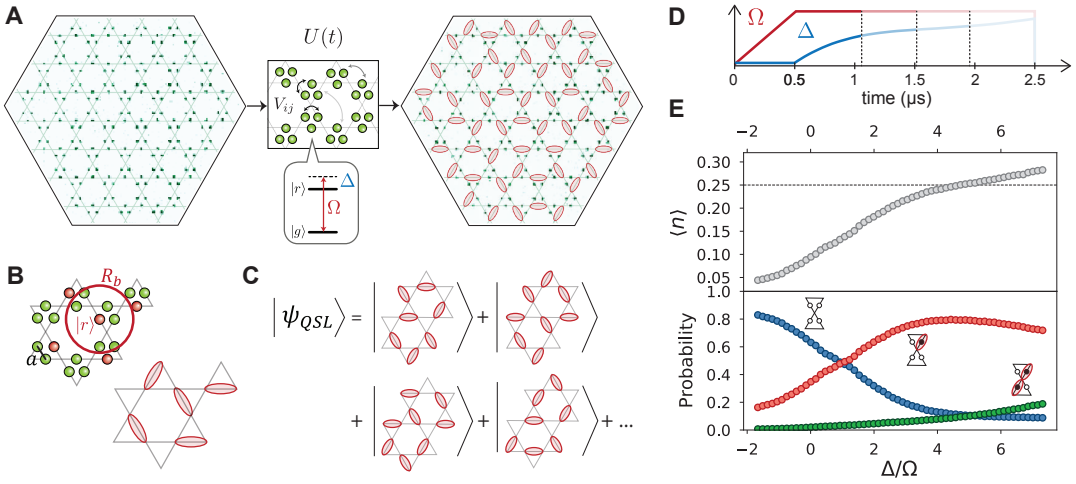


Figure 7.5: Dimer model in Rydberg atoms arrays. (**A**) Fluorescence image of 219 atoms arranged on the links of a kagome lattice. The atoms, initially in the ground state $|g\rangle$, evolve according to the many-body dynamics $U(t)$. The final state of the atoms is determined via fluorescence imaging of ground state atoms. Rydberg atoms are marked with red dimers on the bonds of the kagome lattice. (**B**) We adjust the blockade radius to $R_b/a = 2.4$, by choosing $\Omega = 2\pi \times 1.4$ MHz and $a = 3.9\,\mu$m, such that all six nearest neighbors of an atom in $|r\rangle$ are within the blockade radius R_b. A state consistent with the Rydberg blockade at maximal filling can then be viewed as a dimer covering of the kagome lattice, where each vertex is touched by exactly one dimer. (**C**) In the ideal limit, QSL state corresponds to a coherent superposition of exponentially many dimer coverings. (**D**) Detuning $\Delta(t)$ and Rabi frequency $\Omega(t)$ used for quasi-adiabatic state preparation. (**E**) (Top) Average density of Rydberg excitations $\langle n \rangle$ in the bulk of the system, excluding the outer three layers. (Bottom) Probabilities of empty vertices in the bulk (monomers), vertices attached to a single dimer, or to double dimers (weakly violating blockade). After $\Delta/\Omega \sim 3$, the system reaches $\sim 1/4$ filling, where most vertices are attached to a single dimer, consistent with an approximate dimer phase.

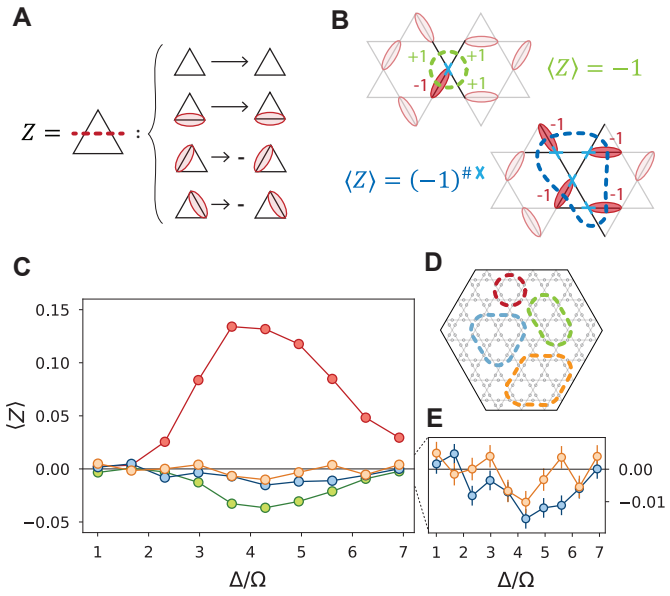


Figure 7.6: Detecting a dimer phase via diagonal string operator. (**A**) The Z string operator measures the parity of dimers along a string. (**B**) A perfect dimer covering always has exactly one dimer touching each vertex of the array, so that $\langle Z \rangle = -1$ around a single vertex and $\langle Z \rangle = (-1)^{\#\text{enclosed vertices}}$ for larger loops. (**C**) Z parity measurements following the quasi-adiabatic sweep of Fig. 4.1D, with the addition of a 200 ns ramp-down of Ω at the end to optimize preparation. At different endpoints of the sweep and for different loop sizes (inset), we measure a finite $\langle Z \rangle$, consistent with an approximate dimer phase.

7.5 Quantum optimization of maximum independent set using Rydberg atom arrays

Additional details for this section can be found in Ref.[48].

Realizing quantum speedup for practically relevant, computationally hard problems is a central challenge in quantum information science. Using Rydberg atom arrays with up to 289 qubits in two spatial dimensions, we experimentally investigate quantum algorithms for solving the Maximum Independent Set problem. We use a hardware-efficient encoding associated with Rydberg blockade, realize closed-loop optimization to test several variational algorithms, and subsequently apply them to systematically explore a class of graphs with programmable connectivity. We find the problem hardness is controlled by the solution degeneracy and number of local minima, and experimentally benchmark the quantum algorithm's performance against classical simulated annealing. On the hardest graphs, we observe a superlinear quantum speedup in finding exact solutions in the deep circuit regime and analyze its origins.

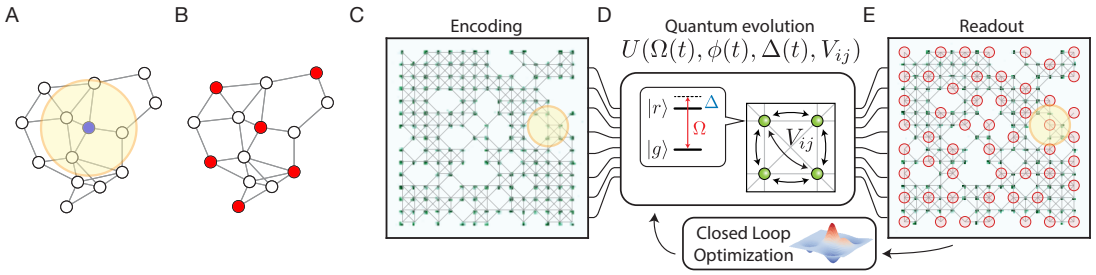


Figure 7.7: Hardware-efficient encoding of the maximum independent set using Rydberg atom arrays. A. An example of a unit disk graph, with any single vertex (e.g. the blue vertex) being connected to all other vertices within a disk of unit radius. **B.** A corresponding MIS solution (denoted by the red nodes). **C.** The MIS problem is encoded with atoms placed at the vertices of the target graph and with interatomic spacing chosen such that the unit disk radius of the graph corresponds to the Rydberg blockade radius. Shown is an example fluorescence image of atoms, with gray lines added to indicate edges between connected vertices. **D.** The system undergoes coherent quantum many-body evolution under a programmable laser drive $(\Omega(t), \varphi(t), \Delta(t))$ and long-range Rydberg interactions V_{ij}. **E.** A site-resolved projective measurement reads out the final quantum many-body state, with atoms excited to the Rydberg state (red circles) corresponding to vertices forming an independent set. A classical optimizer uses the results to update the parameters of the quantum evolution $(\Omega(t), \varphi(t), \Delta(t))$ to maximize a figure of merit for finding an MIS.

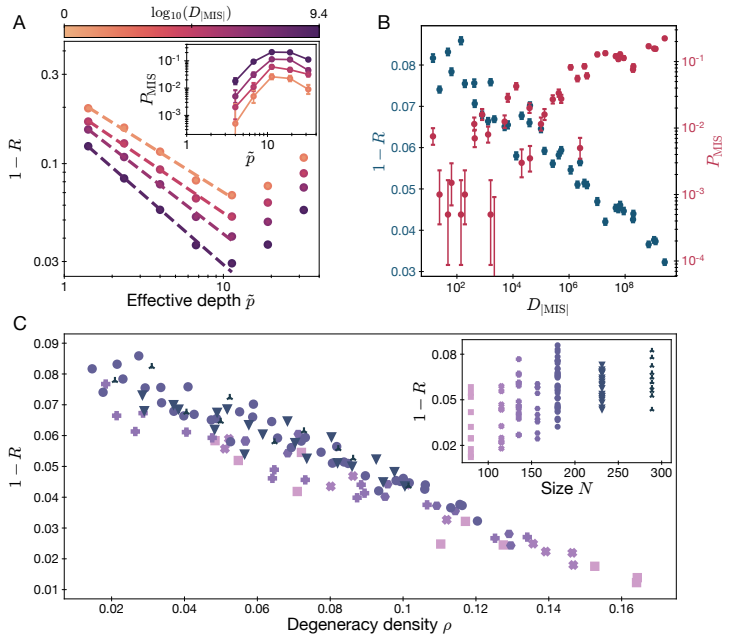


Figure 7.8: Quantum algorithm performance across different graphs. A. The approximation error $1 - R$ for an optimized quasi-adiabatic sweep plotted as a function of effective depth $\tilde{p}$ on four graphs of the same size ($N = 180$ vertices), showing strong dependence on the number of MIS solutions (MIS degeneracy) D_{MIS} (inset: corresponding MIS probability P_{MIS} vs. $\tilde{p}$). **B.** At a fixed depth $\tilde{p} = 20, 1 - R$ and P_{MIS} for various 180-vertex graphs are strongly correlated with D_{MIS}. **C.** At the same effective depth $\tilde{p} = 20, 1 - R$ for 115 graphs of different sizes ($N = 80$–289) and MIS degeneracies D_{MIS} exhibit universal scaling with the degeneracy density $\rho \equiv \log(D_{MIS})/N$ (inset: data plotted as a function of N). Error bars for P_{MIS}, here and throughout the text, denote the 68% confidence interval.

7.6 Quantum coarsening and collective dynamics on a programmable simulator

Additional details for this section can be found in Ref.[103].

Understanding the collective quantum dynamics of nonequilibrium many-body systems is an outstanding challenge in quantum science. In particular, dynamics driven by quantum fluctuations are important for the formation of exotic quantum phases of matter, fundamental high-energy processes, quantum metrology, and quantum algorithms. Here, we use a programmable quantum simulator based on Rydberg atom arrays to experimentally study collective dynamics across a (2+1)D Ising quantum phase transition. After crossing the quantum critical point, we observe a gradual growth of correlations through coarsening of antiferromagnetically ordered domains. By deterministically preparing and following the evolution of ordered domains, we show that the coarsening is driven by the curvature of domain boundaries, and find that the dynamics accelerate with proximity to the quantum critical point. We quantitatively explore these phenomena and further observe long-lived oscillations of the order parameter, corresponding to an amplitude (Higgs) mode. These observations offer a unique viewpoint into emergent collective dynamics in strongly correlated quantum systems and nonequilibrium quantum processes.

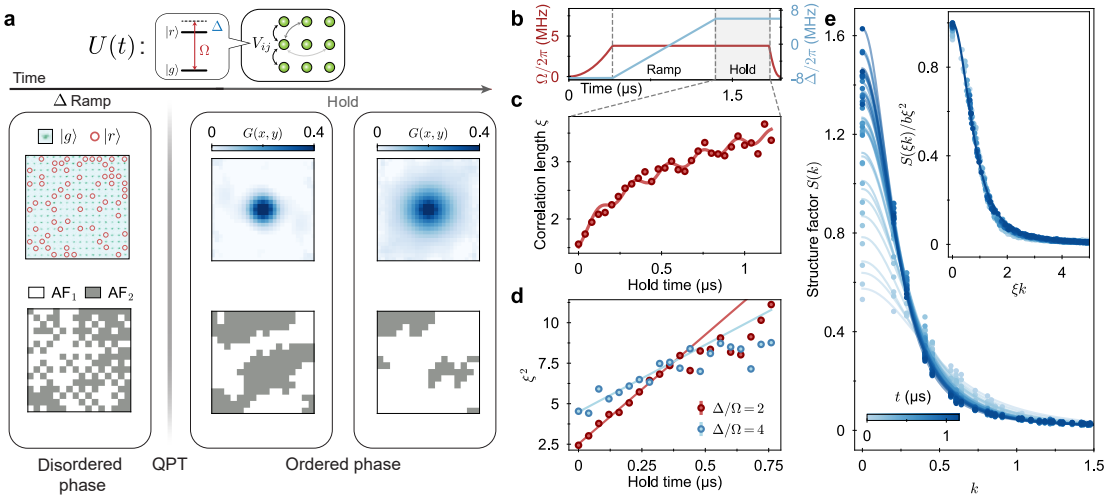


Figure 7.9: Ordering via coarsening. a, Atoms arranged in a square lattice are prepared in the ground state $|g\rangle$ and driven from the disordered phase, shown here as a fluorescence image (atoms in $|r\rangle$ are detected as loss and are indicated with red circles), into an antiferromagnetic ordered phase, across a quantum phase transition. The subsequent coarsening dynamics lead to the progressive development of long-range order (shown here using the real-space correlation function $G(x, y)$ and experimental snapshots). The snapshots show the local single-spin-flip-corrected staggered magnetization, with white (dark grey) shading indicating AF_1 (AF_2) orders (see Methods). **b,** The detuning $\Delta(t)$ is swept linearly, with a fixed Rabi frequency Ω, to a final value Δ in the ordered phase, where it is held constant for the duration of the hold time. **c,** The correlation length ξ grows over the course of the hold time. The solid line is a fit to $(c_0 + c_1 t)^\alpha + c\cos(\omega t + \varphi)$, representing sinusoidal oscillations on top of a power-law growth. Errorbars are smaller than markers. **d,** The early-time growth of the squared correlation length is consistent with a linear increase with time, as expected for coarsening dynamics with a nonconserved order parameter. The dynamics are faster at lower Δ/Ω, closer to the critical point. **e,** The structure factor $S(k, t)$ of the staggered magnetization, for different hold times, at $\Delta/\Omega = 2.0$. Inset: the structure factors collapse onto a single curve when the time axis is rescaled by the correlation length $\xi(t)$ and the magnitude of the structure factor is rescaled by a time-dependent amplitude $\xi^2(t)b(t)$ (see Methods). For the data shown here, $\Omega/2\pi = 3.8$ MHz.

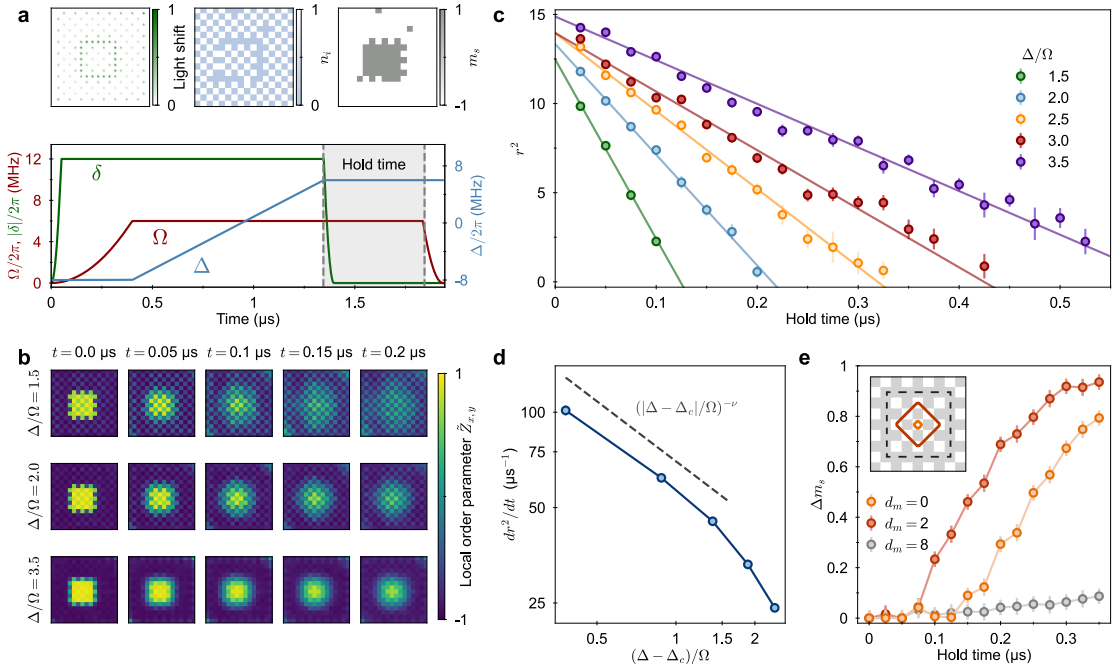


Figure 7.10: Dynamics of seeded domains. a, We deterministically create a centered domain of AF_2 order surrounded by a backdrop of AF_1 order, by "pinning" each ground-state atom in the target ordering using a local light shift δ_i, while sweeping $\Delta(t)$ and maintaining a constant Rabi frequency Ω. Upper left: target amplitude of the local detuning pattern; δ_i is inversely proportional to the number of neighboring Rydberg atoms surrounding each ground-state atom. Center: Rydberg population of a single shot immediately after the local detuning is quenched off. Right: local staggered magnetization, demarcating the prepared domains. **b,** Evolution of the average local staggered magnetization with time, as the prepared state evolves under different values of Δ/Ω. **c,** The squared radius of the central domain, r^2, decreases linearly with time. **d,** The rate of change of the area, r^2, increases with proximity to the critical point, $(\Delta - \Delta_c)/\Omega$. The dashed line is a guide to the eye for the theoretically expected scaling relationship. **e,** Change in the radially averaged local staggered magnetization with Manhattan distances $d_m = 0, 2, 8$ from the center of the injected square domain. The order initially changes near the domain wall and begins to change in the center of the domain only at later times. Atoms is the bulk of the dominant ordering far from the domain wall, in gray, remain close to their initially prepared states. The dynamics are plotted for $\Delta/\Omega = 2.5$ and $\Omega/2\pi = 6.0$ MHz.

7.7 Hybrid analog-digital quantum simulation

Additional details for this section can be found in Ref.[19].

We here explore the applications of atom movement to quantum simulation. In particular, we perform hybrid, modular quantum circuits composed of analog Hamiltonian evolution, reconfiguration, and digital gates. Together, these tools open a wide variety of new possibilities in quantum simulation and many-body physics. As a specific example, we measure the Renyi entanglement entropy after a quantum quench by effectively interfering two copies of a many-body system, and observe the non-monotonic oscillations in entanglement entropy characteristic of non-thermalizing quantum many-body scars[39,78].

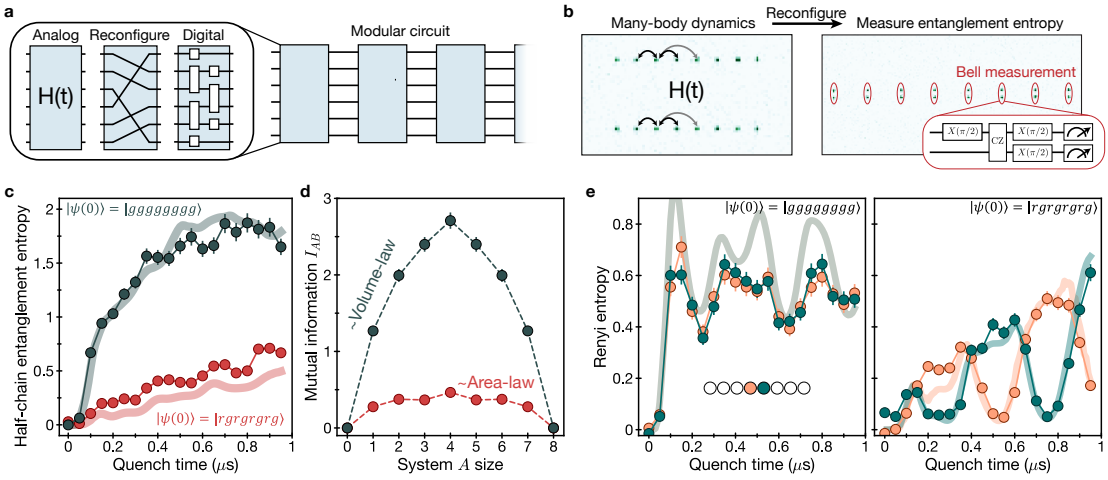


Figure 7.11: Dynamic reconfigurability for hybrid analog-digital quantum simulation. a, Hybrid quantum circuit combining coherent atom transport with analog Hamiltonian evolution and digital quantum gates. **b,** Measuring entanglement entropy in a many-body Rydberg system via two-copy interferometry. **c,** Measured half-chain Renyi entanglement entropy after many-body dynamics following quenches on two 8-atom systems. Quenching from $|ggggg\ldots\rangle$ ($|g\rangle \equiv |1\rangle$) results in rapid entropy growth and saturation, signifying quantum thermalization. Quenching from $|rgrg\ldots\rangle$ reveals a significantly slower growth of entanglement entropy. **d,** Measuring the mutual information at 0.5-μs quench time reveals a volume-law scaling for the thermalizing $|ggggg\ldots\rangle$ state, and an area-law scaling for the scarring $|rgrg\ldots\rangle$ state. **e,** The single-site Renyi entropies for sites in the middle of the chain quickly increase and saturate for the $|ggggg\ldots\rangle$ quench, but show large oscillations for the $|rgrg\ldots\rangle$ quench. Solid curves are results of exact numerical simulations for the isolated quantum system under H_{Ryd} with no free parameters (see Methods for details of data processing). Error bars represent 1 standard deviation.

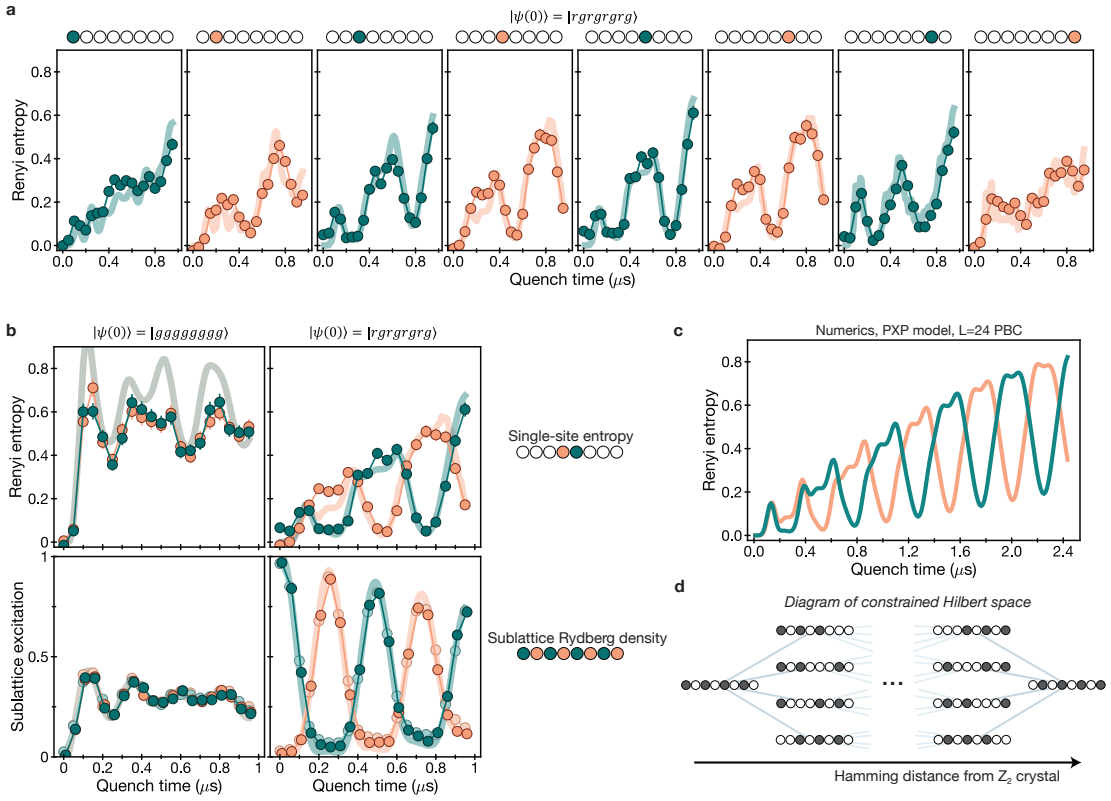


Figure 7.12: Local observables and entanglement entropy for quantum many-body scars. **a,** Experimentally measured single-site entropy for each site in the 8-atom chain when quenching from the scarred $|\mathbb{Z}_2\rangle$ state, including the classical entropy subtraction. Solid curves plot exact, ideal (imperfection-free) numerics of H_{Ryd}; excellent agreement between data and numerics is found for every atom in the chain. **b,** (Top) Same data as Fig. 7.11, showing single-site entropy of the middle two atoms in the chain, for two different initial states[70]. (Bottom) Measurements of the many-body state in the Z-basis with the interferometry circuit turned off. Characteristic of the scars from the $|\mathbb{Z}_2\rangle = |rgrgrgrg\rangle$ state, the Rydberg excitation probability on the sublattices exhibits periodic oscillations[20]. In the bottom row, the dark data points are measured in the $\{|1\rangle, |r\rangle\}$ basis, and the faint data points are measured in the $\{|0\rangle, |1\rangle\}$ basis after the coherent mapping sequence. Measurements in both bases agree well with exact numerics (solid lines), which we emphasize has no free fit parameters and does not account for any experimental imperfections, such as detection infidelity. Moreover, the data indicate the high fidelity of preparation into the $|\mathbb{Z}_2\rangle$ state by use of local Rydberg π pulses. In plotting, we delay the theory curves and the $\{|1\rangle, |r\rangle\}$ basis measurement by 10 ns to account for the fact that the Raman π pulse we apply cuts off the final 10 ns of the Rydberg evolution, when measuring in the $\{|0\rangle, |1\rangle\}$ basis. **c,** Numerical simulations of the single-site Renyi entropy on two adjacent sites in the idealized 'PXP' model of perfect nearest-neighbor blockade[143]. The system size is 24 atoms with periodic boundary conditions, showing the same out-of-phase oscillations in the entanglement entropy of the two sublattices. **d,** Diagram of the constrained Hilbert space of the system[143]. The early-time, out-of-phase entropy oscillations[70] of the scars can be understood in this constrained Hilbert space picture, where the scar dynamics are known to take the state from the left end ($|rgrgrgrg\rangle$) to the right end ($|grgrgrgr\rangle$) (dark circles represent $|r\rangle$ and white circles represent $|g\rangle$)[143]. Near these crystalline ends of this constrained Hilbert space, the Rydberg atoms can fluctuate (high entropy), but the ground state atoms are pinned (low entropy). Our analysis shows that entanglement between atoms on the same sublattice contributes to the eventual degradation of the revival fidelity of the $|\mathbb{Z}_2\rangle$ state.

7.8 Probing topological matter and fermion dynamics on a neutral-atom quantum computer

Additional details for this section can be found in Ref.[54].

Quantum simulations of many-body systems are among the most promising applications of quantum computers. In particular, models based on strongly-correlated fermions are central to our understanding of quantum chemistry and materials problems, and can lead to exotic, topological phases of matter. However, due to the non-local nature of fermions, such models are challenging to simulate with qubit devices. Here we realize a digital quantum simulation architecture for two-dimensional fermionic systems based on reconfigurable atom arrays. We utilize a fermion-to-qubit mapping based on Kitaev's model on a honeycomb lattice, in which fermionic statistics are encoded using long-range entangled states. We prepare these states efficiently using measurement and feed-forward, realize subsequent fermionic evolution through Floquet engineering with tunable entangling gates interspersed with atom rearrangement, and improve results with built-in error detection. Leveraging this fermion description of the Kitaev spin model, we efficiently prepare topological states across its complex phase diagram and verify the non-Abelian spin liquid phase by evaluating an odd Chern number. We further explore this two-dimensional fermion system by realizing tunable dynamics and directly probing fermion exchange statistics. Finally, we simulate strong interactions and study dynamics of the Fermi-Hubbard model on a square lattice. These results pave the way for digital quantum simulations of complex fermionic systems for materials science, chemistry, and high-energy physics.

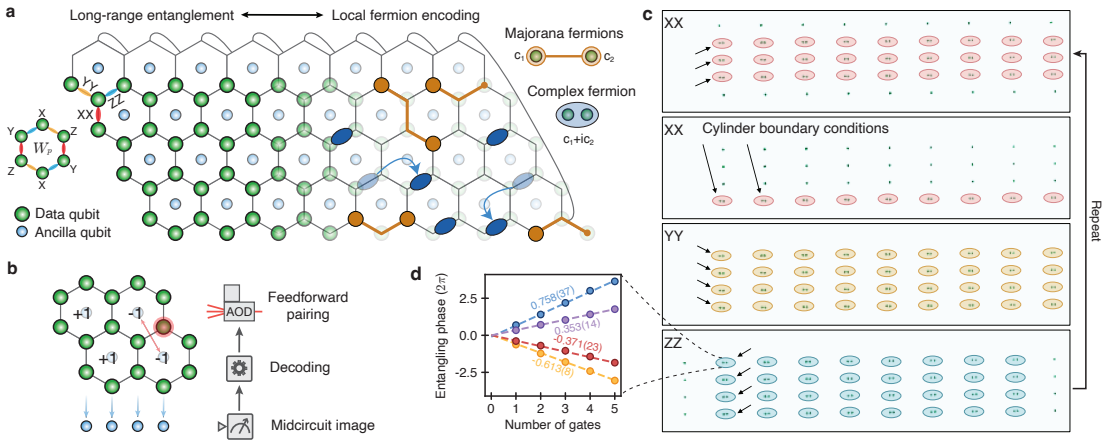


Figure 7.13: Digital quantum simulations with reconfigurable atom arrays. a, The honeycomb lattice used in this work with 104 total atomic qubits and periodic boundary conditions along the shorter direction, forming a cylinder. The qubits are encoded in ^{87}Rb atoms and entangling gates are realized through excitation to interacting Rydberg states. To encode fermion statistics, we prepare a long-range entangled state characterized by hexagonal plaquette operators $W_p = \mathrm{X}_1\mathrm{Z}_2\mathrm{Y}_3\mathrm{X}_4\mathrm{Z}_5\mathrm{Y}_6$ that commute with the unitary evolution and are therefore conserved. The encoded Majorana fermions live on vertices, at the ends of operator strings, and conventional (complex) fermions are formed by combining two Majoranas along a chosen link orientation. **b,** The long-range entangled state is prepared using mid-circuit measurement of ancilla qubits, and plaquettes are deterministically flipped to be +1 using conditional single-qubit gates (red circle). **c,** The Floquet evolution cycle consists of atom reconfiguration interspersed with tunable entangling gates and global basis changes. **d,** Accumulated entangling phase upon repeated application of $\exp(i\pi\theta[\mathrm{Z}\otimes\mathrm{Z}]/4)$ gates with a variable angle θ realized through fast, parameterized laser pulses. The values show extracted values of θ, and error bars represent one standard deviation.

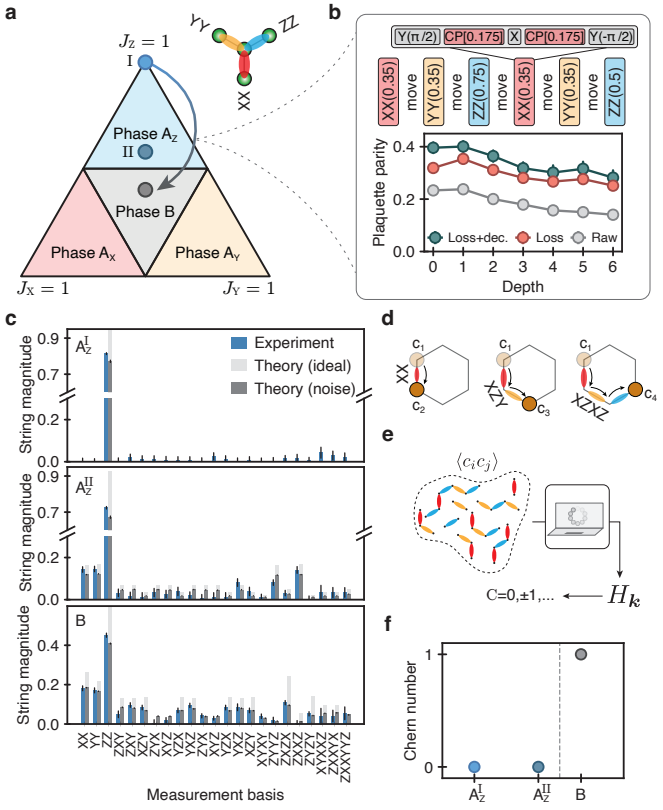


Figure 7.14: Kitaev model on a honeycomb lattice. a, The honeycomb model, consisting of anisotropic spin interactions along the three directions of the honeycomb lattice ($J_X+J_Y+J_Z=1$), exhibits topological order with three Abelian (A) phases and a single non-Abelian phase B. The toric code state is the fixed-point state A_Z^I of the Abelian A_Z phase ($J_Z=1$). **b,** Starting from this initial state, we prepare different states on the phase diagram. Top: numerically optimized sequence of two-qubit gates used to prepare the non-Abelian phase B. Each circuit layer includes CPHASE (CP) gates and global single-qubit rotations. Bottom: plaquette parity during the evolution, with postselection based on atom loss and decoding. **c,** Pauli strings of different lengths measured on the three studied states and averaged over the bulk of the system. Error bars represent 68% confidence intervals. **d,** In the fermion representation, the link operators are proportional to nearest-neighbor Majorana hoppings, $K_{ij}^{X/Y/Z} \propto ic_ic_j$. Longer Pauli strings constructed from their products result in longer-range hopping operators. **e,** The free-fermion parent Hamiltonian $\mathcal{H}_k$ can be reconstructed from measured two-point Majorana correlations. **f,** The Chern number C is evaluated using the learned parent Hamiltonians of the string distributions in **c**, resulting in C=0 in phase A and C=1 in phase B; the robustness of this procedure is explored in Methods. The non-Abelian phase B is characterized by the underlying topological order and an odd Chern number.

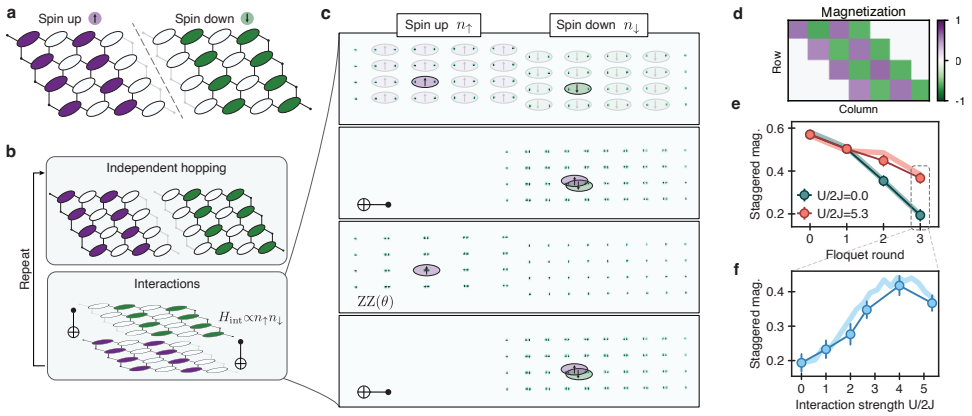


Figure 7.15: Two-component strongly interacting fermions. a, Two independent species of fermions, labeled "spin-up" and "spin-down", are realized by dividing the system into two halves. The XX and YY links across the boundary are turned off to facilitate independent hopping dynamics; this system realizes a 4x4 lattice with two fermion species. **b,** Floquet evolution consists of independent hopping within the two halves followed by density-density interactions between the two spin states. **c,** The interactions are realized through a parallel entangling operation connecting the two halves of the system. Two parallel CNOT gates propagate a ZZ(θ) operator, applied to the spin-up half, into the interaction term. **d,** The initial state is a checkerboard pattern, staggered between the two fermion spin states. **e,** Staggered magnetization dynamics during a quench with and without the interactions. **f,** Final staggered magnetization as a function of interaction strength. Increasing interactions suppress the spin mixing dynamics. The solid lines are the result of noisy numerical simulations of the circuit (Methods). Error bars represent 68% confidence intervals.

References

[1] Aaronson, S. & Gottesman, D. (2004). Improved simulation of stabilizer circuits. *Physical Review A*, 70(5), 052328.

[2] Acharya, R., Aghababaie-Beni, L., Aleiner, I., Andersen, T. I., Ansmann, M., Arute, F., Arya, K., Asfaw, A., Astrakhantsev, N., Atalaya, J., Babbush, R., Bacon, D., Ballard, B., Bardin, J. C., Bausch, J., Bengtsson, A., Bilmes, A., Blackwell, S., Boixo, S., Bortoli, G., Bourassa, A., Bovaird, J., Brill, L., Broughton, M., Browne, D. A., Buchea, B., Buckley, B. B., Buell, D. A., Burger, T., Burkett, B., Bushnell, N., Cabrera, A., Campero, J., Chang, H.-S., Chen, Y., Chen, Z., Chiaro, B., Chik, D., Chou, C., Claes, J., Cleland, A. Y., Cogan, J., Collins, R., Conner, P., Courtney, W., Crook, A. L., Curtin, B., Das, S., Davies, A., De Lorenzo, L., Debroy, D. M., Demura, S., Devoret, M., Di Paolo, A., Donohoe, P., Drozdov, I., Dunsworth, A., Earle, C., Edlich, T., Eickbusch, A., Elbag, A. M., Elzouka, M., Erickson, C., Faoro, L., Farhi, E., Ferreira, V. S., Burgos, L. F., Forati, E., Fowler, A. G., Foxen, B., Ganjam, S., Garcia, G., Gasca, R., Genois, É., Giang, W., Gidney, C., Gilboa, D., Gosula, R., Dau, A. G., Graumann, D., Greene, A., Gross, J. A., Habegger, S., Hall, J., Hamilton, M. C., Hansen, M., Harrigan, M. P., Harrington, S. D., Heras, F. J. H., Heslin, S., Heu, P., Hig-

gott, O., Hill, G., Hilton, J., Holland, G., Hong, S., Huang, H.-Y., Huff, A., Huggins, W. J., Ioffe, L. B., Isakov, S. V., Iveland, J., Jeffrey, E., Jiang, Z., Jones, C., Jordan, S., Joshi, C., Juhas, P., Kafri, D., Kang, H., Karamlou, A. H., Kechedzhi, K., Kelly, J., Khaire, T., Khattar, T., Khezri, M., Kim, S., Klimov, P. V., Klots, A. R., Kobrin, B., Kohli, P., Korotkov, A. N., Kostritsa, F., Kothari, R., Kozlovskii, B., Kreikebaum, J. M., Kurilovich, V. D., Lacroix, N., Landhuis, D., Lange-Dei, T., Langley, B. W., Laptev, P., Lau, K.-M., Guevel, L. L., Ledford, J., Lee, K., Lensky, Y. D., Leon, S., Lester, B. J., Li, W. Y., Li, Y., Lill, A. T., Liu, W., Livingston, W. P., Locharla, A., Lucero, E., Lundahl, D., Lunt, A., Madhuk, S., Malone, F. D., Maloney, A., Mandrá, S., Martin, L. S., Martin, S., Martin, O., Maxfield, C., McClean, J. R., McEwen, M., Meeks, S., Megrant, A., Mi, X., Miao, K. C., Mieszala, A., Molavi, R., Molina, S., Montazeri, S., Morvan, A., Movassagh, R., Mruczkiewicz, W., Naaman, O., Neeley, M., Neill, C., Nersisyan, A., Neven, H., Newman, M., Ng, J. H., Nguyen, A., Nguyen, M., Ni, C.-H., O'Brien, T. E., Oliver, W. D., Opremcak, A., Ottosson, K., Petukhov, A., Pizzuto, A., Platt, J., Potter, R., Pritchard, O., Pryadko, L. P., Quintana, C., Ramachandran, G., Reagor, M. J., Rhodes, D. M., Roberts, G., Rosenberg, E., Rosenfeld, E., Roushan, P., Rubin, N. C., Saei, N., Sank, D., Sankaragomathi, K., Satzinger, K. J., Schurkus, H. F., Schuster, C., Senior, A. W., Shearn, M. J., Shorter, A., Shutty, N., Shvarts, V., Singh, S., Sivak, V., Skruzny, J., Small, S., Smelyanskiy, V., Smith, W. C., Somma, R. D., Springer, S., Sterling, G., Strain, D., Suchard, J., Szasz, A., Sztein, A., Thor, D., Torres, A., Torunbalci, M. M., Vaishnav, A., Vargas, J., Vdovichev, S., Vidal, G., Villalonga, B., Heidweiller, C. V., Waltman, S., Wang, S. X., Ware, B., Weber, K., White, T., Wong, K., Woo, B. W. K., Xing, C., Yao, Z. J., Yeh, P., Ying, B., Yoo, J., Yosri, N., Young, G., Zalcman, A., Zhang, Y., Zhu, N., & Zobrist, N. (2024). Quantum error correction below the surface code threshold. *Nature*, 638(8052), 920–926.

[3] Aharonov, D. & Ben-Or, M. (1999). FAULT-TOLERANT QUANTUM COMPUTA-

TION WITH CONSTANT ERROR RATE *. (pp. 1997).

[4] Andersen, T. I., Lensky, Y. D., Kechedzhi, K., Drozdov, I. K., Bengtsson, A., Hong, S., Morvan, A., Mi, X., Opremcak, A., Acharya, R., Allen, R., Ansmann, M., Arute, F., Arya, K., Asfaw, A., Atalaya, J., Babbush, R., Bacon, D., Bardin, J. C., Bortoli, G., Bourassa, A., Bovaird, J., Brill, L., Broughton, M., Buckley, B. B., Buell, D. A., Burger, T., Burkett, B., Bushnell, N., Chen, Z., Chiaro, B., Chik, D., Chou, C., Cogan, J., Collins, R., Conner, P., Courtney, W., Crook, A. L., Curtin, B., Debroy, D. M., Del Toro Barba, A., Demura, S., Dunsworth, A., Eppens, D., Erickson, C., Faoro, L., Farhi, E., Fatemi, R., Ferreira, V. S., Burgos, L. F., Forati, E., Fowler, A. G., Foxen, B., Giang, W., Gidney, C., Gilboa, D., Giustina, M., Gosula, R., Dau, A. G., Gross, J. A., Habegger, S., Hamilton, M. C., Hansen, M., Harrigan, M. P., Harrington, S. D., Heu, P., Hilton, J., Hoffmann, M. R., Huang, T., Huff, A., Huggins, W. J., Ioffe, L. B., Isakov, S. V., Iveland, J., Jeffrey, E., Jiang, Z., Jones, C., Juhas, P., Kafri, D., Khattar, T., Khezri, M., Kieferová, M., Kim, S., Kitaev, A., Klimov, P. V., Klots, A. R., Korotkov, A. N., Kostritsa, F., Kreikebaum, J. M., Landhuis, D., Laptev, P., Lau, K. M., Laws, L., Lee, J., Lee, K. W., Lester, B. J., Lill, A. T., Liu, W., Locharla, A., Lucero, E., Malone, F. D., Martin, O., McClean, J. R., McCourt, T., McEwen, M., Miao, K. C., Mieszala, A., Mohseni, M., Montazeri, S., Mount, E., Movassagh, R., Mruczkiewicz, W., Naaman, O., Neeley, M., Neill, C., Nersisyan, A., Newman, M., Ng, J. H., Nguyen, A., Nguyen, M., Niu, M. Y., O'Brien, T. E., Omonije, S., Petukhov, A., Potter, R., Pryadko, L. P., Quintana, C., Rocque, C., Rubin, N. C., Saei, N., Sank, D., Sankaragomathi, K., Satzinger, K. J., Schurkus, H. F., Schuster, C., Shearn, M. J., Shorter, A., Shutty, N., Shvarts, V., Skruzny, J., Smith, W. C., Somma, R., Sterling, G., Strain, D., Szalay, M., Torres, A., Vidal, G., Villalonga, B., Heidweiller, C. V., White, T., Woo, B. W., Xing, C., Yao, Z. J., Yeh, P., Yoo, J., Young, G., Zalcman, A., Zhang, Y., Zhu, N., Zobrist, N., Neven, H., Boixo, S., Megrant, A., Kelly, J., Chen, Y., Smelyanskiy,

V., Kim, E. A., Aleiner, I., & Roushan, P. (2023). Non-Abelian braiding of graph vertices in a superconducting processor. *Nature*, 618(7964), 264–269.

[5] Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., Biswas, R., Boixo, S., Brandao, F. G., Buell, D. A., Burkett, B., Chen, Y., Chen, Z., Chiaro, B., Collins, R., Courtney, W., Dunsworth, A., Farhi, E., Foxen, B., Fowler, A., Gidney, C., Giustina, M., Graff, R., Guerin, K., Habegger, S., Harrigan, M. P., Hartmann, M. J., Ho, A., Hoffmann, M., Huang, T., Humble, T. S., Isakov, S. V., Jeffrey, E., Jiang, Z., Kafri, D., Kechedzhi, K., Kelly, J., Klimov, P. V., Knysh, S., Korotkov, A., Kostritsa, F., Landhuis, D., Lindmark, M., Lucero, E., Lyakh, D., Mandrà, S., McClean, J. R., McEwen, M., Megrant, A., Mi, X., Michielsen, K., Mohseni, M., Mutus, J., Naaman, O., Neeley, M., Neill, C., Niu, M. Y., Ostby, E., Petukhov, A., Platt, J. C., Quintana, C., Rieffel, E. G., Roushan, P., Rubin, N. C., Sank, D., Satzinger, K. J., Smelyanskiy, V., Sung, K. J., Trevithick, M. D., Vainsencher, A., Villalonga, B., White, T., Yao, Z. J., Yeh, P., Zalcman, A., Neven, H., & Martinis, J. M. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), 505–510.

[6] Ashkin, A., Dziedzic, J. M., Bjorkholm, J. E., & Chu, S. (1986). Observation of a single-beam gradient force optical trap for dielectric particles. *Optics Letters, Vol. 11, Issue 5, pp. 288-290*, 11(5), 288–290.

[7] Baranes, G., Cain, M., Pablo Bonilla Ataides, J., Bluvstein, D., Sinclair, J., Vuleti´cvuleti´c, V., Zhou, H., & Lukin, M. D. (2025). Leveraging Atom Loss Errors in Fault Tolerant Quantum Algorithms.

[8] Barends, R., Kelly, J., Megrant, A., Veitia, A., Sank, D., Jeffrey, E., White, T. C., Mutus, J., Fowler, A. G., Campbell, B., Chen, Y., Chen, Z., Chiaro, B., Dunsworth, A., Neill, C.,

O'Malley, P., Roushan, P., Vainsencher, A., Wenner, J., Korotkov, A. N., Cleland, A. N., & Martinis, J. M. (2014). Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497), 500–503.

[9] Barnes, K., Battaglino, P., Bloom, B. J., Cassella, K., Coxe, R., Crisosto, N., King, J. P., Kondov, S. S., Kotru, K., Larsen, S. C., Lauigan, J., Lester, B. J., McDonald, M., Megidish, E., Narayanaswami, S., Nishiguchi, C., Notermans, R., Peng, L. S., Ryou, A., Wu, T. Y., & Yarwood, M. (2022). Assembly and coherent control of a register of nuclear spin qubits. *Nature Communications*, 13(1), 2779.

[10] Barredo, D., De Léséleuc, S., Lienhard, V., Lahaye, T., & Browaeys, A. (2016). An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays. *Science*, 354(6315), 1021–1023.

[11] Benioff, P. (1980). The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *Journal of Statistical Physics*, 22(5), 563–591.

[12] Bennett, C. H. (1982). The thermodynamics of computation-a review. *International Journal of Theoretical Physics*, 21(12), 905–940.

[13] Bernien, H., Schwartz, S., Keesling, A., Levine, H., Omran, A., Pichler, H., Choi, S., Zibrov, A. S., Endres, M., Greiner, M., Vuletic, V., & Lukin, M. D. (2017). Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682), 579–584.

[14] Beugnon, J., Tuchendler, C., Marion, H., Gaëtan, A., Miroshnychenko, Y., Sortais, Y. R. P., Lance, A. M., Jones, M. P. A., Messin, G., Browaeys, A., & Grangier, P. (2007). Two-dimensional transport and transfer of a single atomic qubit in optical tweezers. *Nature Physics*, 3(10), 696–699.

[15] Beverland, M. E., Kubica, A., & Svore, K. M. (2021). Cost of Universality: A Comparative Study of the Overhead of State Distillation and Code Switching with Color Codes. *PRX Quantum*, 2(2), 020341.

[16] Blum, L., Shub, M., & Smale, S. (1988). On a theory of computation over the real numbers; NP completeness, recursive functions and universal machines. *Annual Symposium on Foundations of Computer Science (Proceedings)*, (pp. 387–397).

[17] Bluvstein, D., Evered, S. J., Geim, A. A., Li, S. H., Zhou, H., Manovitz, T., Ebadi, S., Cain, M., Kalinowski, M., Hangleiter, D., Bonilla Ataides, J. P., Maskara, N., Cong, I., Gao, X., Sales Rodriguez, P., Karolyshyn, T., Semeghini, G., Gullans, M. J., Greiner, M., Vuletić, V., & Lukin, M. D. (2024). Logical quantum processor based on reconfigurable atom arrays. *Nature*, 626, 58–65.

[18] Bluvstein, D., Geim, A. A., & Et al, . (2025). Physical mechanisms for a fault-tolerant quantum architecture.

[19] Bluvstein, D., Levine, H., Semeghini, G., Wang, T. T., Ebadi, S., Kalinowski, M., Keesling, A., Maskara, N., Pichler, H., Greiner, M., Vuletić, V., & Lukin, M. D. (2022). A quantum processor based on coherent transport of entangled atom arrays. *Nature*, 604(7906), 451–456.

[20] Bluvstein, D., Omran, A., Levine, H., Keesling, A., Semeghini, G., Ebadi, S., Wang, T. T., Michailidis, A. A., Maskara, N., Ho, W. W., Choi, S., Serbyn, M., Greiner, M., Vuletić, V., & Lukin, M. D. (2021). Controlling quantum many-body dynamics in driven Rydberg atom arrays. *Science*, 371(6536), 1355–1359.

[21] Bombín, H. (2015a). Gauge color codes: optimal transversal gates and gauge fixing in topological stabilizer codes. *New Journal of Physics*, 17(8), 083002.

[22] Bombín, H. (2015b). Single-Shot Fault-Tolerant Quantum Error Correction. *Physical Review X*, 5(3), 031043.

[23] Bravyi, S., Browne, D., Calpin, P., Campbell, E., Gosset, D., & Howard, M. (2019). Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum*, 3, 181.

[24] Bravyi, S., Cross, A. W., Gambetta, J. M., Maslov, D., Rall, P., & Yoder, T. J. (2023). High-threshold and low-overhead fault-tolerant quantum memory.

[25] Bravyi, S., Englbrecht, M., König, R., & Peard, N. (2018). Correcting coherent errors with surface codes. *npj Quantum Information*, 4(1), 55.

[26] Bremner, M. J., Montanaro, A., & Shepherd, D. J. (2016). Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations. *Physical Review Letters*, 117(8), 080501.

[27] Brennen, G. K., Caves, C. M., Jessen, P. S., & Deutsch, I. H. (1999). Quantum Logic Gates in Optical Lattices. *Physical Review Letters*, 82(5), 1060.

[28] Breuckmann, N. P. & Eberhardt, J. N. (2021). Quantum Low-Density Parity-Check Codes. *PRX Quantum*, 2(4), 040101.

[29] Brown, B. J. (2020). A fault-tolerant non-clifford gate for the surface code in two dimensions. *Science Advances*, 6(21).

[30] Brown, K. R., Wilson, A. C., Colombe, Y., Ospelkaus, C., Meier, A. M., Knill, E., Leibfried, D., & Wineland, D. J. (2011). Single-qubit-gate error below 10-4 in a trapped ion. *Physical Review A*, 84(3), 030303.

[31] Brown, M. O., Thiele, T., Kiehl, C., Hsu, T. W., & Regal, C. A. (2019). Gray-Molasses Optical-Tweezer Loading: Controlling Collisions for Scaling Atom-Array Assembly. *Physical Review X*, 9(2), 011057.

[32] Cain, M., Zhao, C., Zhou, H., Meister, N., Pablo, J., Ataides, B., Jaffe, A., Bluvstein, D., & Lukin, M. D. (2024). Correlated decoding of logical algorithms with transversal gates.

[33] Calderbank, A. R. & Shor, P. W. (1996). Good quantum error-correcting codes exist. *Physical Review A*, 54(2), 1098–1105.

[34] Campbell, E. T. (2016). The smallest interesting colour code.

[35] Cao, A., Eckner, W. J., Yelin, T. L., Young, A. W., Jandura, S., Yan, L., Kim, K., Pupillo, G., Ye, J., Oppong, N. D., & Kaufman, A. M. (2024). Multi-qubit gates and Schr\"odinger cat states in an optical clock.

[36] Chow, C. H., Ng, B. L., Prakash, V., & Kurtsiefer, C. (2024). Fano resonance in excitation spectroscopy and cooling of an optically trapped single atom. *Physical Review Research*, 6(2), 023154.

[37] Cong, I., Levine, H., Keesling, A., Bluvstein, D., Wang, S. T., & Lukin, M. D. (2022). Hardware-Efficient, Fault-Tolerant Quantum Computation with Rydberg Atoms. *Physical Review X*, 12(2), 021049.

[38] Cummins, H. K., Llewellyn, G., & Jones, J. A. (2003). Tackling systematic errors in quantum logic gates with composite rotations. *Physical Review A*, 67(4), 042308.

[39] Daley, A. J., Pichler, H., Schachenmayer, J., & Zoller, P. (2012). Measuring Entanglement Growth in Quench Dynamics of Bosons in an Optical Lattice. *Physical Review Letters*, 109(2), 020505.

[40] Dalibard, J. & Cohen-Tannoudji, C. (1989). Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *JOSA B*, 6(11), 2023–2045.

[41] Dawson, C. M. & Nielsen, M. A. (2005). The Solovay-Kitaev algorithm. *Quantum Information and Computation*, 6(1), 081–095.

[42] De Léséleuc, S., Barredo, D., Lienhard, V., Browaeys, A., & Lahaye, T. (2018). Analysis of imperfections in the coherent optical excitation of single atoms to Rydberg states. *Physical Review A*, 97(5), 053803.

[43] Deist, E., Lu, Y. H., Ho, J., Pasha, M. K., Zeiher, J., Yan, Z., & Stamper-Kurn, D. M. (2022). Mid-Circuit Cavity Measurement in a Neutral Atom Array. *Physical Review Letters*, 129(20), 203602.

[44] Dennis, E., Kitaev, A., Landahl, A., & Preskill, J. (2002). Topological quantum memory. *Journal of Mathematical Physics*, 43(9), 4452–4505.

[45] Deutsch, D. (1985). QUANTUM THEORY, THE CHURCH-TURING PRINCIPLE AND THE UNIVERSAL QUANTUM COMPUTER. *Proceedings of The Royal Society of London, Series A: Mathematical and Physical Sciences*, 400(1818), 97–117.

[46] Dordević, T., Samutpraphoot, P., Ocola, P. L., Bernien, H., Grinkemeyer, B., Dimitrova, I., Vuletić, V., & Lukin, M. D. (2021). Entanglement transport and a nanophotonic interface for atoms in optical tweezers. *Science*, 373(6562), 1511–1514.

[47] Eastin, B. & Knill, E. (2009). Restrictions on transversal encoded quantum gate sets. *Physical Review Letters*, 102(11), 110502.

[48] Ebadi, S., Keesling, A., Cain, M., Wang, T. T., Levine, H., Bluvstein, D., Semeghini, G., Omran, A., Liu, J. G., Samajdar, R., Luo, X. Z., Nash, B., Gao, X., Barak, B., Farhi, E., Sachdev,

S., Gemelke, N., Zhou, L., Choi, S., Pichler, H., Wang, S. T., Greiner, M., Vuletić, V., & Lukin, M. D. (2022). Quantum optimization of maximum independent set using Rydberg atom arrays. *Science*, 376(6598), 1209–1215.

[49] Ebadi, S., Wang, T. T., Levine, H., Keesling, A., Semeghini, G., Omran, A., Bluvstein, D., Samajdar, R., Pichler, H., Ho, W. W., Choi, S., Sachdev, S., Greiner, M., Vuletić, V., & Lukin, M. D. (2021). Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature*, 595(7866), 227–232.

[50] Egan, L., Debroy, D. M., Noel, C., Risinger, A., Zhu, D., Biswas, D., Newman, M., Li, M., Brown, K. R., Cetina, M., & Monroe, C. (2021). Fault-tolerant control of an error-corrected qubit. *Nature*, 598, 281–286.

[51] Endres, M., Bernien, H., Keesling, A., Levine, H., Anschuetz, E. R., Krajenbrink, A., Senko, C., Vuletic, V., Greiner, M., & Lukin, M. D. (2016). Atom-by-atom assembly of defect-free one-dimensional cold atom arrays. *Science*, 354(6315), 1024–1027.

[52] Evered, S. J., Bluvstein, D., Kalinowski, M., Ebadi, S., Manovitz, T., Zhou, H., Li, S. H., Geim, A. A., Wang, T. T., Maskara, N., Levine, H., Semeghini, G., Greiner, M., Vuletić, V., Lukin, M. D., John, H., Paulson, A., & Kalinowski, M. **?**. (2023a). High-fidelity parallel entangling gates on a neutral-atom quantum computer. *Nature*, 622, 268–272.

[53] Evered, S. J., Bluvstein, D., Kalinowski, M., Ebadi, S., Manovitz, T., Zhou, H., Li, S. H., Geim, A. A., Wang, T. T., Maskara, N., Levine, H., Semeghini, G., Greiner, M., Vuleti´cvuleti´c, V., & Lukin, M. D. (2023b). High-fidelity parallel entangling gates on a neutral atom quantum computer.

[54] Evered, S. J., Kalinowski, M., Geim, A. A., Manovitz, T., Bluvstein, D., Li, S. H., Maskara, N., Zhou, H., Ebadi, S., Xu, M., Campo, J., Cain, M., Ostermann, S., Yelin, S. F., Sachdev,

S., Greiner, M., Vuleti´cvuleti´c, V., & Lukin, M. D. (2025). Probing topological matter and fermion dynamics on a neutral-atom quantum computer.

[55] Feynman, R. P. (1982). Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7).

[56] Flammia, S. T. & Liu, Y. K. (2011). Direct fidelity estimation from few Pauli measurements. *Physical Review Letters*, 106(23), 230501.

[57] Fowler, A. G., Mariantoni, M., Martinis, J. M., & Cleland, A. N. (2012). Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3), 032324.

[58] Fowler, A. G., Wang, D. S., Hill, C. D., Ladd, T. D., Meter, R. V., & Hollenberg, L. C. L. (2010). Surface Code Quantum Communication. *Physical Review Letters*, 104(18), 180503.

[59] Gidney, C. & Ekerå, M. (2021). How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum*, 5, 433.

[60] Goto, H. (2016). Minimizing resource overheads for fault-tolerant preparation of encoded states of the Steane code. *Scientific Reports*, 6, 19578.

[61] Gottesman, D. (1997). Stabilizer Codes and Quantum Error Correction.

[62] Graham, T. M., Phuttitarn, L., Chinnarasu, R., Song, Y., Poole, C., Jooya, K., Scott, J., Scott, A., Eichler, P., & Saffman, M. (2023). Mid-circuit measurements on a neutral atom quantum processor.

[63] Graham, T. M., Song, Y., Scott, J., Poole, C., Phuttitarn, L., Jooya, K., Eichler, P., Jiang, X., Marra, A., Grinkemeyer, B., Kwon, M., Ebert, M., Cherek, J., Lichtman, M. T., Gillette, M., Gilbert, J., Bowman, D., Ballance, T., Campbell, C., Dahl, E. D., Crawford, O., Blunt, N. S.,

Rogers, B., Noel, T., & Saffman, M. (2022). Multi-qubit entanglement and algorithms on a neutral-atom quantum computer. *Nature*, 604(7906), 457–462.

[64] Gullion, T., Baker, D. B., & Conradi, M. S. (1990). New, compensated Carr-Purcell sequences. *Journal of Magnetic Resonance (1969)*, 89(3), 479–484.

[65] Gyger, F., Ammenwerth, M., Tao, R., Timme, H., Snigirev, S., Bloch, I., & Zeiher, J. (2024). Continuous operation of large-scale atom arrays in optical lattices. *Physical Review Research*, 6(3), 033104.

[66] Hangleiter, D. & Gullans, M. J. (2023). Bell sampling from quantum circuits.

[67] Haroche, S. & Raimond, J. M. (1996). Quantum Computing: Dream or Nightmare? *Physics Today*, 49(8), 51–52.

[68] Haug, T. & Kim, M. S. (2023). Scalable Measures of Magic Resource for Quantum Computers. *PRX Quantum*, 4(1), 010301.

[69] Hein, M., Dür, W., Eisert, J., Raussendorf, R., den Nest, M. V., & Briegel, H. J. (2006). Entanglement in Graph States and its Applications. *Proceedings of the International School of Physics "Enrico Fermi"*, 162, 115–218.

[70] Ho, W. W., Choi, S., Pichler, H., & Lukin, M. D. (2019). Periodic Orbits, Entanglement, and Quantum Many-Body Scars in Constrained Models: Matrix Product State Approach. *Physical Review Letters*, 122(4), 040603.

[Honciuc Menendez et al.] Honciuc Menendez, D., Ray, A., & Vasmer, M. Implementing fault-tolerant non-Clifford gates using the [[8,3,2]] color code.

[72] Horsman, C., Fowler, A. G., Devitt, S., & Meter, R. V. (2012). Surface code quantum computing by lattice surgery. *New Journal of Physics*, 14(12), 123011.

[73] Howard, M., Wallman, J., Veitch, V., & Emerson, J. (2014). Contextuality supplies the 'magic' for quantum computation. *Nature*, 510(7505), 351–355.

[74] Hu, B., Sinclair, J., Bytyqi, E., Chong, M., Rudelis, A., Ramette, J., Vendeiro, Z., & Vuletić, V. (2024). Site-selective cavity readout and classical error correction of a 5-bit atomic register.

[75] Huang, H. Y., Broughton, M., Cotler, J., Chen, S., Li, J., Mohseni, M., Neven, H., Babbush, R., Kueng, R., Preskill, J., & McClean, J. R. (2022). Quantum advantage in learning from experiments. *Science*, 376(6598), 1182–1186.

[76] Huang, S. & Brown, K. R. (2021). Between Shor and Steane: A Unifying Construction for Measuring Error Syndromes. *Physical Review Letters*, 127(9).

[77] Hutzler, N. R., Liu, L. R., Yu, Y., & Ni, K. K. (2017). Eliminating light shifts for single atom trapping. *New Journal of Physics*, 19(2), 023007.

[78] Islam, R., Ma, R., Preiss, P. M., Eric Tai, M., Lukin, A., Rispoli, M., & Greiner, M. (2015). Measuring entanglement entropy in a quantum many-body system. *Nature*, 528(7580), 77–83.

[79] Jaksch, D., Cirac, J. I., Zoller, P., Rolston, S. L., Côté, R., & Lukin, M. D. (2000). Fast quantum gates for neutral atoms. *Physical Review Letters*, 85(10), 2208–2211.

[80] Jandura, S. & Pupillo, G. (2022). Time-Optimal Two- And Three-Qubit Gates for Rydberg Atoms. *Quantum*, 6, 712.

[81] Jenkins, A., Lis, J. W., Senoo, A., McGrew, W. F., & Kaufman, A. M. (2022). Ytterbium Nuclear-Spin Qubits in an Optical Tweezer Array. *Physical Review X*, 12(2), 021027.

[82] Johnston, A. H. (2000). Scaling and Technology Issues for Soft Error Rates.

[83] Jones, M. P., Beugnon, J., Gaëtan, A., Zhang, J., Messin, G., Browaeys, A., & Grangier, P. (2007). Fast quantum state control of a single trapped neutral atom. *Physical Review A*, 75(4), 040301.

[84] Kim, Y., Eddins, A., Anand, S., Wei, K. X., van den Berg, E., Rosenblatt, S., Nayfeh, H., Wu, Y., Zaletel, M., Temme, K., & Kandala, A. (2023). Evidence for the utility of quantum computing before fault tolerance. *Nature*, 618(7965), 500–505.

[85] Kitaev, A. Y. (1997). Quantum computations: Algorithms and error correction. *Russian Mathematical Surveys*, 52(6), 1191–1249.

[86] Kitaev, A. Y. (2003). Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1), 2–30.

[87] Knill, E., Leibfried, D., Reichle, R., Britton, J., Blakestad, R. B., Jost, J. D., Langer, C., Ozeri, R., Seidelin, S., & Wineland, D. J. (2008). Randomized benchmarking of quantum gates. *Physical Review A*, 77(1), 012307.

[88] Kubica, A., Yoshida, B., & Pastawski, F. (2015). Unfolding the color code. *New Journal of Physics*, 17(8), 083026.

[89] Kuriyattil, S., Hashizume, T., Bentsen, G., & Daley, A. J. (2023). Onset of Scrambling as a Dynamical Transition in Tunable-Range Quantum Circuits. *PRX Quantum*, 4(3), 030325.

[90] Lampen, J., Nguyen, H., Li, L., Berman, P. R., & Kuzmich, A. (2018). Long-lived coherence between ground and Rydberg levels in a magic-wavelength lattice. *Physical Review A*, 98(3), 033411.

[91] Landauer, R. (1961). Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development*, 5(3), 183–191.

[92] Landauer, R. (1995). Is quantum mechanics useful? *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, 353(1703), 367–376.

[93] Landauer, R. (1996). The physical nature of information. *Physics Letters A*, 217(4-5), 188–193.

[94] Lengwenus, A., Kruse, J., Schlosser, M., Tichelmann, S., & Birkl, G. (2010). Coherent transport of atomic quantum states in a scalable shift register. *Physical Review Letters*, 105(17), 170502.

[95] Levine, H., Bluvstein, D., Keesling, A., Wang, T. T., Semeghini, G., Ebadi, S., Omran, A., Greiner, M., Vuletić, V., & Lukin, M. D. (2022). Dispersive optics for scalable Raman driving of hyperfine qubits. *Physical Review A*, 105, 032618.

[96] Levine, H., Keesling, A., Semeghini, G., Omran, A., Wang, T. T., Ebadi, S., Bernien, H., Greiner, M., Vuletić, V., Pichler, H., & Lukin, M. D. (2019). Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms. *Physical Review Letters*, 123(17), 170503.

[97] Lin, W.-H., Tan, D. B., & Cong, J. (2024). Reuse-Aware Compilation for Zoned Quantum Architectures Based on Neutral Atoms.

[98] Lis, J. W., Senoo, A., McGrew, W. F., Rönchen, F., Jenkins, A., & Kaufman, A. M. (2023). Mid-circuit operations using the omg-architecture in neutral atom arrays.

[99] Litinski, D., Psiquantum, N. N., & Alto, P. (2022). Active volume: An architecture for efficient fault-tolerant quantum computers with limited non-local connections.

[100] Ma, S., Liu, G., Peng, P., Zhang, B., Jandura, S., Claes, J., Burgers, A. P., Pupillo, G., Puri, S., & Thompson, J. D. (2023a). High-fidelity gates with mid-circuit erasure conversion in a metastable neutral atom qubit.

[101] Ma, S., Liu, G., Peng, P., Zhang, B., Jandura, S., Claes, J., Burgers, A. P., Pupillo, G., Puri, S., Thompson, J. D., & Peng, P. . (2023b). High-fidelity gates and mid-circuit erasure conversion in an atomic qubit. *Nature*, 622, 279–284.

[102] Manetsch, H. J., Nomura, G., Bataille, E., Leung, K. H., Lv, X., & Endres, M. (2024). A tweezer array with 6100 highly coherent atomic qubits.

[103] Manovitz, T., Li, S. H., Ebadi, S., Samajdar, R., Geim, A. A., Evered, S. J., Bluvstein, D., Zhou, H., Köylüoğlu, N. U., Feldmeier, J., Dolgirev, P. E., Maskara, N., Kalinowski, M., Sachdev, S., Huse, D. A., Greiner, M., Vuletić, V., & Lukin, M. D. (2025). Quantum coarsening and collective dynamics on a programmable quantum simulator. *Nature*, 638(8049), 86–92.

[104] Mezher, R., Ghalbouni, J., Dgheim, J., & Markham, D. (2020). Fault-tolerant quantum speedup from constant depth quantum circuits. *Physical Review Research*, 2(3), 033444.

[105] Mi, X., Roushan, P., Quintana, C., Mandrà, S., Marshall, J., Neill, C., Arute, F., Arya, K., Atalaya, J., Babbush, R., Bardin, J. C., Barends, R., Basso, J., Bengtsson, A., Boixo, S., Bourassa, A., Broughton, M., Buckley, B. B., Buell, D. A., Burkett, B., Bushnell, N., Chen, Z., Chiaro, B., Collins, R., Courtney, W., Demura, S., Derk, A. R., Dunsworth, A., Eppens, D., Erickson, C., Farhi, E., Fowler, A. G., Foxen, B., Gidney, C., Giustina, M., Gross, J. A., Harrigan, M. P., Harrington, S. D., Hilton, J., Ho, A., Hong, S., Huang, T., Huggins, W. J., Ioffe, L. B., Isakov, S. V., Jeffrey, E., Jiang, Z., Jones, C., Kafri, D., Kelly, J., Kim, S., Kitaev, A., Klimov, P. V., Korotkov, A. N., Kostritsa, F., Landhuis, D., Laptev, P., Lucero, E., Mar-

tin, O., McClean, J. R., McCourt, T., McEwen, M., Megrant, A., Miao, K. C., Mohseni, M., Montazeri, S., Mruczkiewicz, W., Mutus, J., Naaman, O., Neeley, M., Newman, M., Niu, M. Y., O'Brien, T. E., Opremcak, A., Ostby, E., Pato, B., Petukhov, A., Redd, N., Rubin, N. C., Sank, D., Satzinger, K. J., Shvarts, V., Strain, D., Szalay, M., Trevithick, M. D., Villalonga, B., White, T., Yao, Z. J., Yeh, P., Zalcman, A., Neven, H., Aleiner, I., Kechedzhi, K., Smelyanskiy, V., & Chen, Y. (2021). Information scrambling in quantum circuits. *Science*, 374(6574), 1479–1483.

[106] Monroe, C., Swann, W., Robinson, H., & Wieman, C. (1990). Very cold trapped atoms in a vapor cell. *Physical Review Letters*, 65(13), 1571.

[107] Norcia, M. A., Cairncross, W. B., Barnes, K., Battaglino, P., Brown, A., Brown, M. O., Cassella, K., Chen, C.-A., Coxe, R., Crow, D., Epstein, J., Griger, C., Jones, A. M. W., Kim, H., Kindem, J. M., King, J., Kondov, S. S., Kotru, K., Lauigan, J., Li, M., Lu, M., Megidish, E., Marjanovic, J., Mcdonald, M., Mittiga, T., Muniz, J. A., Narayanaswami, S., Nishiguchi, C., Notermans, R., Paule, T., Pawlak, K. A., Peng, L. S., Ryou, A., Smull, A., Stack, D., Stone, M., Sucich, A., Urbanek, M., Van De Veerdonk, R. J. M., Vendeiro, Z., Wilkason, T., Wu, T.-Y., Xie, X., Zhang, X., & Bloom, B. J. (2023). Mid-circuit qubit measurement and rearrangement in a 171 Yb atomic array.

[108] Norcia, M. A., Kim, H., Cairncross, W. B., Stone, M., Ryou, A., Jaffe, M., Brown, M. O., Barnes, K., Battaglino, P., Bohdanowicz, T. C., Brown, A., Cassella, K., Chen, C. A., Coxe, R., Crow, D., Epstein, J., Griger, C., Halperin, E., Hummel, F., Jones, A. M., Kindem, J. M., King, J., Kotru, K., Lauigan, J., Li, M., Lu, M., Megidish, E., Marjanovic, J., McDonald, M., Mittiga, T., Muniz, J. A., Narayanaswami, S., Nishiguchi, C., Paule, T., Pawlak, K. A., Peng, L. S., Pudenz, K. L., Rodríguez Pérez, D., Smull, A., Stack, D., Urbanek, M., Van De Veerdonk, R. J., Vendeiro, Z., Wadleigh, L., Wilkason, T., Wu, T. Y., Xie, X., Zalys-Geller, E.,

Zhang, X., & Bloom, B. J. (2024). Iterative Assembly of 171 Yb Atom Arrays with Cavity-Enhanced Optical Lattices. *PRX Quantum*, 5(3), 030316.

[109] Pagano, A., Weber, S., Jaschke, D., Pfau, T., Meinert, F., Montangero, S., & Büchler, H. P. (2022). Error budgeting for a controlled-phase gate with strontium-88 Rydberg atoms. *Physical Review Research*, 4(3), 033019.

[110] Paletta, L., Leverrier, A., Sarlette, A., Mirrahimi, M., & Vuillot, C. (2023). Robust sparse IQP sampling in constant depth.

[111] Patterson, D. A. & Hennessy, J. L. (2018). *Computer Organization and Design: The Hardware Software Interface - RISC-V Edition*.

[112] Postler, L., Heußen, S., Pogorelov, I., Rispler, M., Feldker, T., Meth, M., Marciniak, C. D., Stricker, R., Ringbauer, M., Blatt, R., Schindler, P., Müller, M., & Monz, T. (2022). Demonstration of fault-tolerant universal quantum gate operations. *Nature*, 605, 675–680.

[113] Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. *Quantum*, 2, 79.

[114] Quantum, G. (2023). Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614(7949), 676–681.

[115] Raab, E. L., Prentiss, M., Cable, A., Chu, S., & Pritchard, D. E. (1987). Trapping of Neutral Sodium Atoms with Radiation Pressure. *Physical Review Letters*, 59(23), 2631.

[116] Raussendorf, R. (2012). Key ideas in quantum error correction. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1975), 4541–4565.

[117] Raussendorf, R. & Briegel, H. J. (2001). A One-Way Quantum Computer. *Physical Review Letters*, 86(22), 5188.

[118] Rolston, S. L., Phillips, W. D., Walhout, M., & Dalibard, J. (1992). $\sigma+-\sigma-$ Optical molasses in a longitudinal magnetic field. *JOSA B*, 9(11), 1997–2007.

[119] Ryan-Anderson, C., Brown, N. C., Allman, M. S., Arkin, B., Asa-Attuah, G., Baldwin, C., Berg, J., Bohnet, J. G., Braxton, S., Burdick, N., Campora, J. P., Chernoguzov, A., Esposito, J., Evans, B., Francois, D., Gaebler, J. P., Gatterman, T. M., Gerber, J., Gilmore, K., Gresh, D., Hall, A., Hankin, A., Hostetter, J., Lucchetti, D., Mayer, K., Myers, J., Neyenhuis, B., Santiago, J., Sedlacek, J., Skripka, T., Slattery, A., Stutz, R. P., Tait, J., Tobey, R., Vittorini, G., Walker, J., & Hayes, D. (2022). Implementing Fault-tolerant Entangling Gates on the Five-qubit Code and the Color Code.

[120] Saffman, M., Walker, T. G., & Mølmer, K. (2010). Quantum information with Rydberg atoms. *Reviews of Modern Physics*, 82(3), 2313–2363.

[121] Schlosser, M., Tichelmann, S., Kruse, J., & Birkl, G. (2011). Scalable architecture for quantum information processing with atoms in optical micro-structures. *Quantum Information Processing*, 10(6), 907–924.

[122] Scholl, P., Schuler, M., Williams, H. J., Eberharter, A. A., Barredo, D., Schymik, K. N., Lienhard, V., Henry, L. P., Lang, T. C., Lahaye, T., Läuchli, A. M., & Browaeys, A. (2021). Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms. *Nature*, 595(7866), 233–238.

[123] Scholl, P., Shaw, A. L., Bing-Shiun Tsai, R., Finkelstein, R., Choi, J., & Endres, M. (2023). Erasure conversion in a high-fidelity Rydberg quantum simulator. *Nature*, 622, 273–278.

[124] Schwartz, I., Cogan, D., Schmidgall, E. R., Don, Y., Gantz, L., Kenneth, O., Lindner, N. H., & Gershoni, D. (2016). Deterministic generation of a cluster state of entangled photons. *Science*, 354(6311), 434–437.

[125] Sekino, Y. & Susskind, L. (2008). Fast scramblers. *Journal of High Energy Physics*, 2008(10).

[126] Self, C. N., Benedetti, M., & Amaro, D. (2022). Protecting Expressive Circuits with a Quantum Error Detection Code.

[127] Semeghini, G., Levine, H., Keesling, A., Ebadi, S., Wang, T. T., Bluvstein, D., Verresen, R., Pichler, H., Kalinowski, M., Samajdar, R., Omran, A., Sachdev, S., Vishwanath, A., Greiner, M., Vuletić, V., & Lukin, M. D. (2021a). Probing topological spin liquids on a programmable quantum simulator. *Science*, 374(6572), 1242–1247.

[128] Semeghini, G., Levine, H., Keesling, A., Ebadi, S., Wang, T. T., Bluvstein, D., Verresen, R., Pichler, H., Kalinowski, M., Samajdar, R., Omran, A., Sachdev, S., Vishwanath, A., Greiner, M., Vuletic, V., & Lukin, M. D. (2021b). Probing Topological Spin Liquids on a Programmable Quantum Simulator.

[129] Shea, M. E., Baker, P. M., Joseph, J. A., Kim, J., & Gauthier, D. J. (2020). Submillisecond, nondestructive, time-resolved quantum-state readout of a single, trapped neutral atom. *Physical Review A*, 102(5), 053101.

[130] Sheng, C., He, X., Xu, P., Guo, R., Wang, K., Xiong, Z., Liu, M., Wang, J., & Zhan, M. (2018). High-Fidelity Single-Qubit Gates on Neutral Atoms in a Two-Dimensional Magic-Intensity Optical Dipole Trap Array. *Physical Review Letters*, 121(24), 240501.

[131] Shor, P. W. (1995). Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52(4), R2493.

[132] Shor, P. W. (1996). Fault-tolerant quantum computation. In *Annual Symposium on Foundations of Computer Science - Proceedings* (pp. 56–65).: IEEE.

[133] Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5).

[134] Sinclair, J., Ramette, J., Grinkemeyer, B., Bluvstein, D., Lukin, M. D., & Vuletić, V. (2025). Fault-tolerant optical interconnects for neutral-atom arrays. *Physical Review Research*, 7(1), 013313.

[135] Singh, K., Bradley, C. E., Anand, S., Ramesh, V., White, R., & Bernien, H. (2023). Mid-circuit correction of correlated phase errors using an array of spectator qubits. *Science*, 380(6651), 1265–1269.

[136] Steane, A. (1996). Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 452(1954), 2551–2577.

[137] Steane, A. M. (1997). Active stabilization, quantum computation, and quantum state synthesis. *Physical Review Letters*, 78(11), 2252.

[138] Tan, D. B., Bluvstein, D., Lukin, M. D., & Cong, J. (2024a). Compiling Quantum Circuits for Dynamically Field-Programmable Neutral Atoms Array Processors. *Quantum*, 8, 1281.

[139] Tan, D. B., Lin, W.-H., & Cong, J. (2024b). Compilation for Dynamically Field-Programmable Qubit Arrays with Efficient and Provably Near-Optimal Scheduling. *Proceedings of the 30th Asia and South Pacific Design Automation Conference*, (pp. 921–929).

[140] Tao, R., Ammenwerth, M., Gyger, F., Bloch, I., & Zeiher, J. (2023). High-fidelity detection of large-scale atom arrays in an optical lattice.

[141] Tóth, G. & Gühne, O. (2005). Entanglement detection in the stabilizer formalism. *Physical Review A*, 72(2), 022340.

[142] Tsai, R. B. S., Sun, X., Shaw, A. L., Finkelstein, R., & Endres, M. (2025). Benchmarking and Fidelity Response Theory of High-Fidelity Rydberg Entangling Gates. *PRX Quantum*, 6(1), 010331.

[143] Turner, C. J., Michailidis, A. A., Abanin, D. A., Serbyn, M., & Papić, Z. (2018). Weak ergodicity breaking from quantum many-body scars. *Nature Physics*, 14(7), 745–749.

[144] Unruh, W. G. (1995). Maintaining coherence in quantum computers. *Physical Review A*, 51(2), 992.

[145] Urban, E., Johnson, T. A., Henage, T., Isenhower, L., Yavuz, D. D., Walker, T. G., & Saffman, M. (2009). Observation of Rydberg blockade between two atoms. *Nature Physics*, 5(2), 110–114.

[146] Vasmer, M. & Kubica, A. (2022). Morphing Quantum Codes. *Physical Review Applied*, 10, 030319.

[Wang et al.] Wang, Y., Simsek, S., Gatterman, T. M., Gerber, J. A., Gilmore, K., Gresh, D., Hewitt, N., Horst, C. V., Matheny, M., Mengle, T., Neyenhuis, B., & Criger, B. Fault-Tolerant One-Bit Addition with the Smallest Interesting Colour Code.

[148] Wimperis, S. (1994). Broadband, Narrowband, and Passband Composite Pulses for Use in Advanced NMR Experiments. *Journal of Magnetic Resonance, Series A*, 109(2), 221–231.

[149] Wu, T.-Y., Kumar, A., Giraldo, F., & Weiss, D. S. (2019). Stern–Gerlach detection of neutral-atom qubits in a state-dependent optical lattice. *Nature Physics*, 15(6), 538–542.

[150] Wu, Y., Bao, W. S., Cao, S., Chen, F., Chen, M. C., Chen, X., Chung, T. H., Deng, H., Du, Y., Fan, D., Gong, M., Guo, C., Guo, C., Guo, S., Han, L., Hong, L., Huang, H. L., Huo, Y. H., Li, L., Li, N., Li, S., Li, Y., Liang, F., Lin, C., Lin, J., Qian, H., Qiao, D., Rong, H.,

Su, H., Sun, L., Wang, L., Wang, S., Wu, D., Xu, Y., Yan, K., Yang, W., Yang, Y., Ye, Y., Yin, J., Ying, C., Yu, J., Zha, C., Zhang, C., Zhang, H., Zhang, K., Zhang, Y., Zhao, H., Zhao, Y., Zhou, L., Zhu, Q., Lu, C. Y., Peng, C. Z., Zhu, X., & Pan, J. W. (2021). Strong Quantum Computational Advantage Using a Superconducting Quantum Processor. *Physical Review Letters*, 127(18), 180501.

[151] Wu, Y., Kolkowitz, S., Puri, S., & Thompson, J. D. (2022). Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays. *Nature Communications*, 13(1), 1–7.

[152] Xu, Q., Bonilla Ataides, J. P., Pattison, C. A., Raveendran, N., Bluvstein, D., Wurtz, J., Vasić, B., Lukin, M. D., Jiang, L., & Zhou, H. (2024a). Constant-overhead fault-tolerant quantum computation with reconfigurable atom arrays. *Nature Physics*, 20(7), 1084–1090.

[Xu et al.] Xu, Q., Pablo, J., Ataides, B., Pattison, C. A., Raveendran, N., Bluvstein, D., Wurtz, J., Vasi´cvasi´c, B., Lukin, M. D., Jiang, L., & Zhou, H. Constant-Overhead Fault-Tolerant Quantum Computation with Reconfigurable Atom Arrays.

[154] Xu, Q., Zhou, H., Zheng, G., Bluvstein, D., Pablo, J., Ataides, B., Lukin, M. D., & Jiang, L. (2024b). Fast and Parallelizable Logical Computation with Homological Product Codes.

[155] Xu, W., Venkatramani, A. V., Cantú, S. H., Šumarac, T., Klüsener, V., Lukin, M. D., & Vuletić, V. (2021). Fast Preparation and Detection of a Rydberg Qubit Using Atomic Ensembles. *Physical Review Letters*, 127(5), 050501.

[156] Yang, J., He, X., Guo, R., Xu, P., Wang, K., Sheng, C., Liu, M., Wang, J., Derevianko, A., & Zhan, M. (2016). Coherence Preservation of a Single Neutral Atom Qubit Transferred between Magic-Intensity Optical Traps. *Physical Review Letters*, 117(12), 123201.

[157] Yavuz, D. D., Kulatunga, P. B., Urban, E., Johnson, T. A., Proite, N., Henage, T., Walker, T. G., & Saffman, M. (2006). Fast ground state manipulation of neutral atoms in microscopic optical traps. *Physical Review Letters*, 96(6), 063001.

[158] Zhou, H., Zhao, C., Cain, M., Bluvstein, D., Duckering, C., Hu, H.-Y., Wang, S.-T., Kubica, A., & Lukin, M. D. (2024). Algorithmic Fault Tolerance for Fast Quantum Computing.

[159] Ziegler, J. F., Curtis, H. W., Muhlfeld, H. P., Montrose, C. J., Chin, B., Nicewicz, M., Russell, C. A., Wang, W. Y., Freeman, L. B., Hosier, P., LaFave, L. E., Walsh, J. L., Orro, J. M., Unger, G. J., Ross, J. M., O'Gorman, T. J., Messina, B., Sullivan, T. D., Sykes, A. J., Yourke, H., Enger, T. A., Tolat, V., Scott, T. S., Taber, A. H., Sussman, R. J., Klein, W. A., & Wahaus, C. W. (1996). IBM experiments in soft fails in computer electronics (1978-1994). *IBM Journal of Research and Development*, 40(1), 3–16.