

Tour de gross: A modular quantum computer based on bivariate bicycle codes

Theodore Yoder **Eddie Schoute** Patrick Rall Emily Pritchett
Jay Gambetta Andrew Cross Malcolm Carroll Michael Beverland¹

Feb 2026
arXiv:2506.03094

¹In reverse alphabetical order

Introduction

Scaling beyond the surface code

Surface code qubit

1 logical qubit requires $2(d + 1)^2$
physical qubits.

Scaling beyond the surface code

Surface code qubit

1 logical qubit requires $2(d + 1)^2$
physical qubits.

Bivariate bicycle codes

Gross ($d = 12$): 12 logical qubits per
288 physical qubits.

Scaling beyond the surface code

Surface code qubit

1 logical qubit requires $2(d + 1)^2$
physical qubits.

Bivariate bicycle codes

Gross ($d = 12$): 12 logical qubits per
288 physical qubits.

Two-gross ($d = 18$): 12 logical qubits
per 576 physical qubits.

Scaling beyond the surface code

Surface code qubit

1 logical qubit requires $2(d + 1)^2$ physical qubits.

Bivariate bicycle codes

Gross ($d = 12$): 12 logical qubits per 288 physical qubits.

Two-gross ($d = 18$): 12 logical qubits per 576 physical qubits.

Surface codes require 288 ($d = 11$) and 648 ($d = 17$) physical qubits.

Scaling beyond the surface code

Surface code qubit

1 logical qubit requires $2(d + 1)^2$ physical qubits.

Bivariate bicycle codes

Gross ($d = 12$): 12 logical qubits per 288 physical qubits.

Two-gross ($d = 18$): 12 logical qubits per 576 physical qubits.

Surface codes require 288 ($d = 11$) and 648 ($d = 17$) physical qubits.

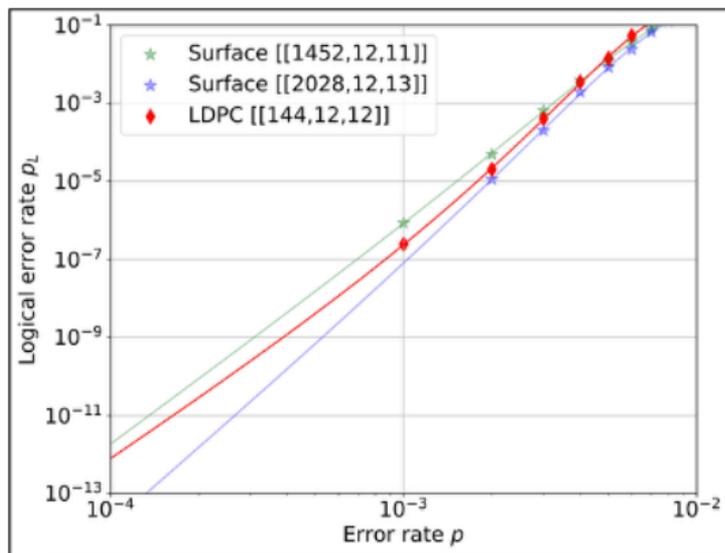


Figure: Compared to surface code, the qLDPC code requires about 10x fewer physical qubits at comparable k and d [Bravyi et al., Nature 627 (2024)].

The gross code and friends

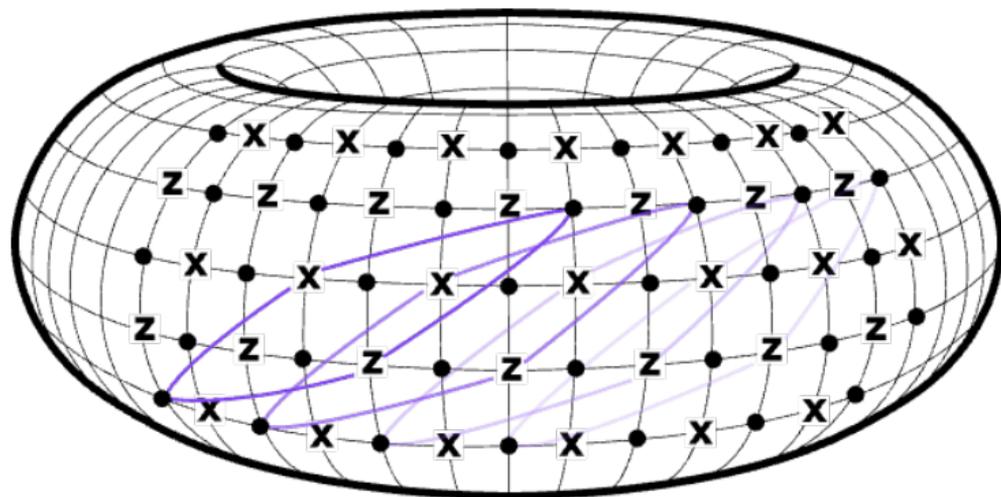


Figure: The gross code is visually represented by a torus due to its Tanner graph embedding plus two long-range connections.

HW demo: 6-way couplers

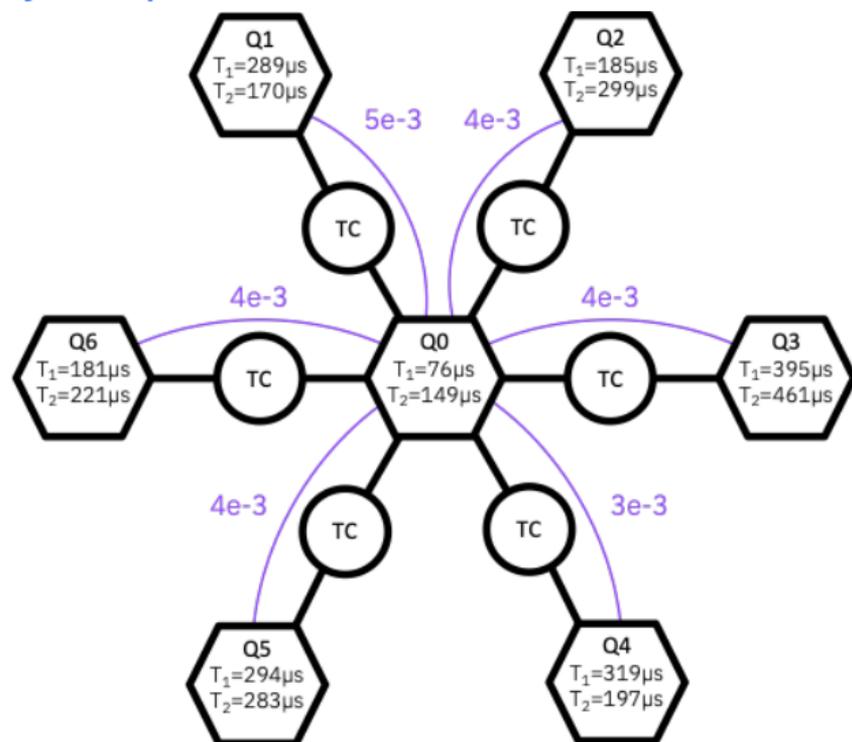


Figure: Degree-6 connectivity without loss of gate fidelity (purple) using tunable couplers (TC)

HW demo: Long-range gates

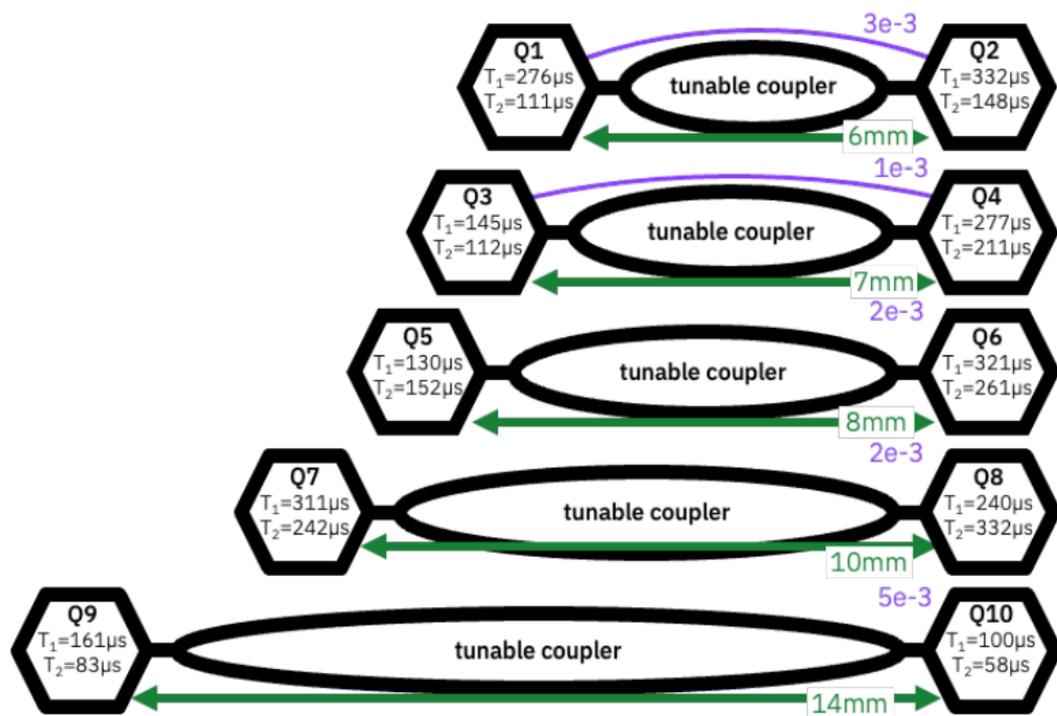


Figure: Demonstration of long-range gates required for the gross code.

HW demo: Non-planar interconnect

Single additional layer of low-loss interconnect

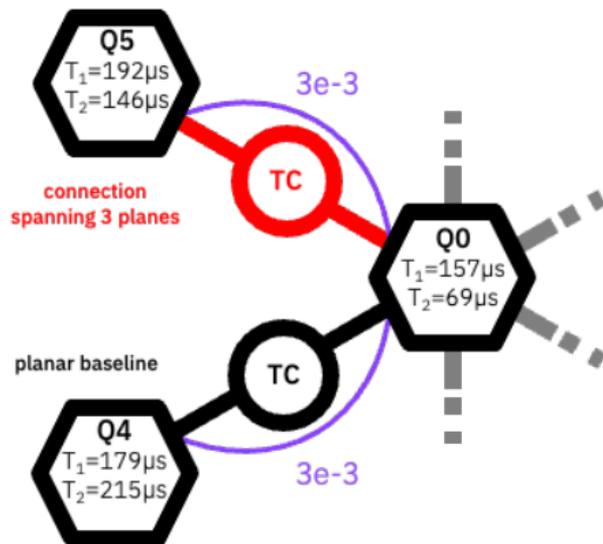
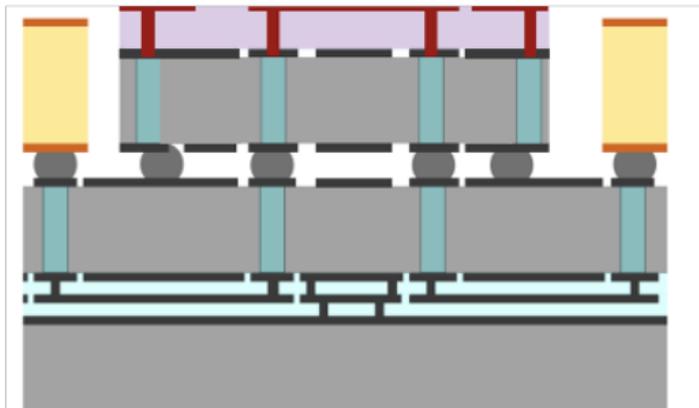
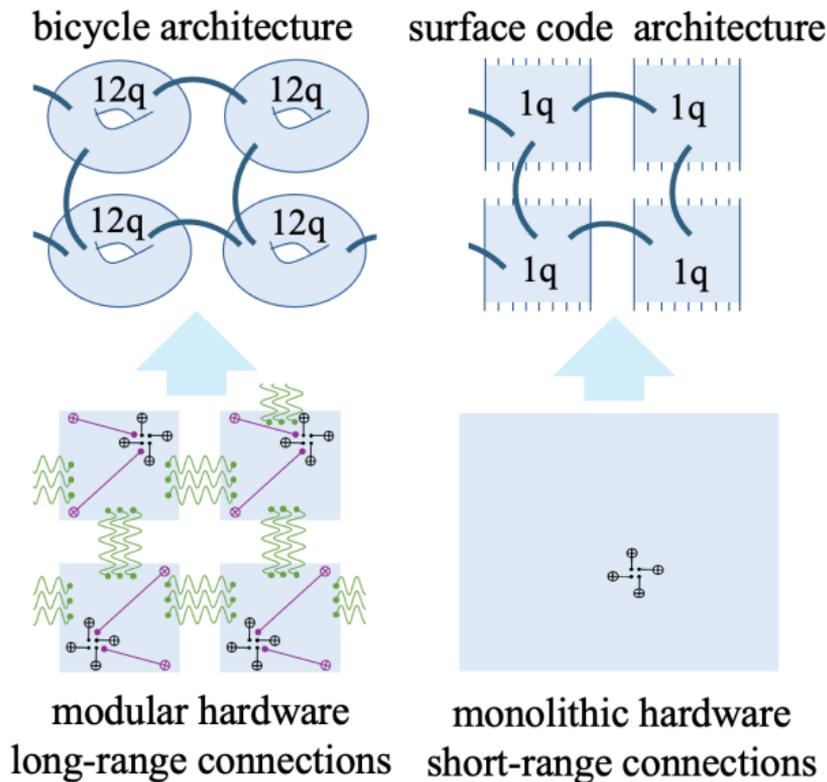


Figure: Demonstration of multi-layer interconnect required for non-planar connectivity in the gross code

Long-range connections enable modularity and scale



Logical capability estimates

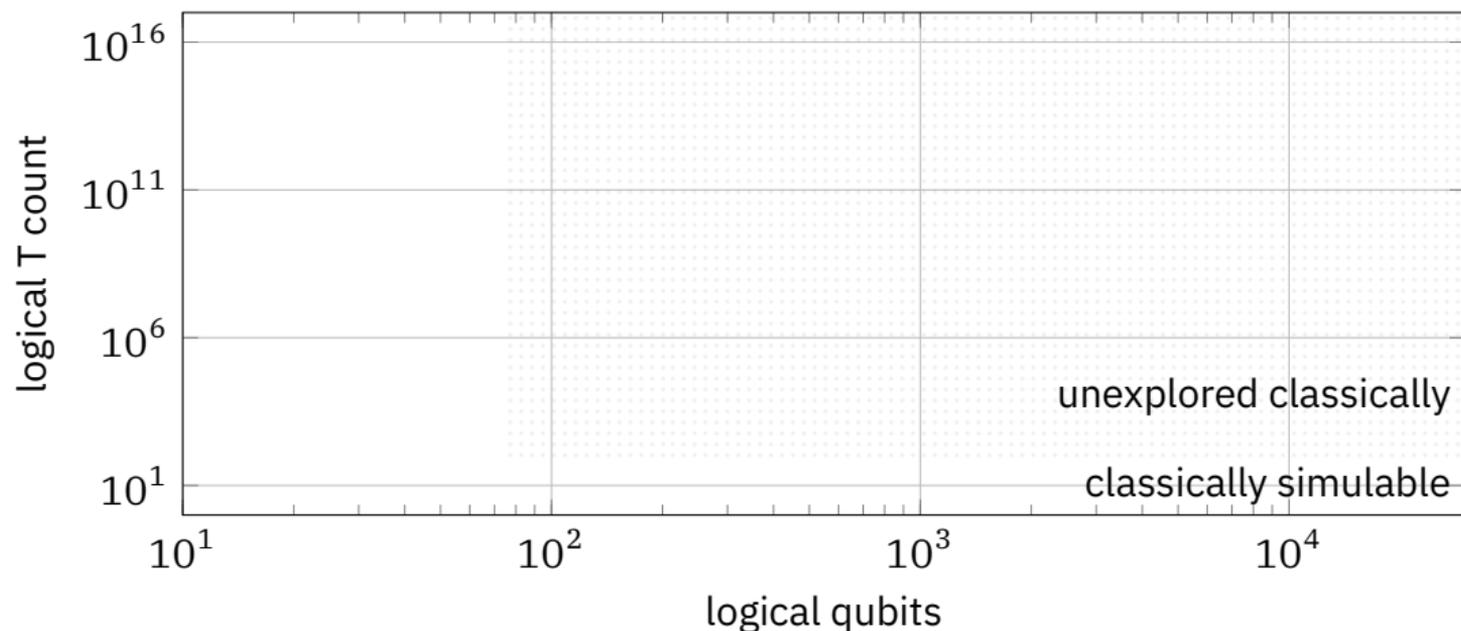


Figure: Bicycle architecture capabilities using gross (red filled ellipse) and two-gross (blue hollow ellipse) codes when compared to surface code architectures (black square) given q physical qubits and p physical error rate.

Logical capability estimates

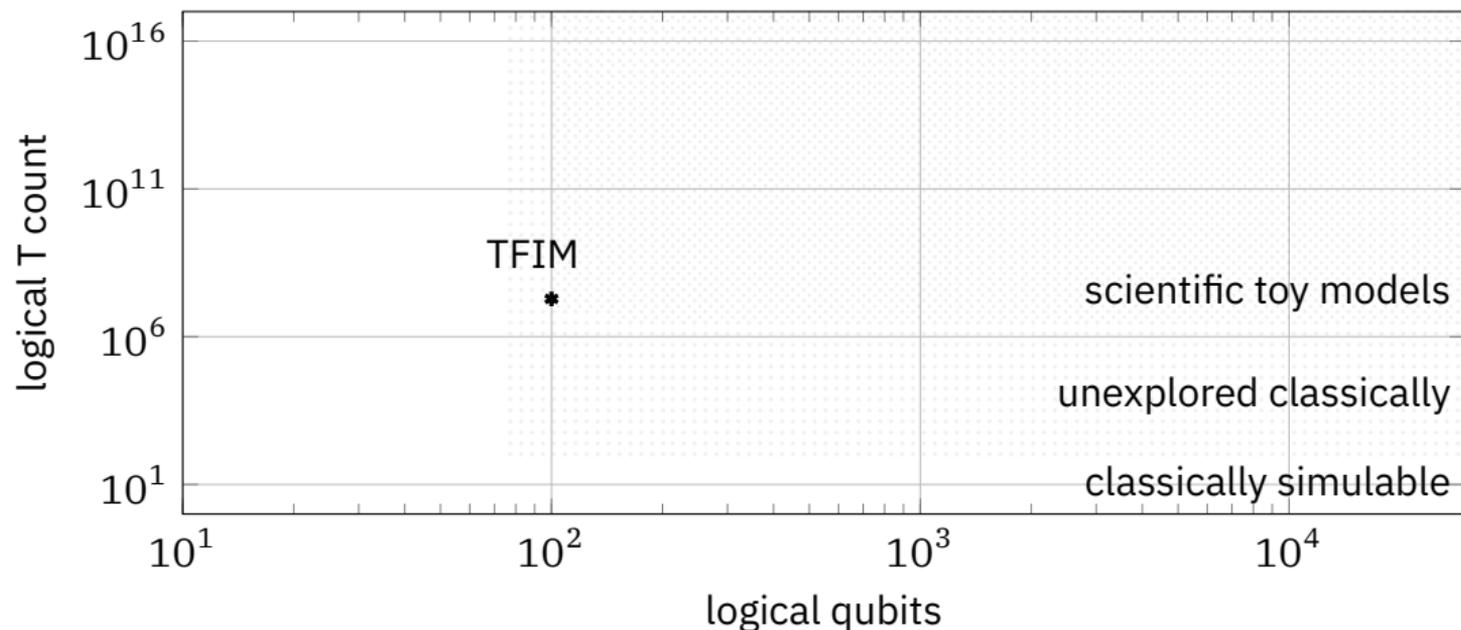


Figure: Bicycle architecture capabilities using gross (red filled ellipse) and two-gross (blue hollow ellipse) codes when compared to surface code architectures (black square) given q physical qubits and p physical error rate.

Logical capability estimates

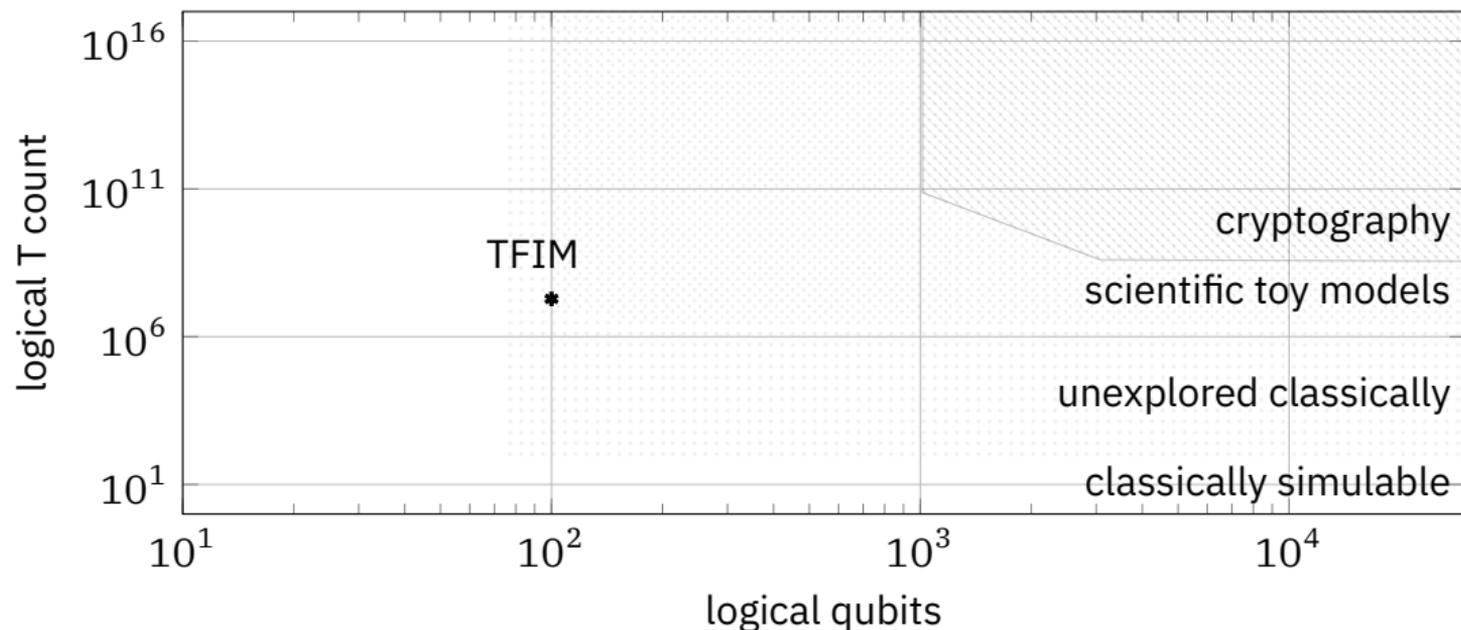


Figure: Bicycle architecture capabilities using gros (red filled ellipse) and two-gros (blue hollow ellipse) codes when compared to surface code architectures (black square) given q physical qubits and p physical error rate.

Logical capability estimates

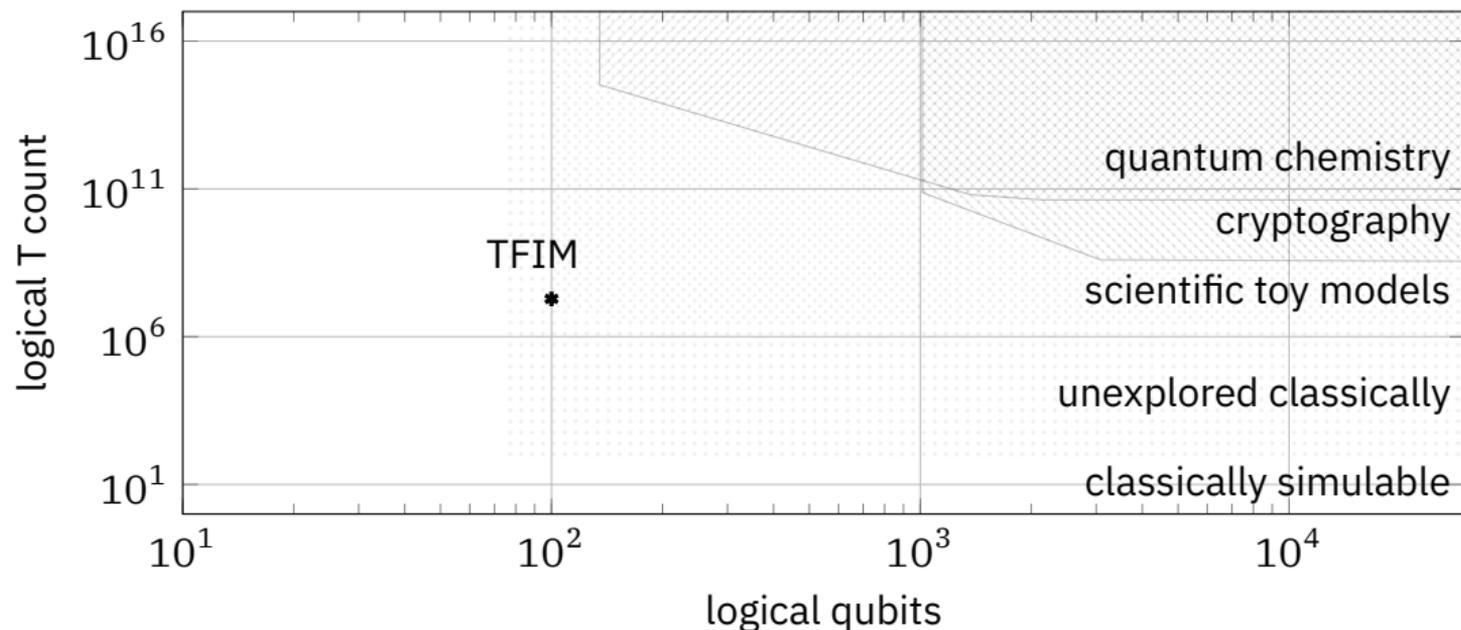


Figure: Bicycle architecture capabilities using gross (red filled ellipse) and two-gross (blue hollow ellipse) codes when compared to surface code architectures (black square) given q physical qubits and p physical error rate.

Logical capability estimates

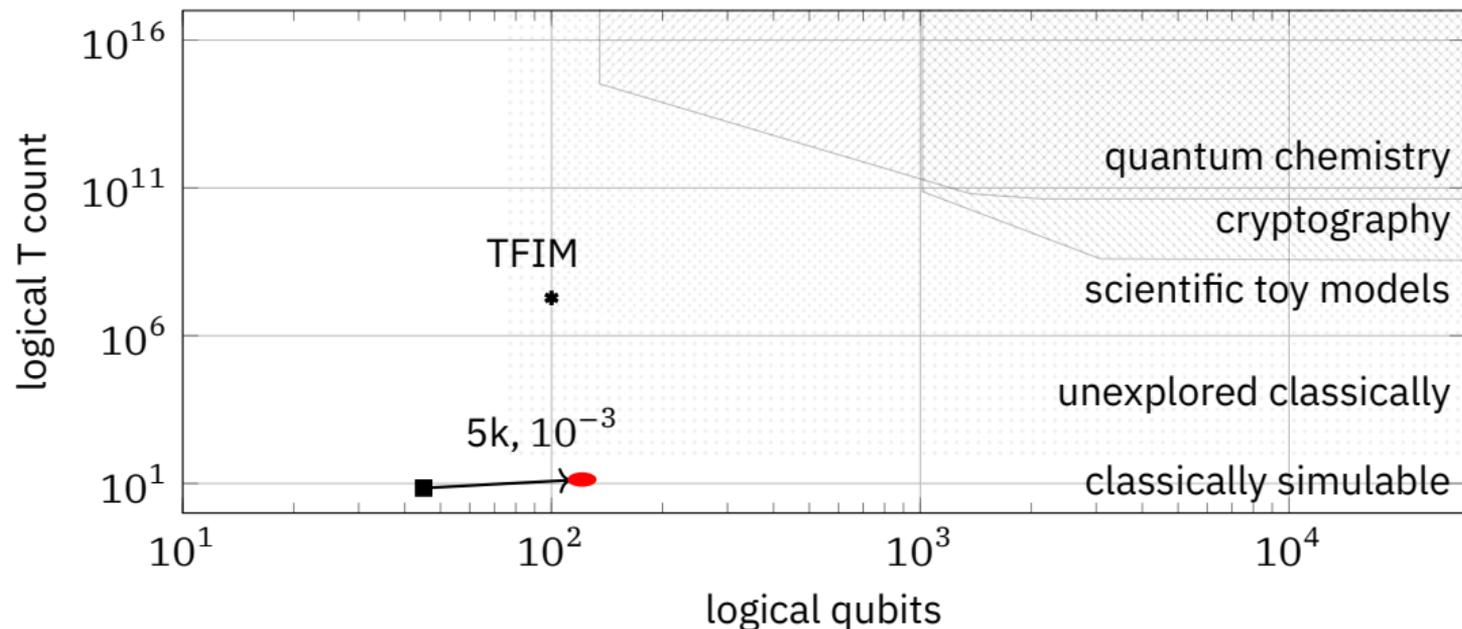


Figure: Bicycle architecture capabilities using gross (red filled ellipse) and two-gross (blue hollow ellipse) codes when compared to surface code architectures (black square) given q physical qubits and p physical error rate.

Logical capability estimates

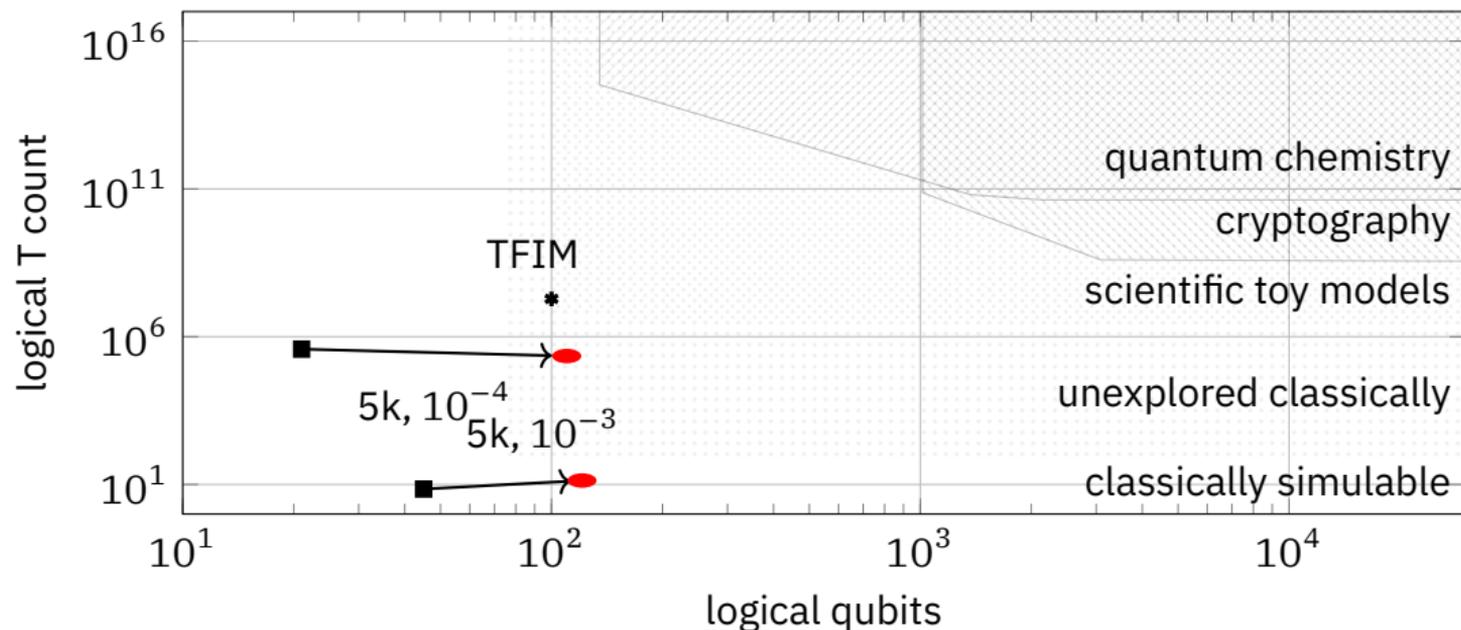


Figure: Bicycle architecture capabilities using gross (red filled ellipse) and two-gross (blue hollow ellipse) codes when compared to surface code architectures (black square) given q physical qubits and p physical error rate.

Logical capability estimates

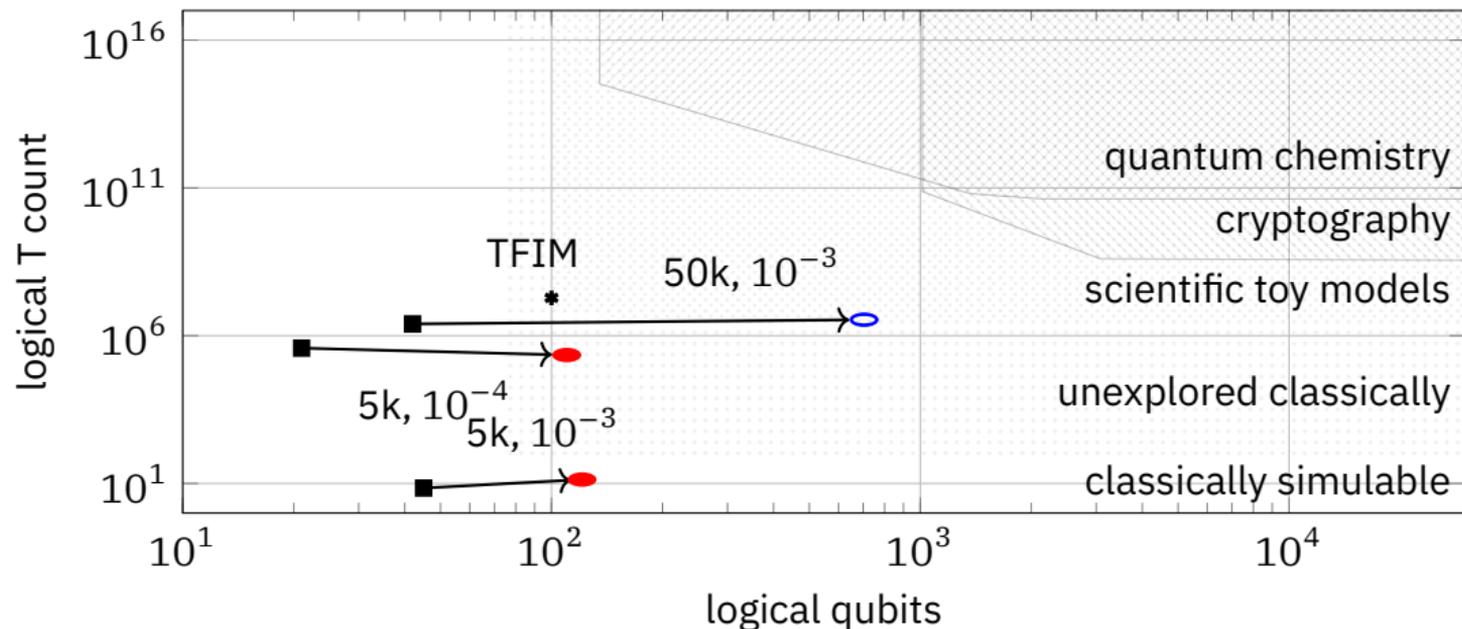


Figure: Bicycle architecture capabilities using gross (red filled ellipse) and two-gross (blue hollow ellipse) codes when compared to surface code architectures (black square) given q physical qubits and p physical error rate.

Logical capability estimates

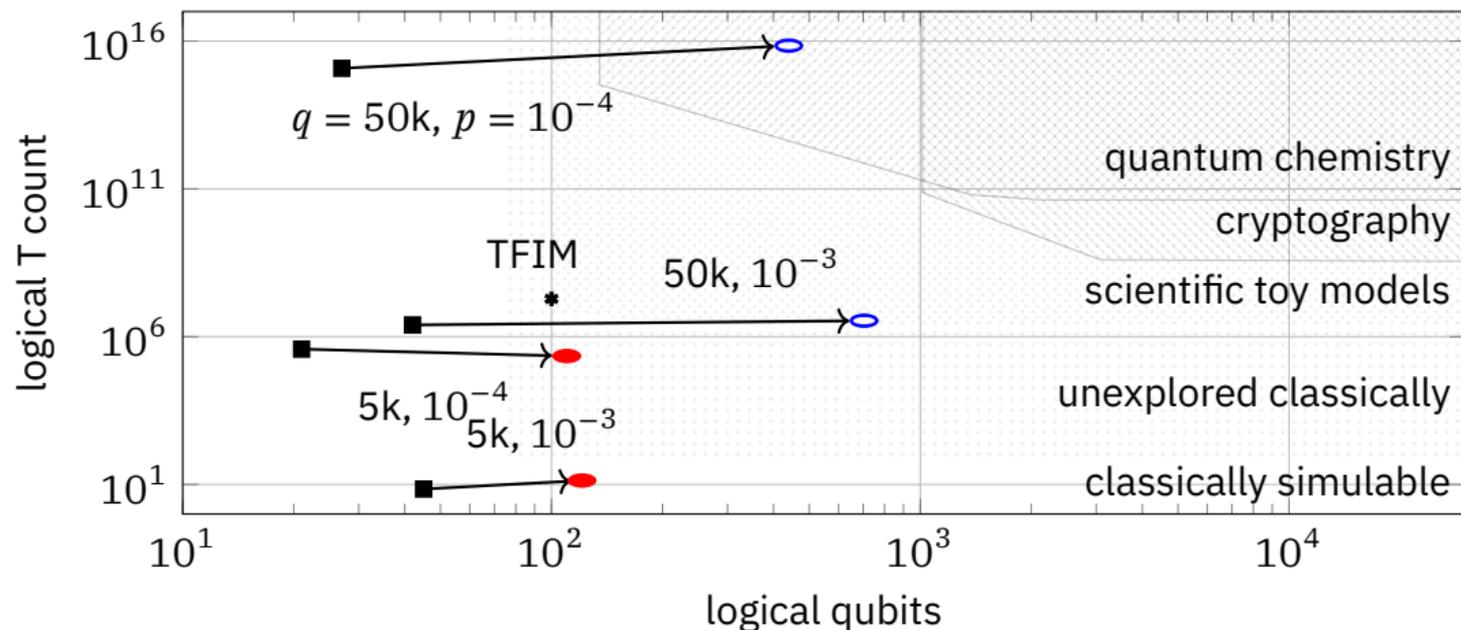


Figure: Bicycle architecture capabilities using gross (red filled ellipse) and two-gross (blue hollow ellipse) codes when compared to surface code architectures (black square) given q physical qubits and p physical error rate.

Logical capability estimates

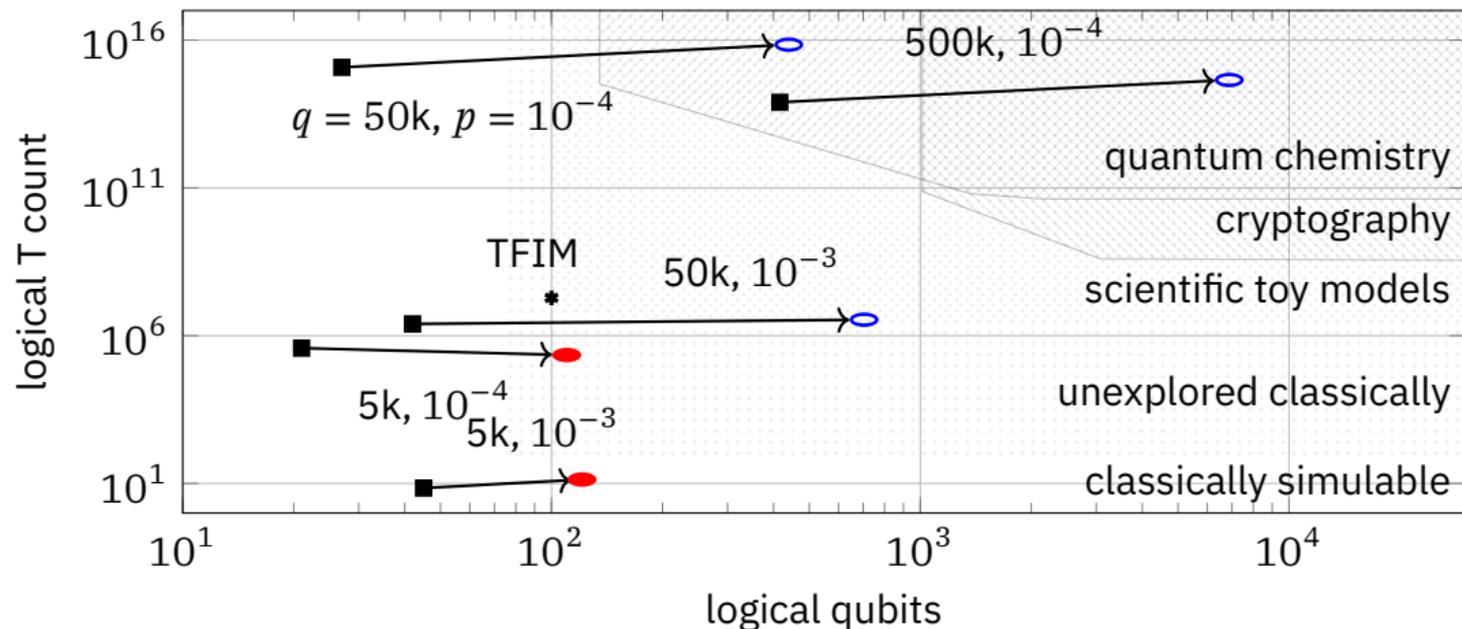


Figure: Bicycle architecture capabilities using gross (red filled ellipse) and two-gross (blue hollow ellipse) codes when compared to surface code architectures (black square) given q physical qubits and p physical error rate.

Bicycle architecture

Two code modules

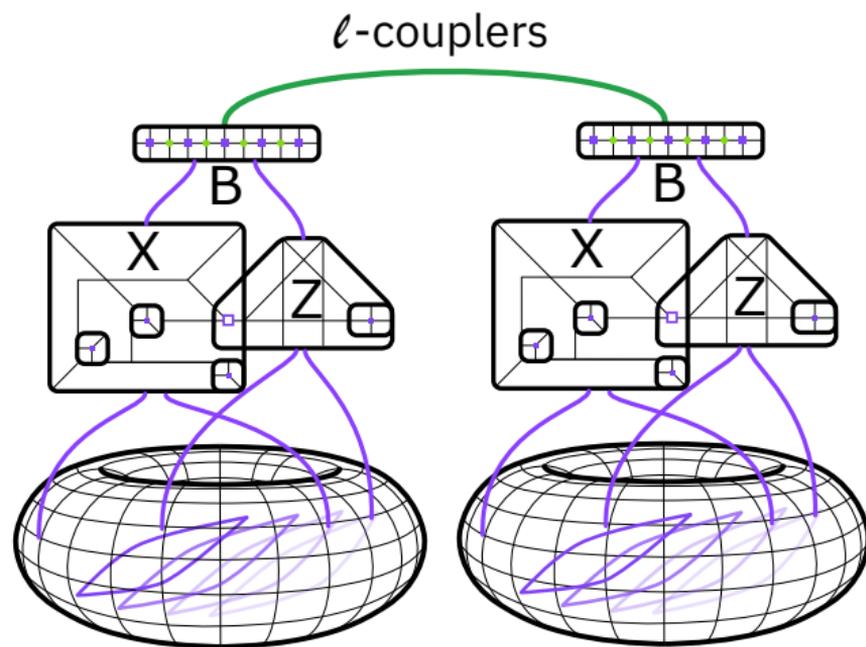


Figure: A gross code is just a memory. We can perform computation by attaching an ancilla system called the *logical processing unit* (LPU) [Cross, He, Rall, Yoder (2024); Williamson, Yoder (2024)]. By connecting LPUs through a bridge system, we can perform joint measurements between two codes.

Bicycle architecture

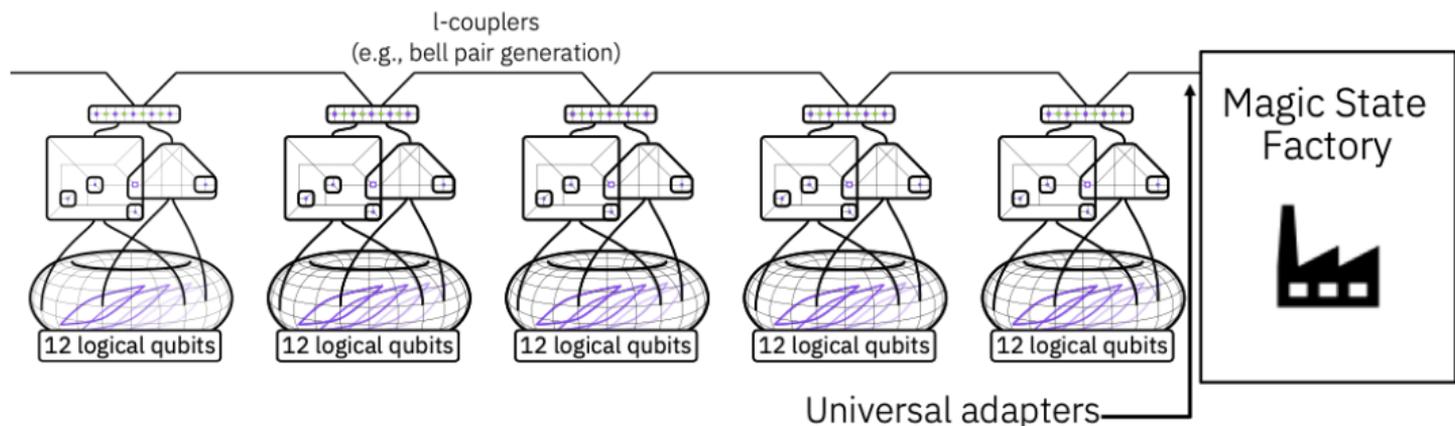


Figure: The bicycle architecture consists of many connected code modules and a (magic state) factory module.

HW demo: L-coupler

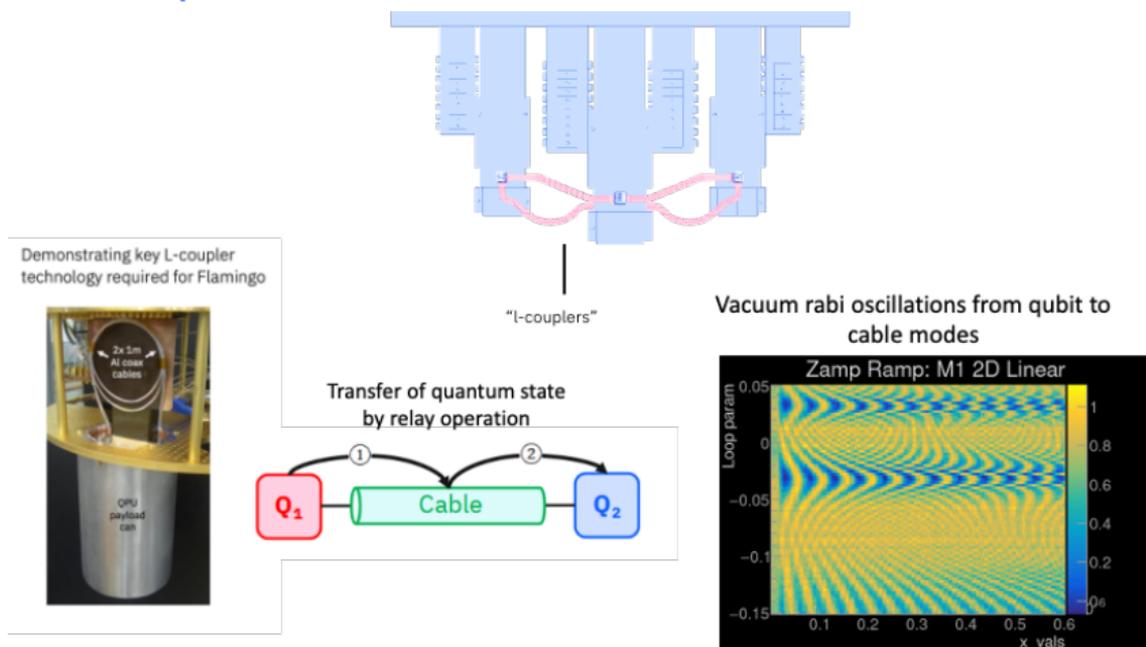


Figure: Demonstration of entangling gate operation over L-coupled interface of up to 96.5% fidelity through interconnect of about 1 meter [Shawn Hall, Kentaro Heya, Yadav Prasad Kandel, Moein Malek, Yves Martin, Jae-Woong Nah, Jason Orcutt, Timothy Phung, Rachel Steiner, Neereja Sundaresan].

Physical qubit costs

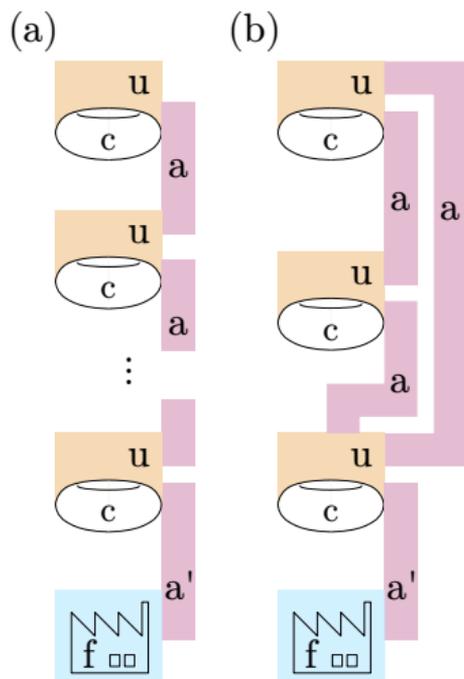


Figure: Code modules (c) with logical processing units (u) and a factory module (f) connected in two ways via adapters (a, a').

Physical qubit costs

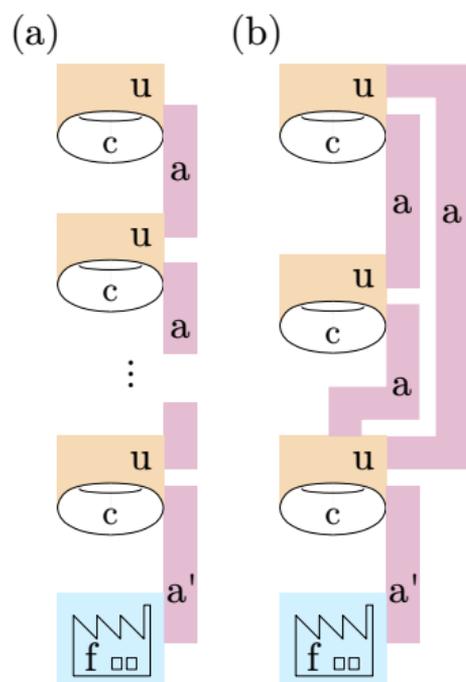


Figure: Code modules (c) with logical processing units (u) and a factory module (f) connected in two ways via adapters (a, a').

system	p	gross	two-gross
c		288	576
u		90	158
a		22	34
f	10^{-3}	454	810
	10^{-4}	463	18,600
a'	10^{-3}	29	13
	10^{-4}	29	49

Table: Physical qubit counts for each system in the architecture

Bicycle instructions

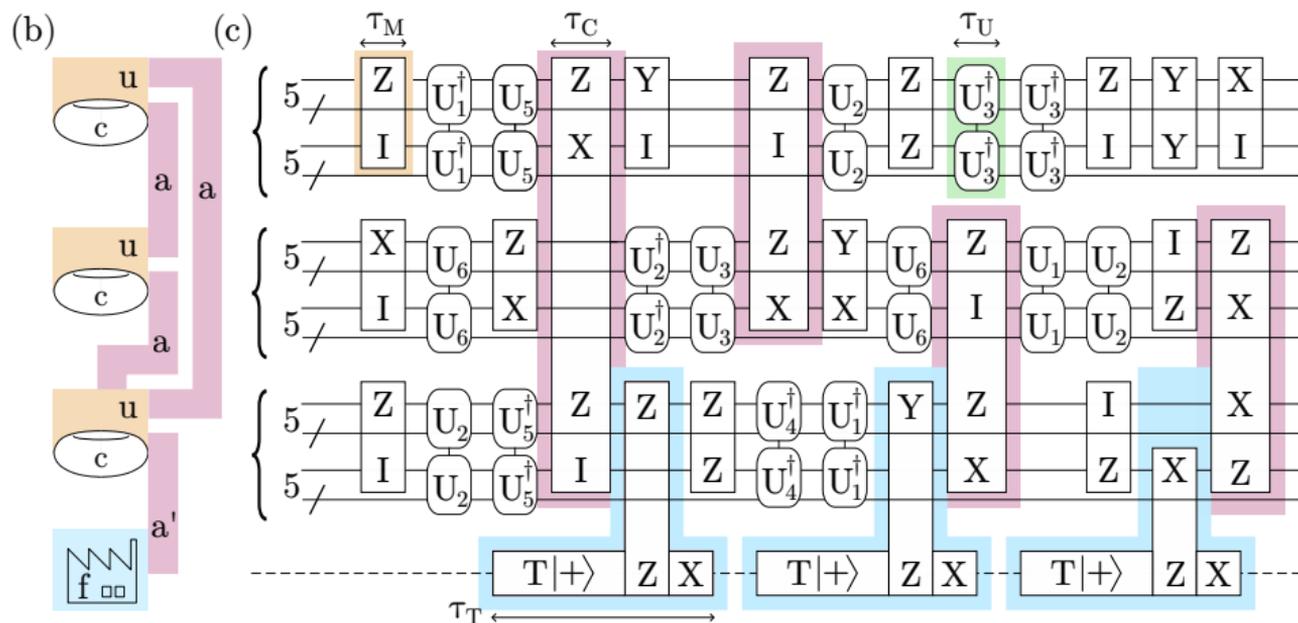


Figure: Given connectivity (b) we define the available *bicycle instructions* (c) on the bicycle architecture: In-module measurement (yellow) and inter-module measurement (red); Automorphisms $U_i \otimes U_i$ (green) and T state injection (blue).

Bicycle instruction error rates

Real-time decoding: Relay-BP

We use Relay-BP, which can achieve a orders of magnitude improvement over BP+OSD+CS-10 and is suitable for real-time implementation [Müller et al., arXiv:2506.01779]

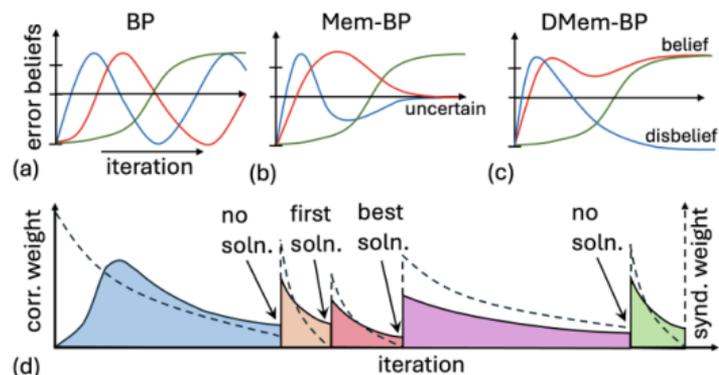


Figure: Relay-BP chains multiple runs of belief propagation with disordered memory strengths.

Real-time decoding: Relay-BP

We use Relay-BP, which can achieve a orders of magnitude improvement over BP+OSD+CS-10 and is suitable for real-time implementation [Müller et al., arXiv:2506.01779]

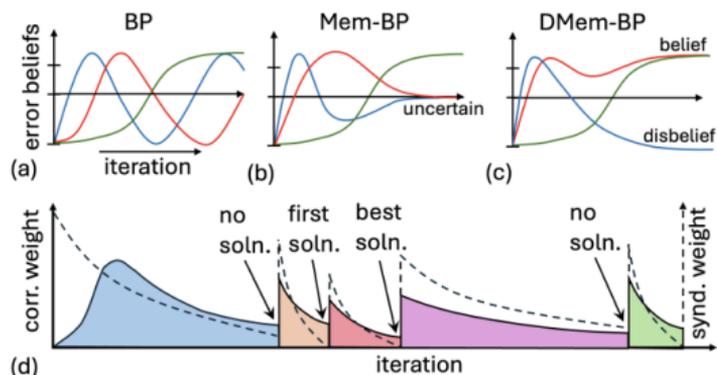


Figure: Relay-BP chains multiple runs of belief propagation with disordered memory strengths.

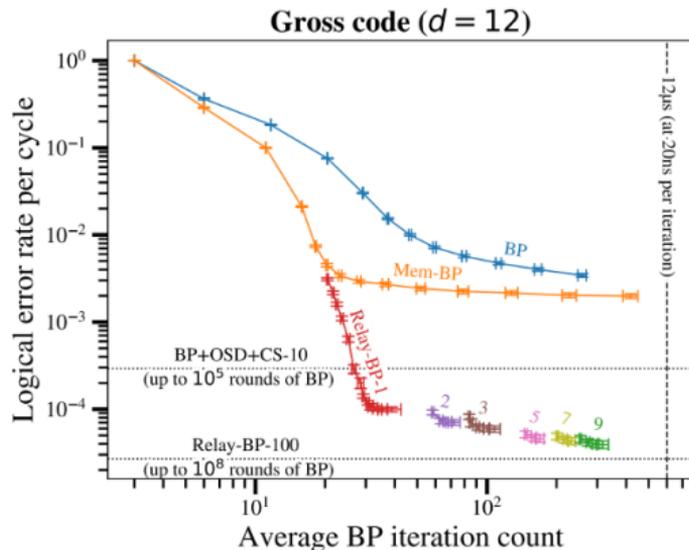
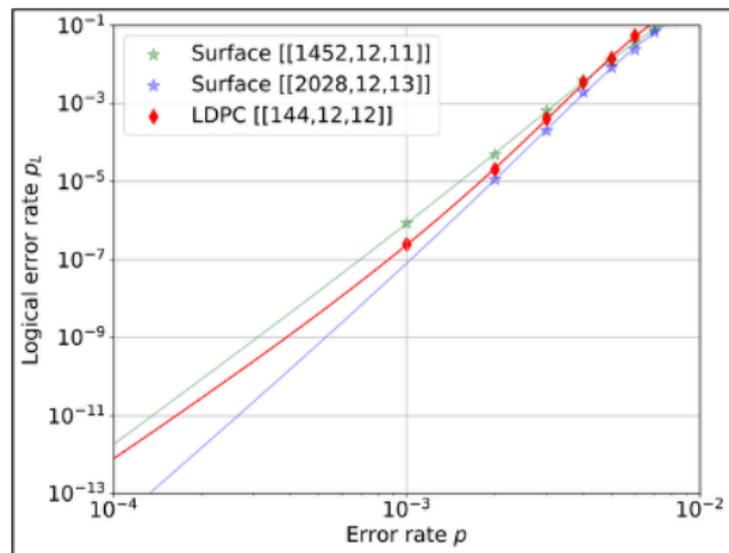


Figure: Substantially lower error rates within 600 iterations

Monte Carlo limitations



- ▶ Monte Carlo can sample at high error rates
- ▶ How do we figure out what's going on at low p ?

Extrapolation from low-weight failure fraction (gross)

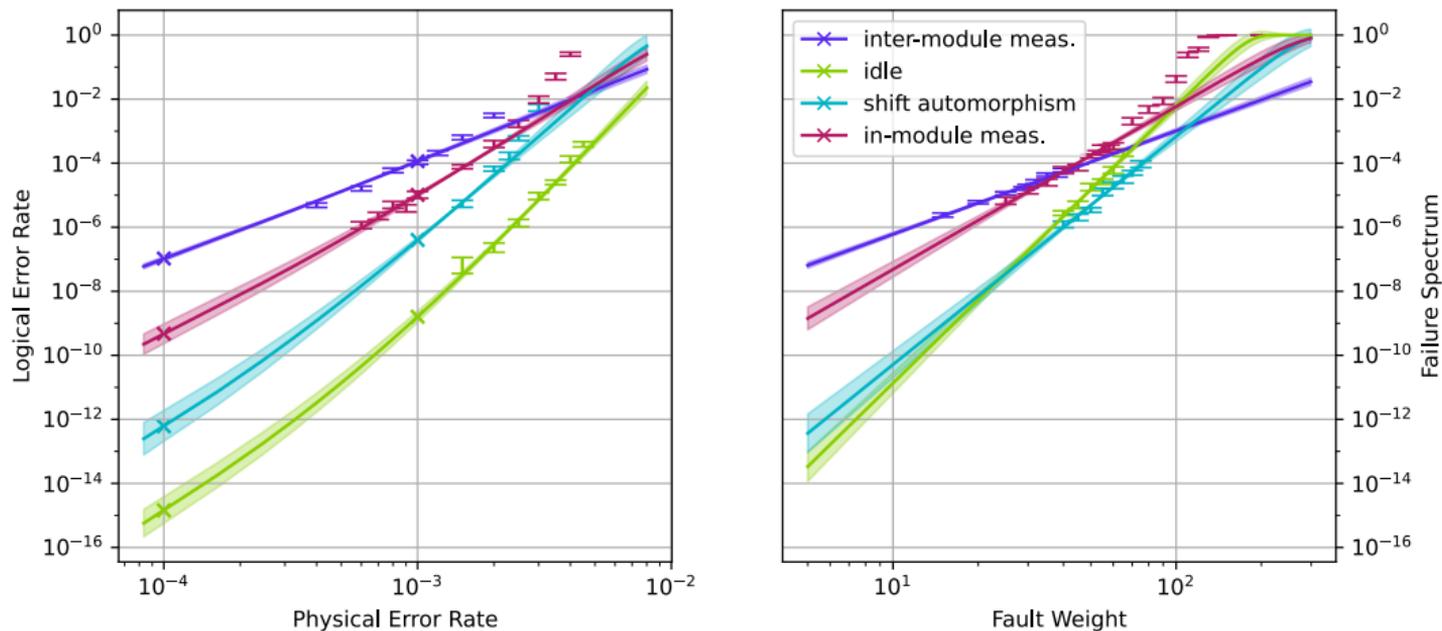


Figure: We sample sets of exactly w circuit faults and estimate failure rate (right). We then fit ansatz function and translate back to logical error rate (left) and observe agreement with independent Monte Carlo data. [Beverland et al. (2025)]

Extrapolation from low-weight failure fraction (two-gross)

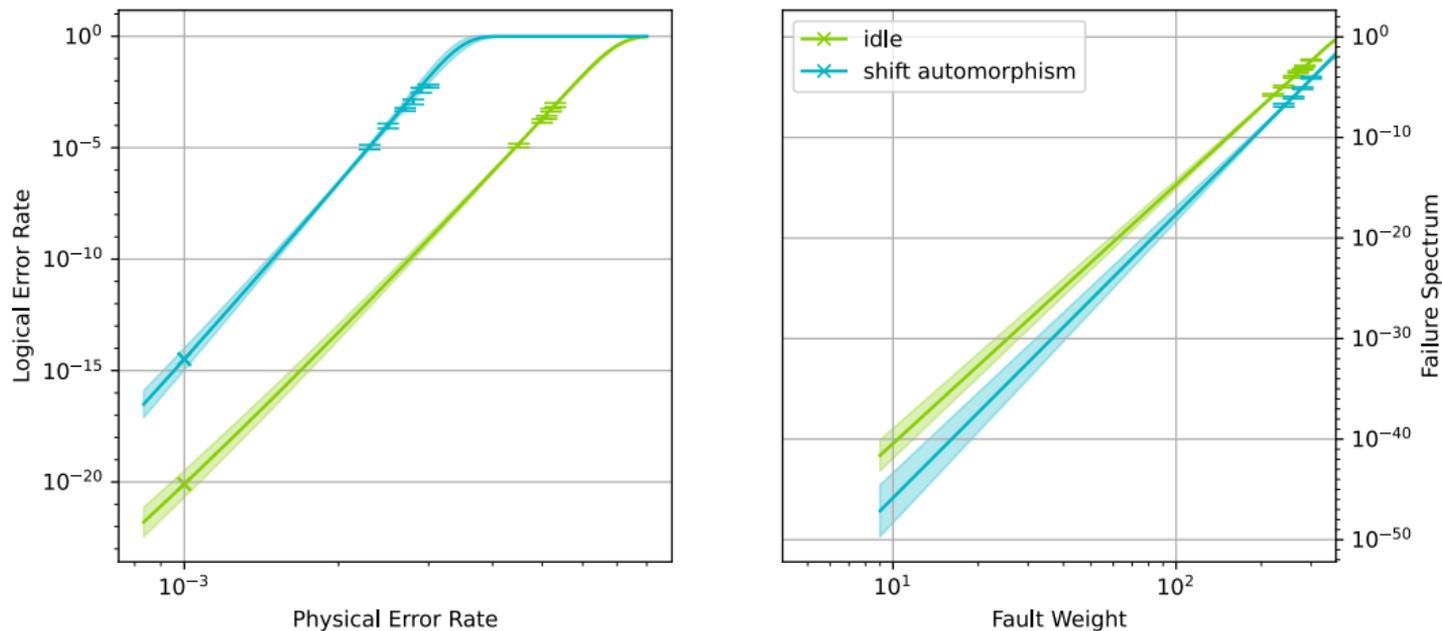


Figure: We sample sets of exactly w circuit faults and estimate failure rate (right). We then fit ansatz function and translate back to logical error rate (left) and observe agreement with independent Monte Carlo data. [Beverland et al. (2025)]

Bicycle instruction properties

instruction	module	τ_i (timesteps)	P_i (logical error rate)	
			$p = 10^{-3}$	$p = 10^{-4}$
idle	gross	8	$10^{-8.8 \pm 0.2}$	$10^{-14.8 \pm 0.4}$
	two-gross	8	$10^{-20.1 \pm 0.5}$	$10^{-38.3 \pm 0.9}$
shift automorphism	gross	14	$10^{-6.4 \pm 0.2}$	$10^{-12.2 \pm 0.5}$
	two-gross	14	$10^{-14.5 \pm 0.4}$	$10^{-37 \pm 1}$
in-module meas.	gross	120	$10^{-5.0 \pm 0.1}$	$10^{-9.0 \pm 0.2}$
	two-gross	216	[10^{-11}]	[10^{-20}]
inter-module meas.	gross	120	$10^{-2.7 \pm 0.1}$	$10^{-7.3 \pm 0.3}$
	two-gross	216	[10^{-9}]	[10^{-18}]

Table: Time duration and logical error rate properties of bicycle instructions. We omit here T state injection, using magic state cultivation [Gidney, Shutty, Jones (2024)] or distillation [Litinski (2019)]. Bracketed numbers are assumptions without data.

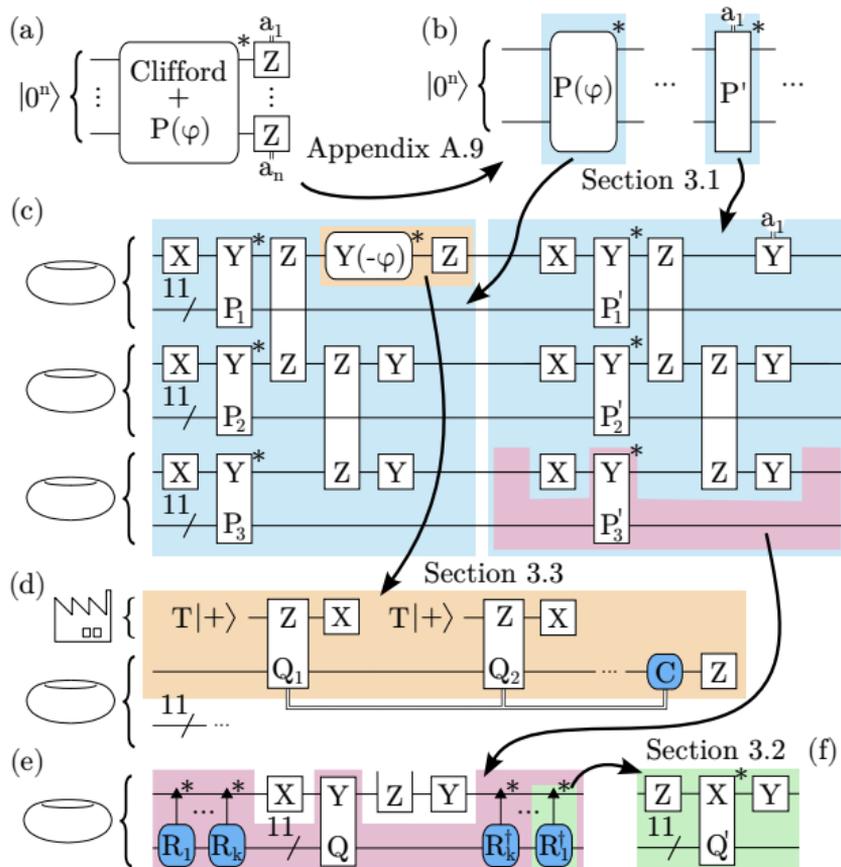
Bicycle instruction properties

instruction	module	τ_i (timesteps)	P_i (logical error rate)	
			$p = 10^{-3}$	$p = 10^{-4}$
idle	gross	8	$10^{-8.8 \pm 0.2}$	$10^{-14.8 \pm 0.4}$
	two-gross	8	$10^{-20.1 \pm 0.5}$	$10^{-38.3 \pm 0.9}$
shift automorphism	gross	14	$10^{-6.4 \pm 0.2}$	$10^{-12.2 \pm 0.5}$
	two-gross	14	$10^{-14.5 \pm 0.4}$	$10^{-37 \pm 1}$
in-module meas.	gross	120	$10^{-5.0 \pm 0.1}$	$10^{-9.0 \pm 0.2}$
	two-gross	216	[10^{-11}]	[10^{-20}]
inter-module meas.	gross	120	$10^{-4.0 \pm 0.2}$	$10^{-7.3 \pm 0.3}$
	two-gross	216	[10^{-9}]	[10^{-18}]

Table: Time duration and logical error rate properties of bicycle instructions. We omit here T state injection, using magic state cultivation [Gidney, Shutty, Jones (2024)] or distillation [Litinski (2019)]. Bracketed numbers are assumptions without data.

Compiling for the bicycle architecture

Compiler summary



Compiler implementation

Compiler implementation available at
`github.com/qiskit-community/bicycle-architecture-compiler`

Graphical language

Round boxes are unitary gates

$$V := \text{---} \boxed{V} \text{---}, \quad e^{i\frac{\pi}{4}P} := \text{---} \boxed{P} \text{---}, \quad (1)$$

and square boxes are measurements

$$\frac{\mathbb{1} + (-1)^a P}{2} := \text{---} \boxed{P} \text{---}, \quad (2)$$

for measurement outcome $a \in \{0, 1\}$.

Measurement projection

Frequently, we don't care about the measurement outcome because it's random.

Measurement projection

Frequently, we don't care about the measurement outcome because it's random. We draw a measurement without an outcome wire and define it as a *measurement projection*,

$$\text{[P]} := \text{[P]}\text{[Q]} \quad (3)$$

Effectively, it is the projection $\frac{1+P}{2}$.

Measurement projection

Frequently, we don't care about the measurement outcome because it's random. We draw a measurement without an outcome wire and define it as a *measurement projection*,

$$\text{[P]} := \text{[P]}\text{[Q]} \quad (3)$$

Effectively, it is the projection $\frac{\mathbb{1}+P}{2}$.

This is only possible if we know an anti-commuting stabilizer Q before measuring P so that

$$Q^a \frac{\mathbb{1} + (-1)^a P}{2} |\psi\rangle = \frac{\mathbb{1} + P}{2} Q^a |\psi\rangle = \frac{\mathbb{1} + P}{2} |\psi\rangle. \quad (4)$$

Measurement ancilla

Let us focus on Pauli-generated rotations,

$$P(\varphi) := e^{i\varphi P}. \quad (5)$$

Measurement ancilla

Let us focus on Pauli-generated rotations,

$$P(\varphi) := e^{i\varphi P}. \quad (5)$$

A Pauli-generated rotation can be implemented by

$\text{---} \boxed{P(\varphi)} \text{---} = \text{---} \begin{array}{c} \boxed{X} \text{---} \bullet \text{---} \boxed{X(\varphi)} \text{---} \boxed{Z} \text{---} \\ | \\ \boxed{P} \text{---} \end{array} \text{---}$ (6)

Distributing across data modules

We can add some ancillas to

$$\text{---} \boxed{P(\varphi)} \text{---} = \text{---} \boxed{X} \text{---} \bullet \text{---} \boxed{X(\varphi)} \text{---} \boxed{Z} \text{---} \\ \text{---} \boxed{P} \text{---}$$

(7)

and assume $P = P_1 \otimes P_2 \otimes P_3$

Distributing across data modules

We can add some ancillas to

$$\text{---} \boxed{P(\varphi)} \text{---} = \text{---} \begin{array}{c} \boxed{X} \text{---} \bullet \text{---} \boxed{X(\varphi)} \text{---} \boxed{Z} \text{---} \\ | \\ \boxed{P} \text{---} \end{array} \text{---} \quad (7)$$

and assume $P = P_1 \otimes P_2 \otimes P_3$

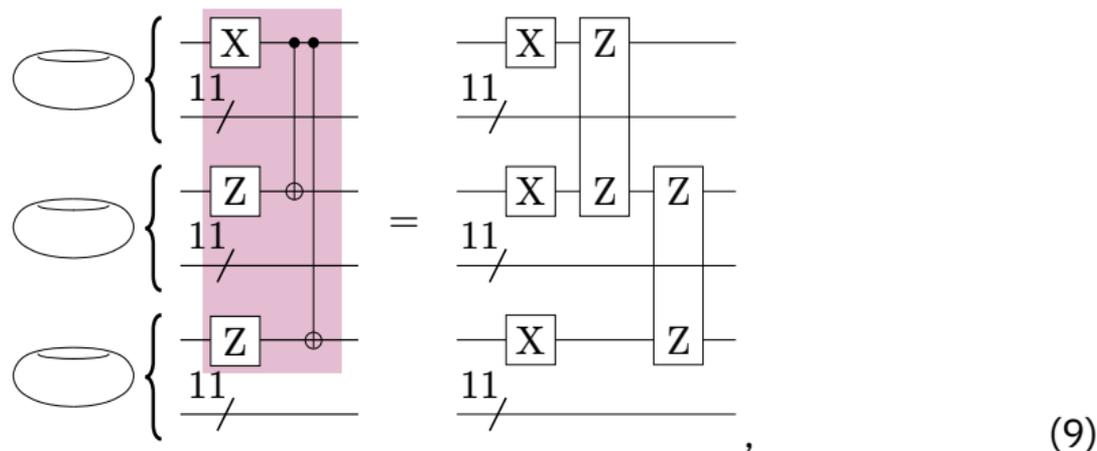
The diagram illustrates the decomposition of the gate $P(\varphi)$ into three ancilla qubits P_1 , P_2 , and P_3 . The circuit is shown in four stages, connected by equals signs:

- The first stage shows the gate $P(\varphi)$ acting on three qubits. The ancilla qubits P_1 , P_2 , and P_3 are introduced. The qubits are labeled with X , Z , and P .
- The second stage shows the decomposition of P into P_1 , P_2 , and P_3 . The qubits are labeled with Z and P .
- The third stage shows the decomposition of $X(\varphi)$ into X and P_1 . The qubits are labeled with Z and P .
- The fourth stage shows the final decomposition of the entire gate into X , Z , and P_1 , P_2 , and P_3 . The qubits are labeled with X , Z , and P .

(8)

It's a GHZ state

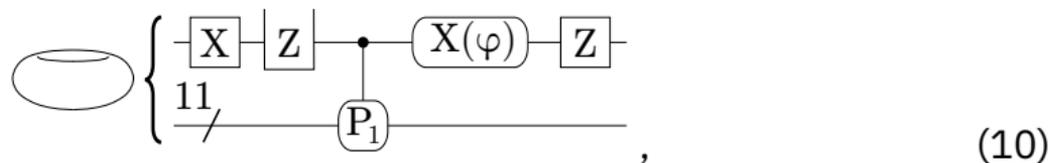
We identify the state preparation as a GHZ state



which uses just bicycle instructions (assuming connected modules).

The result

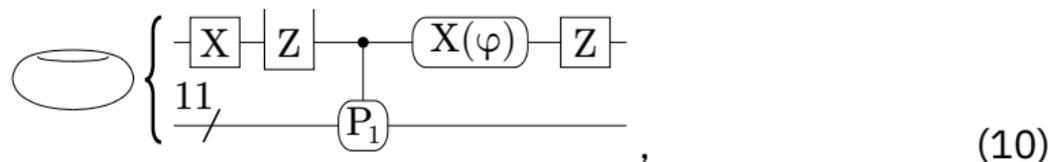
Thus, the first code module has the following circuit



with one end of a ZZ -measurement depicted.

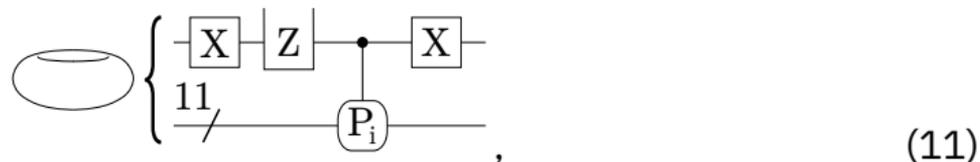
The result

Thus, the first code module has the following circuit



with one end of a ZZ -measurement depicted.

Code modules $i = 2, \dots, M$ have

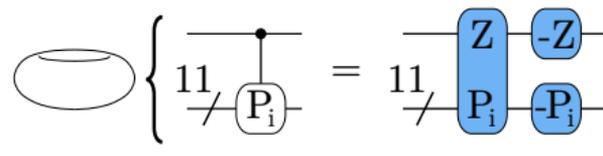


where we depict only one ZZ measurement.

(If we were compiling a Pauli measurement, then there are no instances of (10).)

Decomposing controlled- P_i

We use the identity, for any Pauli P_i ,



The diagram shows an identity for a controlled- P_i gate. On the left, a loop representing a controlled- P_i gate is shown. This is equal to a circuit with two qubits. The top qubit has a control point connected to a P_i gate on the bottom qubit. This is further decomposed into a circuit where the top qubit has a Z gate and the bottom qubit has a P_i gate, with a $-Z$ gate on the top qubit and a P_i gate on the bottom qubit. The Z and P_i gates on the top qubit are highlighted in blue.

$$\text{Controlled-}P_i = \text{Circuit with } Z, -Z, P_i, P_i \text{ gates} \quad (12)$$

Decomposing controlled- P_i

We use the identity, for any Pauli P_i ,

$$\text{torus} \left\{ \begin{array}{c} \text{---} \bullet \text{---} \\ | \\ \text{---} \text{P}_i \text{---} \end{array} \right. = \text{torus} \left\{ \begin{array}{c} \text{---} \text{Z} \text{---} \text{-Z} \text{---} \\ | \\ \text{---} \text{P}_i \text{---} \text{-P}_i \text{---} \end{array} \right. \quad (12)$$

to simplify

$$\text{torus} \left\{ \begin{array}{c} \text{---} \text{X} \text{---} \text{Z} \text{---} \bullet \text{---} \text{X}(\varphi) \text{---} \text{Z} \text{---} \\ | \\ \text{---} \text{P}_i \text{---} \end{array} \right. = \text{torus} \left\{ \begin{array}{c} \text{---} \text{X} \text{---} \text{Z} \text{---} \text{Z} \text{---} \text{-Z} \text{---} \text{X}(\varphi) \text{---} \text{Z} \text{---} \\ | \\ \text{---} \text{P}_i \text{---} \text{-P}_i \text{---} \end{array} \right. \quad (13)$$

$$= \text{torus} \left\{ \begin{array}{c} \text{---} \text{X} \text{---} \text{Z} \text{---} \text{Z} \text{---} \text{Y}(-\varphi) \text{---} \text{Z} \text{---} \\ | \\ \text{---} \text{P}_i \text{---} \text{-P}_i \text{---} \end{array} \right. \quad (14)$$

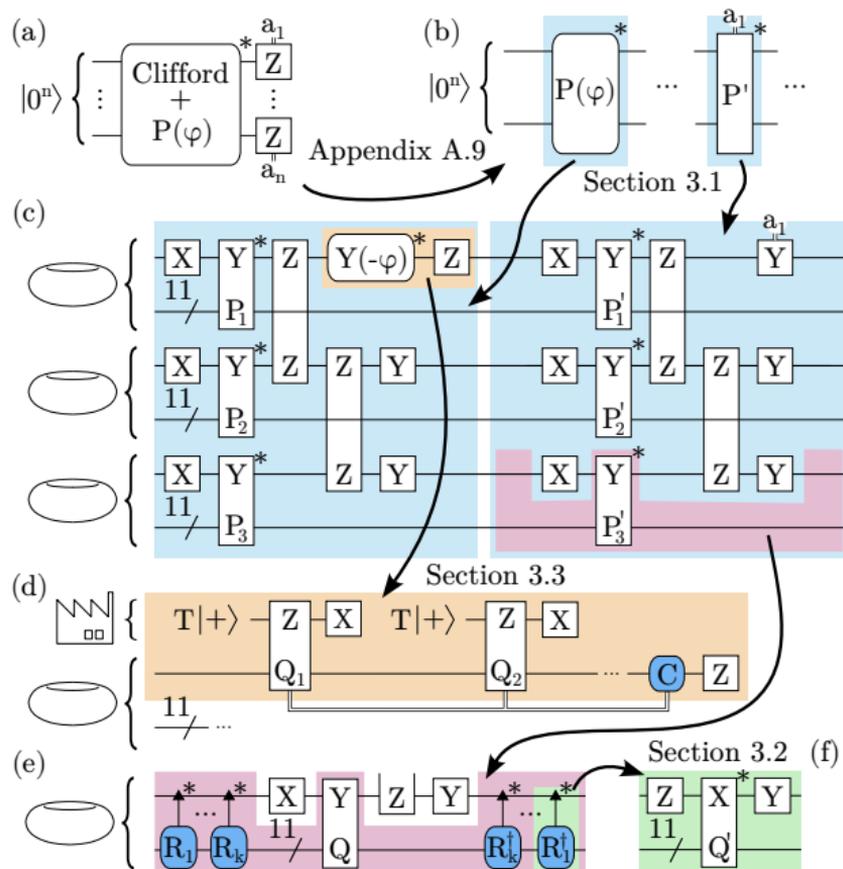
Measurement to rotation identity

For anti-commuting Paulis Q and R , it holds that

$$\text{---} \boxed{\text{R}} \text{---} \boxed{\text{Q}} \text{---} = \text{---} \boxed{\text{R}} \text{---} \boxed{\text{QR}} \text{---} \quad (15)$$

The diagram shows two equivalent quantum circuit segments. On the left, a wire passes through a rotation gate R (represented by a square box with 'R' and a parameter 'a' below it) followed by a Pauli gate Q (represented by a square box with 'Q'). On the right, the same wire passes through the rotation gate R followed by a Pauli gate QR (represented by a rounded rectangular box with 'QR'). The two segments are separated by an equals sign, and the entire equation is labeled (15) on the right.

Recap



Measuring arbitrary Paulis in a module

We are faced with implementing the circuit

The diagram shows a quantum circuit. On the left, there is a module represented by an oval with a horizontal line through its center. A large curly brace groups the circuit components. The top wire of the circuit consists of a sequence of operations: a box labeled 'X', a box labeled 'Y', a box labeled 'Z', a rounded rectangle labeled 'Y(-φ)', and another box labeled 'Z'. The bottom wire consists of a box labeled 'P₁'. A slash '/' is positioned between the two wires, indicating a measurement operation between the 'Y' and 'P₁' gates. The label '(17)' is positioned to the right of the circuit.

$$\left\{ \begin{array}{l} \text{---} \boxed{X} \text{---} \boxed{Y} \text{---} \boxed{Z} \text{---} \text{Y}(-\varphi) \text{---} \boxed{Z} \text{---} \\ \text{---} \boxed{P_1} \text{---} \end{array} \right. \quad (17)$$

The measurement $Y \otimes P_1$ is not a bicycle instruction.

Measuring arbitrary Paulis in a module

We are faced with implementing the circuit

The diagram shows a quantum circuit with two qubits. The top qubit starts with an X gate, followed by a Y gate, a Z gate, a $Y(-\varphi)$ gate, and finally a Z gate. The bottom qubit starts with a P_1 gate. A large curly brace on the left groups the X and Y gates on the top qubit and the P_1 gate on the bottom qubit. To the left of the top qubit is a symbol representing a module, which is a horizontal oval with a vertical line through its center. To the right of the circuit is the label (17).

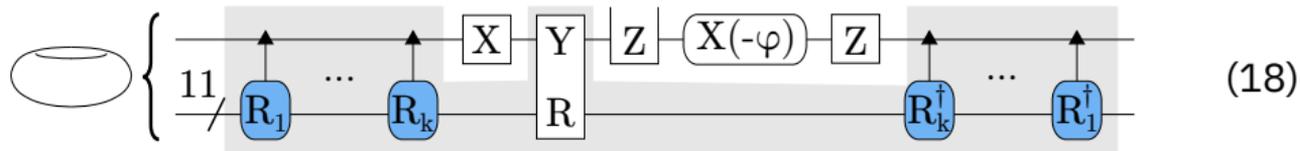
$$\left\{ \begin{array}{l} \text{X} \\ \text{11} \\ \text{Y} \\ \text{P}_1 \end{array} \right. \begin{array}{l} \text{Z} \\ \text{Y}(-\varphi) \\ \text{Z} \end{array} \quad (17)$$

The measurement $Y \otimes P_1$ is not a bicycle instruction.

But we can conjugate P_1 with some 11-qubit Clifford gate so that it becomes easy to implement [Cross, He, Rall, Yoder (2024)].

Implementing a Clifford rotation

We need at most three native Clifford rotations for any R_1 in



The diagram shows a quantum circuit with two horizontal qubit lines. On the left, a torus symbol is connected to a large curly brace. The top line starts with a 11 gate, followed by a sequence of rotation gates R_1, \dots, R_k (indicated by blue circles with upward arrows). This is followed by a X gate, a Y gate, a Z gate, a $X(-\varphi)$ gate, and another Z gate. The bottom line starts with a R gate. The final part of the circuit consists of a sequence of rotation gates $R_k^\dagger, \dots, R_1^\dagger$ (indicated by blue circles with upward arrows). The entire circuit is enclosed in a large curly brace on the left, and the label (18) is on the right.

$$\left\{ \begin{array}{c} \text{---} \\ \text{---} \end{array} \right. \left[\begin{array}{c} \text{---} \\ \text{---} \end{array} \right] \quad (18)$$

Overhead of arbitrary Pauli addressing

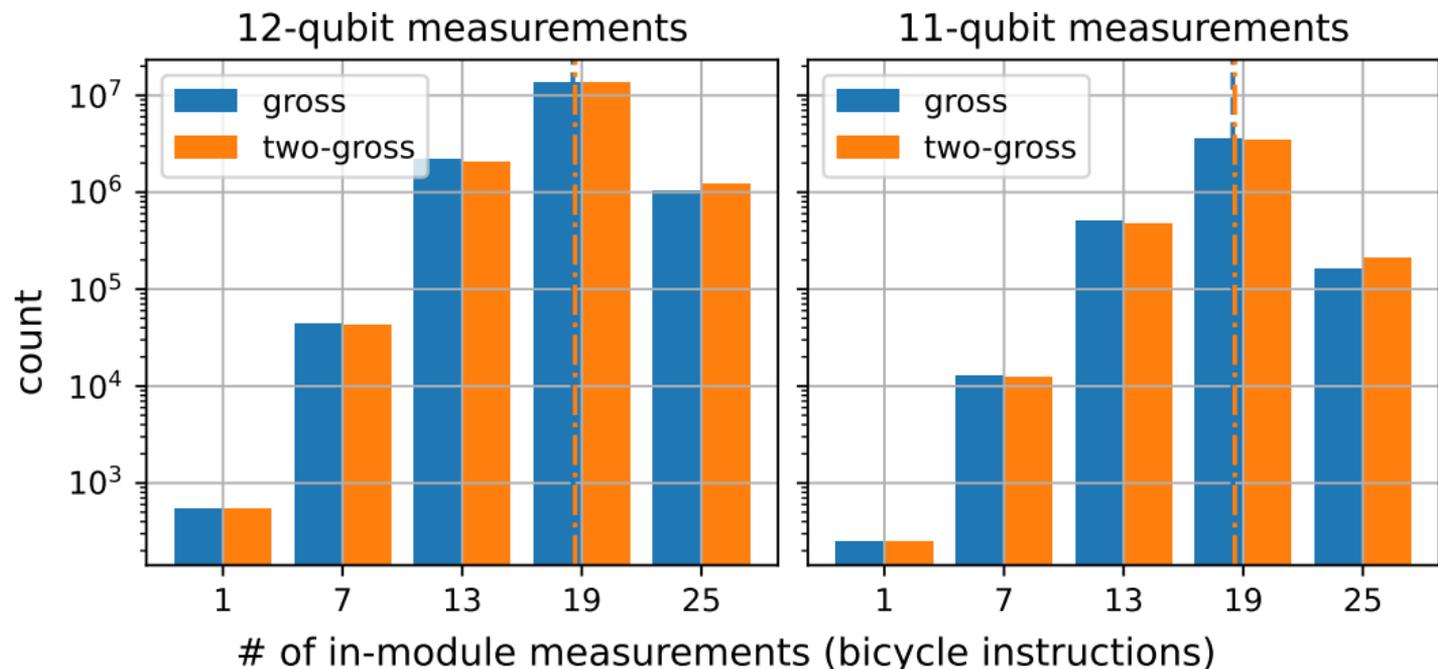


Figure: Number of bicycle measurements needed to measure the 12-qubit Paulis, $Q \otimes R$, (left) with mean values 18.55 (gross) and 18.67 (two-gross). In practice, we can minimize over Q (right) to get mean values 18.48 (gross) and 18.58 (two-gross).

Automorphism group structure

The automorphism unitaries are representations of a group. The parity matrices

$$A_x^{\text{gross}} = \begin{bmatrix} 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 1 \end{bmatrix}, \quad A_y^{\text{gross}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 1 \end{bmatrix} \quad (20)$$

are each of order 6.

Automorphism group structure

The automorphism unitaries are representations of a group. The parity matrices

$$A_x^{\text{gross}} = \begin{bmatrix} 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 1 \end{bmatrix}, \quad A_y^{\text{gross}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 1 \end{bmatrix} \quad (20)$$

are each of order 6.

This gives us $15 \times 6^2 = 540$ easy measurements and 510 easy Clifford rotations.

Automorphism group structure

The automorphism unitaries are representations of a group. The parity matrices

$$A_x^{\text{gross}} = \begin{bmatrix} 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 1 \end{bmatrix}, \quad A_y^{\text{gross}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 1 \end{bmatrix} \quad (20)$$

are each of order 6.

This gives us $15 \times 6^2 = 540$ easy measurements and 510 easy Clifford rotations.

Combinatorial open question: Reducing overhead

Given measurements $\{P_1, \dots, P_k\}$ can we map each P_j to easier measurements through $l \ll k$ easy Clifford rotations?

End-to-end resource estimates

Logical capability estimates

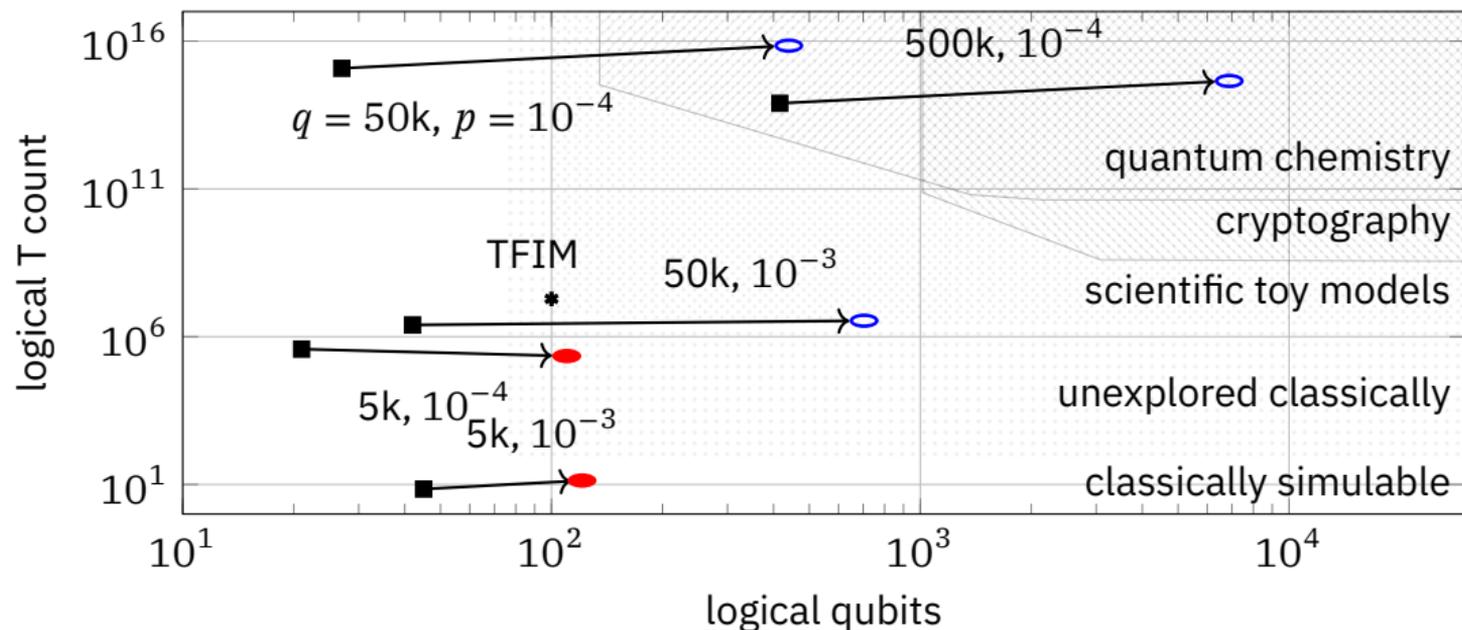


Figure: Bicycle architecture capabilities using gross (red filled ellipse) and two-gross (blue hollow ellipse) codes when compared to surface code architectures (black square) given q physical qubits and p physical error rate.

Methodology

We consider circuits of the form

$$\prod_j e^{i\frac{\theta_j}{2}Q_j}.$$

After compiling, we associate identical logical errors and timing information to all bicycle instructions and compute a simple additive failure probability,

$$\sum_j N_j P_j,$$

including any required idling.

Capability estimates: Random circuits

p	q	code	n	N_T	$\frac{\text{timesteps}}{\text{T gate}}$
10^{-3}	5k	gross	121	1.4×10^1	3367
		surface ($d = 5$)	45	7.0	381
	50k	two-gross	704	3.5×10^6	7057
		surface ($d = 17$)	42	2.5×10^6	2269
10^{-4}	5k	gross	110	2.2×10^5	3091
		surface ($d = 7$)	21	3.8×10^5	115
	50k	gross	1342	3.7×10^4	3135
		surface ($d = 7$)	250	3.2×10^4	115
		two-gross	440	7.0×10^{15}	5436
		surface ($d = 17$)	27	1.2×10^{15}	509
	500k	two-gross	6886	4.5×10^{14}	5437
		surface ($d = 17$)	416	7.9×10^{13}	509

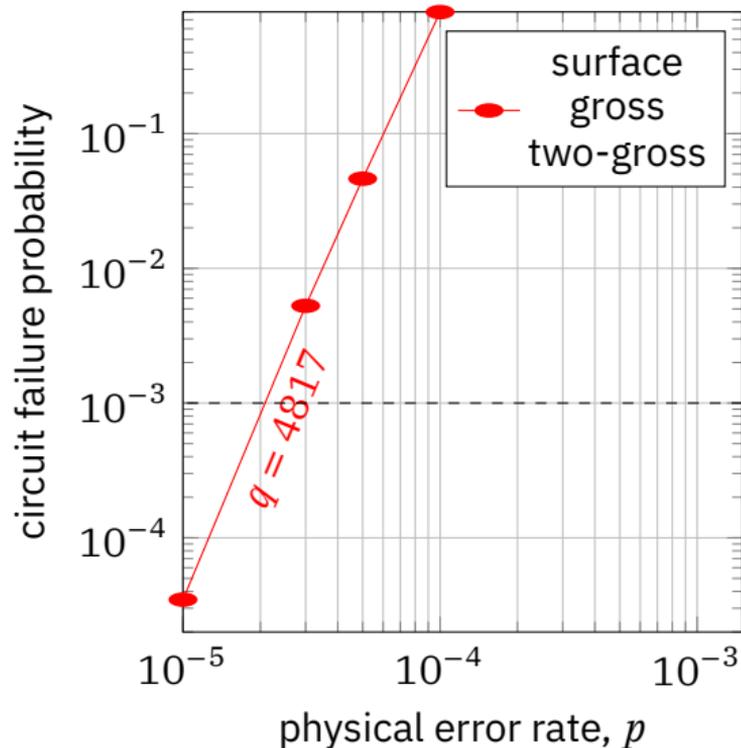
Table: The T count, N_T , estimates the number of $e^{i\frac{\pi}{8}Q}$ gates with full support on n qubits such that circuit failure probability $\leq \frac{1}{3}$.

Resource requirements for transverse-field Ising model

- ▶ We choose the same parameters for the 10×10 2D TFIM Hamiltonian as [Beverland et al. (2022)] and require $\epsilon \leq 10^{-3}$.
- ▶ The resulting circuit has 1.8×10^5 Pauli-generated rotations (of small-angle).

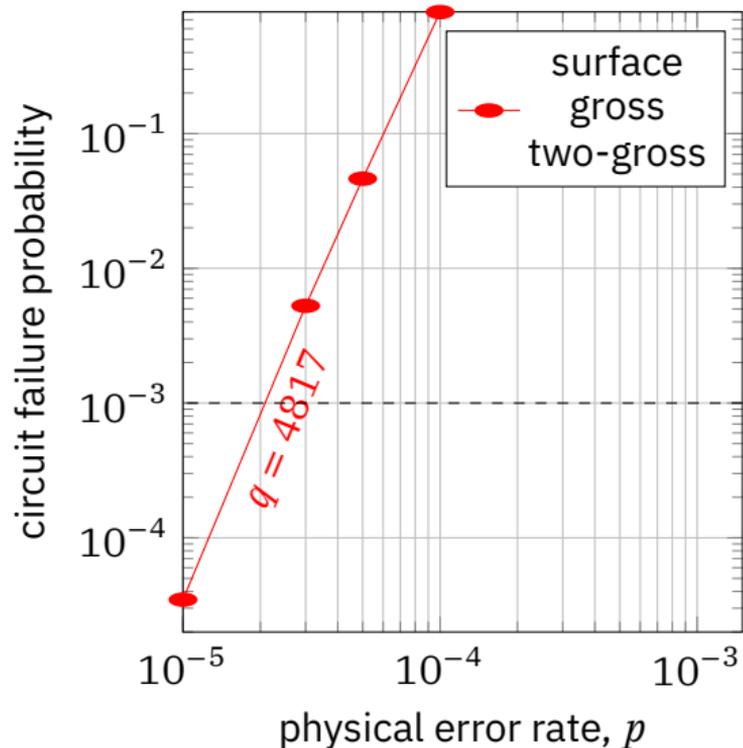
Resource requirements for transverse-field Ising model

- ▶ We choose the same parameters for the 10×10 2D TFIM Hamiltonian as [Beverland et al. (2022)] and require $\epsilon \leq 10^{-3}$.
- ▶ The resulting circuit has 1.8×10^5 Pauli-generated rotations (of small-angle).



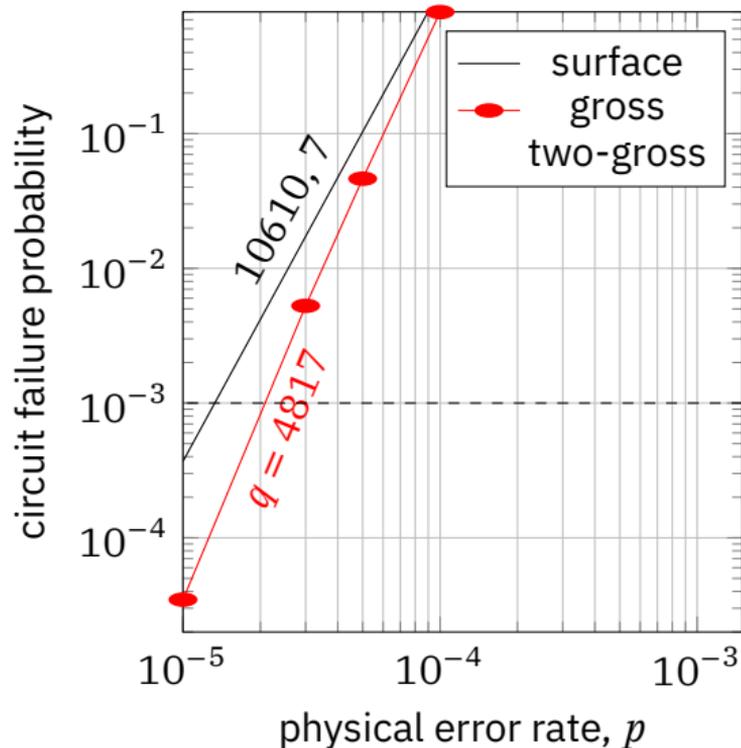
Resource requirements for transverse-field Ising model

- ▶ We choose the same parameters for the 10×10 2D TFIM Hamiltonian as [Beverland et al. (2022)] and require $\epsilon \leq 10^{-3}$.
- ▶ The resulting circuit has 1.8×10^5 Pauli-generated rotations (of small-angle).
- ▶ Gross requires $p \leq 2.1 \times 10^{-5}$



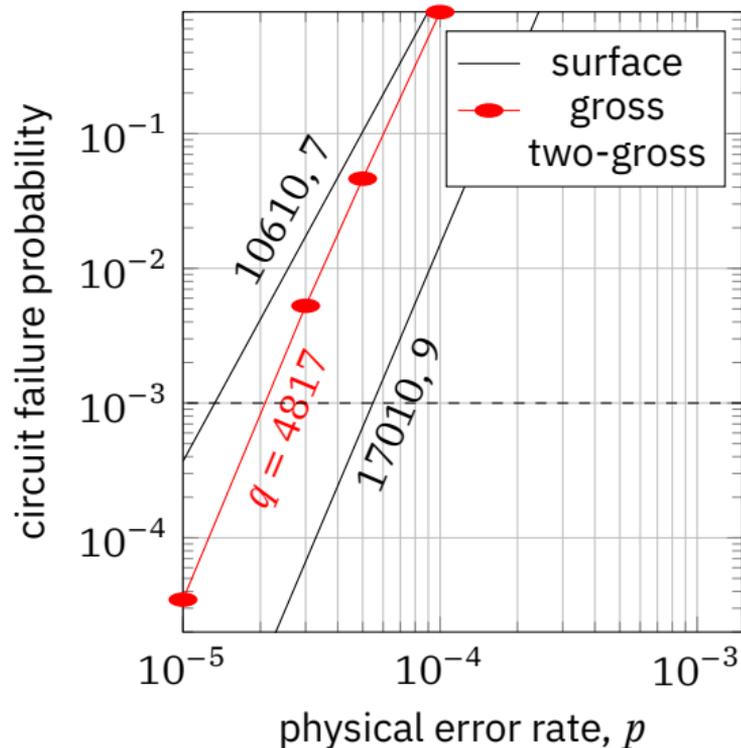
Resource requirements for transverse-field Ising model

- ▶ We choose the same parameters for the 10×10 2D TFIM Hamiltonian as [Beverland et al. (2022)] and require $\epsilon \leq 10^{-3}$.
- ▶ The resulting circuit has 1.8×10^5 Pauli-generated rotations (of small-angle).
- ▶ Gross requires $p \leq 2.1 \times 10^{-5}$



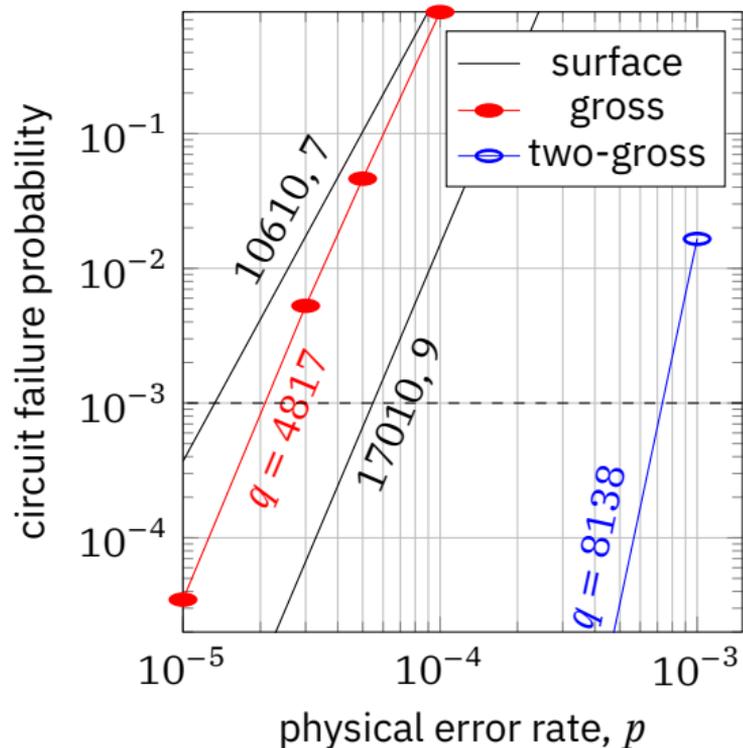
Resource requirements for transverse-field Ising model

- ▶ We choose the same parameters for the 10×10 2D TFIM Hamiltonian as [Beverland et al. (2022)] and require $\epsilon \leq 10^{-3}$.
- ▶ The resulting circuit has 1.8×10^5 Pauli-generated rotations (of small-angle).
- ▶ Gross requires $p \leq 2.1 \times 10^{-5}$



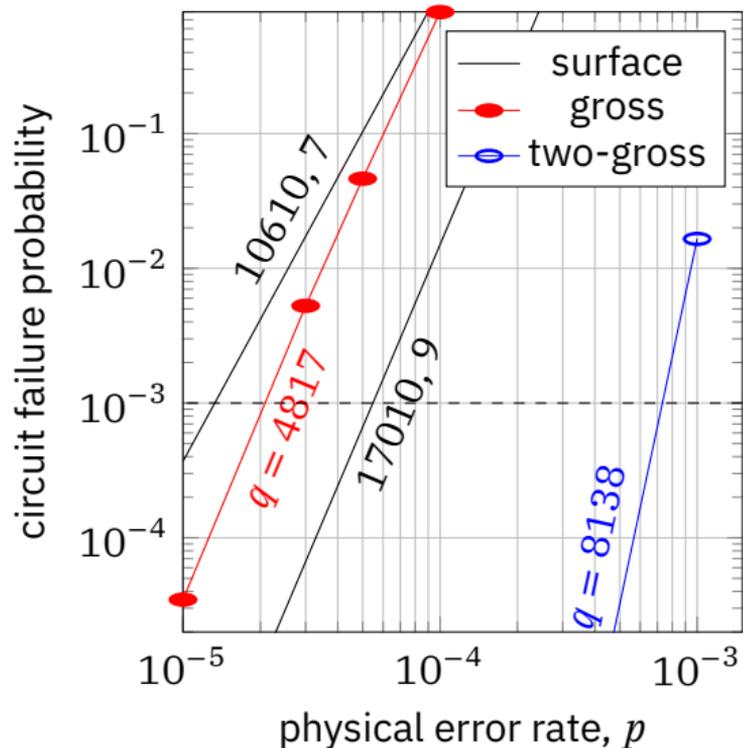
Resource requirements for transverse-field Ising model

- ▶ We choose the same parameters for the 10×10 2D TFIM Hamiltonian as [Beverland et al. (2022)] and require $\epsilon \leq 10^{-3}$.
- ▶ The resulting circuit has 1.8×10^5 Pauli-generated rotations (of small-angle).
- ▶ Gross requires $p \leq 2.1 \times 10^{-5}$



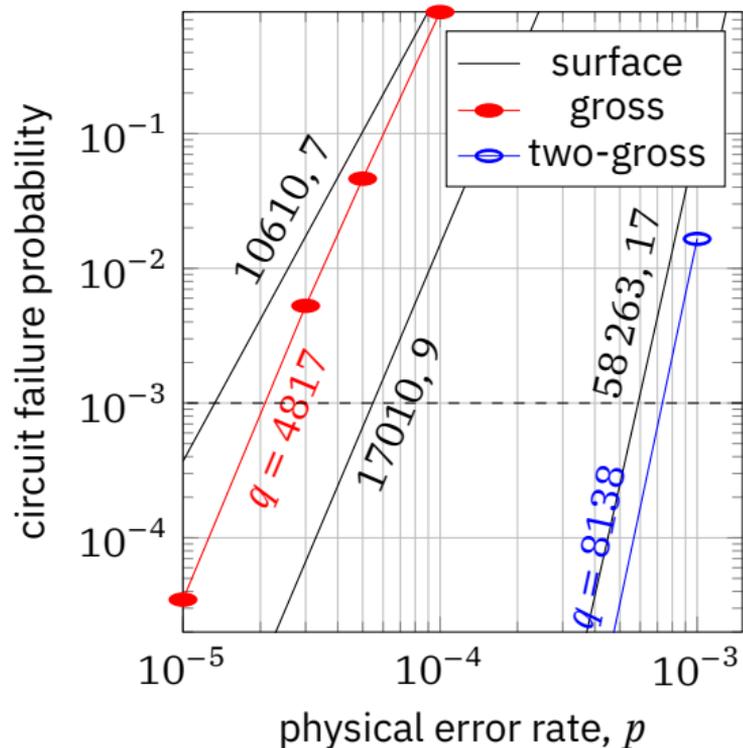
Resource requirements for transverse-field Ising model

- ▶ We choose the same parameters for the 10×10 2D TFIM Hamiltonian as [Beverland et al. (2022)] and require $\epsilon \leq 10^{-3}$.
- ▶ The resulting circuit has 1.8×10^5 Pauli-generated rotations (of small-angle).
- ▶ Gross requires $p \leq 2.1 \times 10^{-5}$
- ▶ Two-gross requires $p \leq 7.3 \times 10^{-4}$



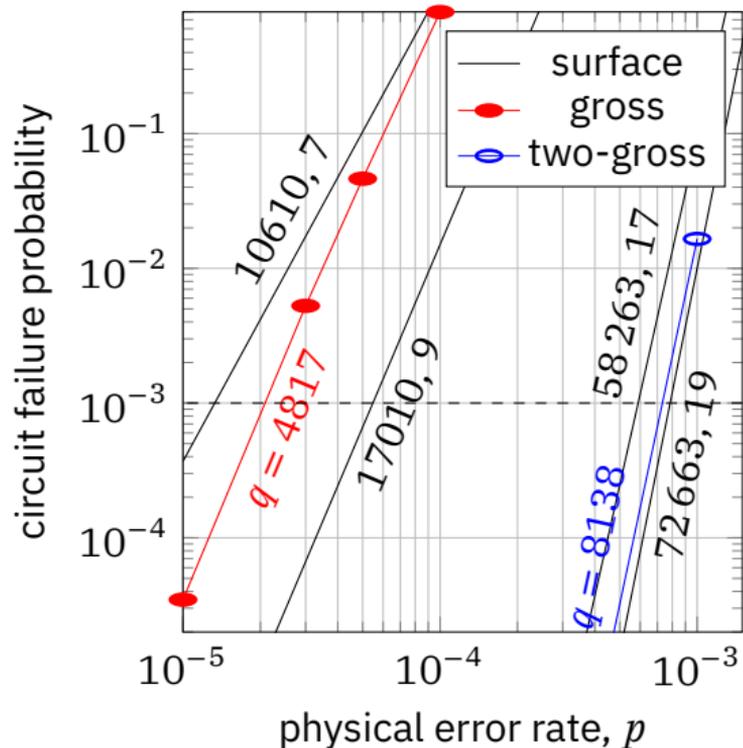
Resource requirements for transverse-field Ising model

- ▶ We choose the same parameters for the 10×10 2D TFIM Hamiltonian as [Beverland et al. (2022)] and require $\epsilon \leq 10^{-3}$.
- ▶ The resulting circuit has 1.8×10^5 Pauli-generated rotations (of small-angle).
- ▶ Gross requires $p \leq 2.1 \times 10^{-5}$
- ▶ Two-gross requires $p \leq 7.3 \times 10^{-4}$



Resource requirements for transverse-field Ising model

- ▶ We choose the same parameters for the 10×10 2D TFIM Hamiltonian as [Beverland et al. (2022)] and require $\epsilon \leq 10^{-3}$.
- ▶ The resulting circuit has 1.8×10^5 Pauli-generated rotations (of small-angle).
- ▶ Gross requires $p \leq 2.1 \times 10^{-5}$
- ▶ Two-gross requires $p \leq 7.3 \times 10^{-4}$



Conclusion

Scalable fault-tolerance criteria

The bicycle architecture satisfies our six sufficient criteria for a scalable fault-tolerant architecture:

- (i) *Fault-tolerant* – logical errors are suppressed enough for meaningful algorithms to succeed.

Scalable fault-tolerance criteria

The bicycle architecture satisfies our six sufficient criteria for a scalable fault-tolerant architecture:

- (i) *Fault-tolerant* – logical errors are suppressed enough for meaningful algorithms to succeed.
- (ii) *Addressable* – individual logical qubits can be prepared or measured throughout the computation.

Scalable fault-tolerance criteria

The bicycle architecture satisfies our six sufficient criteria for a scalable fault-tolerant architecture:

- (i) *Fault-tolerant* – logical errors are suppressed enough for meaningful algorithms to succeed.
- (ii) *Addressable* – individual logical qubits can be prepared or measured throughout the computation.
- (iii) *Universal* – a universal set of quantum instructions can be applied to the logical qubits.

Scalable fault-tolerance criteria

The bicycle architecture satisfies our six sufficient criteria for a scalable fault-tolerant architecture:

- (i) *Fault-tolerant* – logical errors are suppressed enough for meaningful algorithms to succeed.
- (ii) *Addressable* – individual logical qubits can be prepared or measured throughout the computation.
- (iii) *Universal* – a universal set of quantum instructions can be applied to the logical qubits.
- (iv) *Adaptive* – measurements are real-time decoded and can alter subsequent quantum instructions.

Scalable fault-tolerance criteria

The bicycle architecture satisfies our six sufficient criteria for a scalable fault-tolerant architecture:

- (i) *Fault-tolerant* – logical errors are suppressed enough for meaningful algorithms to succeed.
- (ii) *Addressable* – individual logical qubits can be prepared or measured throughout the computation.
- (iii) *Universal* – a universal set of quantum instructions can be applied to the logical qubits.
- (iv) *Adaptive* – measurements are real-time decoded and can alter subsequent quantum instructions.
- (v) *Modular* – the hardware is distributed across a set of replaceable modules connected quantumly.

Scalable fault-tolerance criteria

The bicycle architecture satisfies our six sufficient criteria for a scalable fault-tolerant architecture:

- (i) *Fault-tolerant* – logical errors are suppressed enough for meaningful algorithms to succeed.
- (ii) *Addressable* – individual logical qubits can be prepared or measured throughout the computation.
- (iii) *Universal* – a universal set of quantum instructions can be applied to the logical qubits.
- (iv) *Adaptive* – measurements are real-time decoded and can alter subsequent quantum instructions.
- (v) *Modular* – the hardware is distributed across a set of replaceable modules connected quantumly.
- (vi) *Efficient* – meaningful algorithms can be executed with reasonable physical resources.

Opportunities for improvement

Reduce the time overhead The synthesis of arbitrary Pauli measurements incurs significant overhead.

Opportunities for improvement

Reduce the time overhead The synthesis of arbitrary Pauli measurements incurs significant overhead.

Increase parallelism The Pauli-based compilation scheme leads to sequential T gate execution.

Opportunities for improvement

Reduce the time overhead The synthesis of arbitrary Pauli measurements incurs significant overhead.

Increase parallelism The Pauli-based compilation scheme leads to sequential T gate execution.

Decoding with speed and accuracy For instance, decoding speed and error rates of inter-module measurement [Maurya et al. (2025)].

Opportunities for improvement

Reduce the time overhead The synthesis of arbitrary Pauli measurements incurs significant overhead.

Increase parallelism The Pauli-based compilation scheme leads to sequential T gate execution.

Decoding with speed and accuracy For instance, decoding speed and error rates of inter-module measurement [Maurya et al. (2025)].

Increasing code and circuit distances With long-range coupling, can consider a wide variety of qLDPC codes.

Opportunities for improvement

Reduce the time overhead The synthesis of arbitrary Pauli measurements incurs significant overhead.

Increase parallelism The Pauli-based compilation scheme leads to sequential T gate execution.

Decoding with speed and accuracy For instance, decoding speed and error rates of inter-module measurement [Maurya et al. (2025)].

Increasing code and circuit distances With long-range coupling, can consider a wide variety of qLDPC codes.

Improved magic state factories Similarly, long range connectivity can be used for better magic state factories [Jacob et al. (2025); Li et al. (2025); Guanyu Zhu (2025)].

Opportunities for improvement

Reduce the time overhead The synthesis of arbitrary Pauli measurements incurs significant overhead.

Increase parallelism The Pauli-based compilation scheme leads to sequential T gate execution.

Decoding with speed and accuracy For instance, decoding speed and error rates of inter-module measurement [Maurya et al. (2025)].

Increasing code and circuit distances With long-range coupling, can consider a wide variety of qLDPC codes.

Improved magic state factories Similarly, long range connectivity can be used for better magic state factories [Jacob et al. (2025); Li et al. (2025); Guanyu Zhu (2025)].

Modifying the bicycle instructions Simplify circuits for bicycle instructions, especially for complicated and error-prone instructions.