A Heterogeneous Distributed Architecture for Quantum Simulation

John Stack[1,2,*], Sitong Liu[3,2], Abhinav Anand[2], Inder Monga[2], Yuan Liu[1,5,6],
Erhan Saglamyurek[2], Pedro L. S. Lopes[4,*], Frank Mueller[1,*], and Katherine Klymko[2,*]
[1] Department of Computer Science, North Carolina State University, Raleigh, North Carolina, USA
[2] Lawrence Berkeley National Laboratory, Berkeley, California, USA
[3] Department of Electrical and Computer Engineering, Duke University, Durham, North Carolina, USA
[4] QuEra Computing Inc., Boston, Massachusetts, USA
[5] Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, North Carolina, USA
[6] Department of Physics and Astronomy, North Carolina State University, Raleigh, North Carolina, USA
[*]Corresponding authors: jstack@ncsu.edu, plopes@quera.com, fmuelle@ncsu.edu, kklymko@lbl.gov

***Abstract*—Architectural specialization and distribution can help scale fault-tolerant quantum computers, but may also introduce substantial overheads from communication, routing, and resource duplication. We introduce a heterogeneous distributed architecture in which a magic core is connected to an extensible storage system composed of one-dimensional lanes of specialized cold-storage nodes. The storage system supports parallel random access to Pauli string parities.**

This organization is particularly well suited to fermionic quantum simulation, enabling parallel execution of the highly non-local Pauli strings arising from these systems. We evaluate the architecture on fault-tolerant simulations of the dynamics of the Fermi–Hubbard and sparse Sachdev–Ye–Kitaev (SYK) models on systems of up to 450 logical qubits. These workloads exhibit complementary communication structures: Fermi–Hubbard produces a spectrum of interactions from local to non-local shaped by lattice geometry, whereas sparse SYK produces highly non-local and overlapping Pauli operators. For a Trotter step of a 450-logical-qubit Fermi–Hubbard workload, a six-lane system with 30 T-state factories is within approximately $1.4\times$ the wall-clock time of a homogeneous distributed architecture with 4 times as many T-state factories and substantially greater connectivity and sites for injecting magic. For matched T-factory counts, our architecture is $\sim 2\times$ faster.

I. Introduction

Architectural specialization and distribution are promising approaches for scaling fault-tolerant quantum computers. By partitioning computation across modules, a distributed architecture reduces the burden on individual devices and allows memory, computation and communication to be optimized independently. A natural design philosophy [14], [21], [24], [36], [46], [51] is to equip each module with tightly integrated memory, computation and communication resources, maximizing local capability while minimizing the need for remote operations. However, this organization duplicates expensive resources across the system and can introduce substantial overhead in control, routing, and scheduling. These overheads are particularly significant in early fault-tolerant systems, where distributed communication via Bell pairs may be slow and expensive, making contention for these operations a significant performance bottleneck.

At the same time, fault-tolerant [17], [18], [39] workloads impose additional architectural requirements. Non-Clifford resources [11] must be available where and when they are needed, and the qubit interactions induced by the compiled algorithm must be implemented efficiently.

Many quantum simulation workloads [5], [27] exhibit an asymmetric structure in which parity is accumulated using Clifford operations across a large set of qubits (the support of a Pauli string) but non-Clifford resources are applied only to one of these qubits. This is analogous to a classical workload with high "arithmetic/operational intensity" [48], which refers to the number of floating-point operations per byte accessed from memory. In the architecture we introduce we can describe quantum simulation tasks as having high magic intensity: The amount of T-gates applied is much higher than the number of accesses from the extensible storage system. The effectiveness of a distributed architecture therefore depends not only on the underlying hardware, but also on how well its organization matches these workload characteristics.

Quantum simulation of electronic systems [4], [15], [31], a task where quantum computers are used to simulate the properties of quantum systems including molecules, strongly correlated materials, and lattice models of interacting fermions [12], [33], is a particularly useful setting in which to study these trade-offs. The communication structure of these workloads is shaped not only by lattice symmetries and flavor couplings, but also by the chosen fermion-to-qubit mapping. For example, the Jordan–Wigner [23] mapping produces geometrically structured but high-weight Pauli strings, whereas Bravyi–Kitaev [8] yields lighter-weight strings with more non-local and overlapping supports. The suitability of a distributed architecture therefore depends not only on problem size, but also on how the chosen mapping translates fermionic structure into communication, movement, and parallelism requirements.

Here, we introduce a heterogeneous distributed architecture designed to exploit this structure while maintaining a simple physical organization. The architecture combines a single magic core with one-dimensional chains of cold-storage

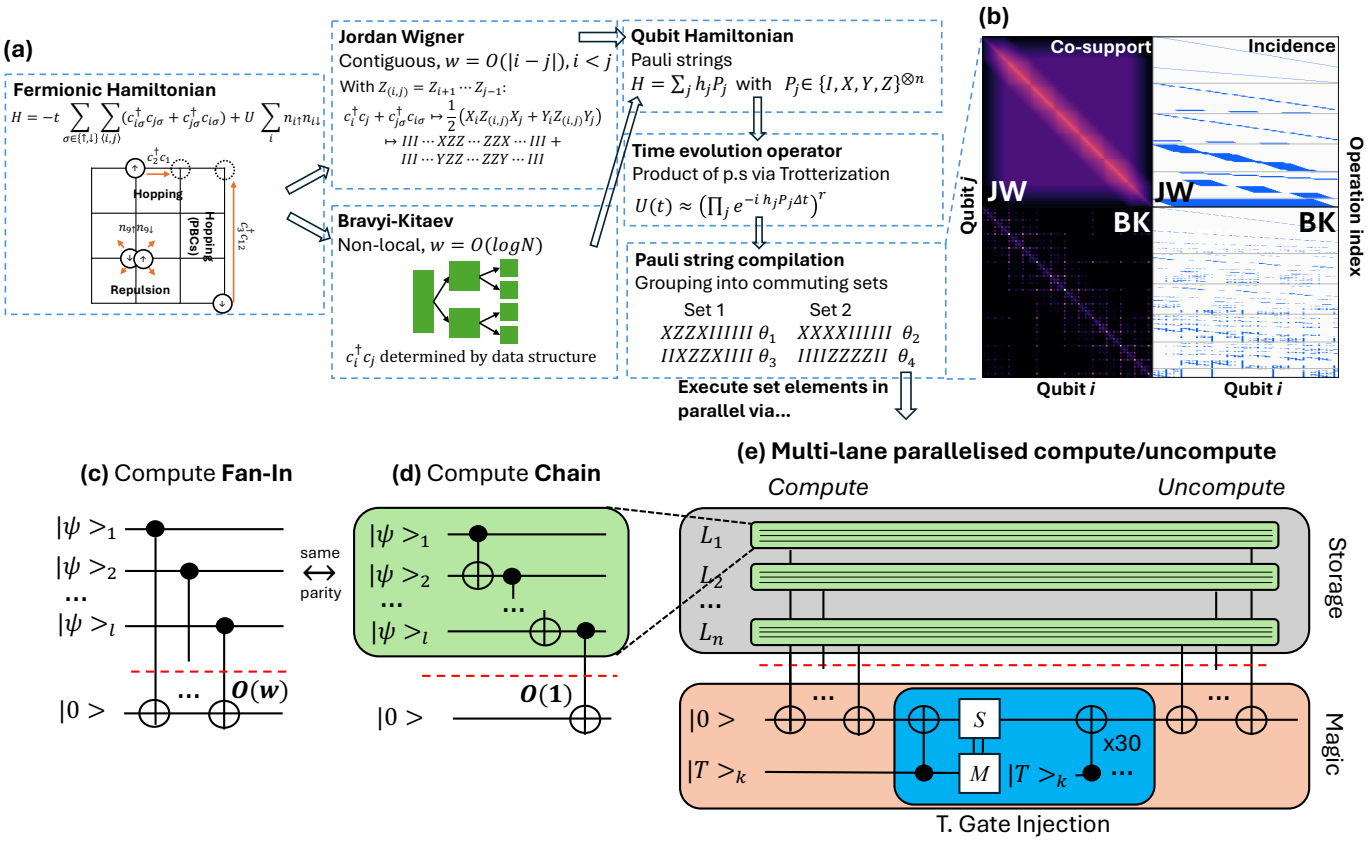

Fig. 1: **Simulation workflow, Pauli-string structure and device implementation.** **(a)** Simulation workflow. **(b)** Co-support and incidence matrices for the Jordan–Wigner (JW, top) and Bravyi–Kitaev (BK, bottom) transformations on a 7×7 Fermi-Hubbard Hamiltonian. The co-support matrices indicate the probability that qubits i and j both occur in the support of a Pauli string. The incidence matrices indicate which qubits occur in the support of each Pauli string. **(c)** A method of computing parity that is suboptimal on distributed devices, requiring $O(w)$ non-local CNOTs for a Pauli string of weight w. **(d)** A communication-efficient method of computing parity. For each lane on which a Pauli string has support, only one non-local CNOT is required between the lane and the compute node. **(e)** Logical circuit overview of the implementation of Pauli rotations on the distributed device. Single qubit Clifford gates applied during rotation synthesis omitted for clarity.

modules, separating magic-state generation and consumption from logical data storage. This organization eliminates the need for a complex communication network, reduces the control and calibration burden on storage modules, simplifies routing and scheduling, exposes parallelism across independent storage lanes, and enables the machine to scale by extending simpler storage rather than replicating full execution capability. We note that this architecture can also be realized within a single device. We evaluate the architecture on two physically relevant applications, compiled single-band square lattice Fermi–Hubbard [19] and sparse Sachdev–Ye–Kitaev (SYK) [49] simulation workloads, which span communication patterns ranging from structured local interactions to highly non-local overlapping Pauli strings. We show that the proposed architecture supports efficient execution of these workloads with experimentally modest module sizes and entanglement-generation rates. We also evaluate the performance of the extensible storage system on random Pauli strings with varying degrees of non-locality. More broadly, our results demonstrate that useful distributed fault-tolerant quantum simulation does not require resource-intensive highly-connected homogeneous modules or complex network fabrics. Instead, efficient execution can be achieved by matching the architecture to the communication structure of the compiled workload.

II. Distributed Architectures and Quantum Simulation

We consider a local Fermi-Hubbard model [19] on a square lattice with periodic boundary conditions and a non-local sparse SYK model [49], both ranging in size from 98 to 450 qubits (see Table II for all sizes considered). These models have broad interest across condensed matter physics, high-energy theory, and quantum information science [10]. The fermionic Hamiltonians are mapped to qubit operators using either the Jordan–Wigner (JW) or Bravyi–Kitaev (BK) transformations (see Appendix X-B for details) [8], [23], and the time evolution operator is implemented using a first-order Trotter-Suzuki [44] decomposition, resulting in a sequence of exponentiated Pauli strings that can be grouped into commuting sets. See Fig. 1(a) for a summary of this workflow. As

illustrated in Fig. 1(e), executing exponentiated Pauli strings on a quantum computer involves a sequence of two operations: Pauli string compute (uncompute) operations and associated Pauli rotation [15]. Compute in this case means using CNOTs to map the parity of qubits supported by an operator onto an ancilla qubit to perform a rotation. Parities can be computed in many ways, some of which are more or less amenable to distributed settings (compare Fig. 1(c) and (d)). Uncompute disentangles the qubits from the ancilla after the rotation.

The choice of fermion-to-qubit mapping strongly influences the architectural requirements. The JW transformation generates nearest-neighbor local qubit operators that are often disjoint and thus readily parallelizable [32]. However, the weight of the resulting Pauli strings scales with $O(N)$, where N is the number of fermionic modes. By contrast, the BK transformation generates significantly lower weight Pauli strings with $O(\log N)$ scaling, but at the cost of introducing highly non-local Pauli strings, usually with overlapping supports.

These differences produce substantial architectural consequences even at modest system sizes. For the 98-qubit Fermi–Hubbard instance the average BK Pauli weight (i.e., the non-identity gates in a Pauli string) is approximately 5.7, compared with approximately 11.1 under JW; at 450 qubits, the corresponding averages are approximately 6.6 and 22.8. The benefit of BK therefore grows with problem size, provided that the architecture can efficiently execute its less regular and more spatially distributed supports. Fig. 1(b) illustrates these differences: the co-support matrix of JW is centered around the diagonal, indicating that neighboring strings feature qubits close to each other. At the same time, the matrix is densely populated due to the presence of long Pauli strings that involve almost all qubits. For the case of the Fermi-Hubbard model, the longest strings come from the periodic boundary conditions (PBCs). This is because fermionic nearest-neighbor at the lattice boundaries cross the entire lattice (they are wrapped around) and in JW weight is proportional to the distance between sites in the lattice (assuming standard ordering).

In contrast, the BK cosupport matrix is substantially sparser but with significant cosupport between distant qubits. This is because BK represents these long-range hopping terms with a tree (typically Fenwick [13]) encoding, which saves weight but produces non-locality. The corresponding incidence matrices further highlight the difference in regularity. JW produces highly regular lines of support, whereas BK generates sparse and irregular support.

Historically, the JW transformation has been preferred for fault-tolerant simulation because of the difficulty of accessing and interacting highly non-local supports [32], [36]. However, the long Pauli strings generated by the JW transformation can be expensive to execute, incurring large movement costs for larger systems and more complex fermionic models, especially on distributed architectures.

Conversely, although BK substantially reduces the Pauli-string weights, the high communication costs (evaluated in Sec. VI) associated with non-local operators with supports spread across module(s) can offset these benefits, especially

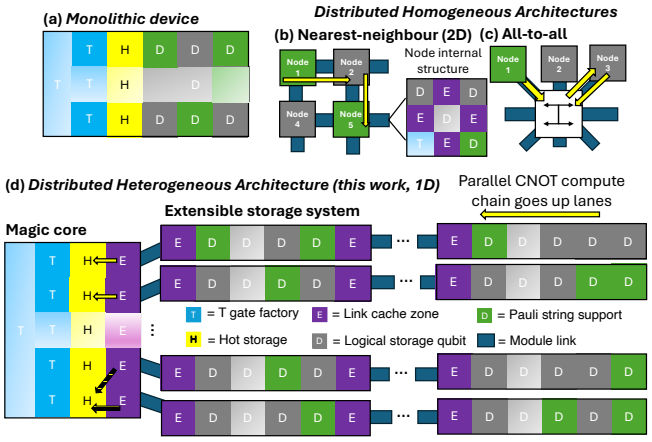

Fig. 2: **Quantum architectures.** Gradient coloring represents design space for sizing. **(a)** Monolithic device with column(s) of T factories (T), hot-storage zones (H), connected via ebits (E) to data (D) / cold storage qubits. **(b-c)** Distributed homogeneous architectures. Yellow lines indicate node-to-node interactions required to implement distributed Pauli strings. Example internal homogeneous node structure shown. **(d)** Distributed heterogeneous architecture. Yellow and cross-hatched yellow arrows at the Magic core indicate inter- and intra-string parallelism, respectively. Note that while nodes are shown as squares or 1D rectangles, this is just illustrative. We model nodes as rectangles with a standard aspect ratio.

in distributed architectures where different and typically more expensive modalities (Bell pairs, shuttling, long movements, swaps, etc) are needed to interact with distant qubits. This tradeoff motivates the development of architectures that can efficiently support non-local Pauli strings without incurring prohibitive communication overheads.

To address this challenge, we propose a heterogeneous lane architecture shown in Fig. 2(d). The architecture consists of a centralized magic core connected to multiple cold-storage lanes. The magic core hosts the resources required for active execution, including T-state factories (T), hot-storage zones (H) for frequently accessed logical qubits, and a communication zone for caching Bell pairs (E) and for connection to cold-storage lanes consisting of modules which hold data (logical) qubits (D). The cold-storage lanes contain the majority of logical data and are extended at low relative cost as required to accommodate larger problem instances, without increasing network/bus complexity. This organization separates active computation from long-lived logical storage and allows each component of the system to be optimized for its specific role.

A key advantage of this design is that communication contention is largely eliminated: each lane in the heterogeneous architecture is dedicated to the Pauli-string task currently being executed. Consequently, the communication cost is largely independent of the spatial distribution of the Pauli-string support. For comparison, we also evaluate a homogeneous distributed architecture (Fig. 2(b), the “2D design”), which requires substantially greater communication resources. While all-to-all or highly connected approaches (Fig. 2(c)) have

been proposed [34], these have high experimental costs, with approaches such as including a switch in the optical generation of Bell pairs generating infidelities [40], [52]. Maintaining Bell pair caches in an all-to-all architecture would mean that each device must have $O(N_{nodes})$ Bell pair caches, which would take up a significant amount of space.

Another key consideration in our architectural model is that magic state generation is substantially more resource-intensive than logical memory. Devices that generate magic states have significantly greater complexity and usage requirements, as they must continuously produce magic states during computation. This means that a substantial fraction of the device must be dedicated to T-state factory operation, limiting the resources available for other functions. In current neutral-atom platforms, for example, the number of Acousto-Optical Deflectors (AODs) is limited, restricting the degree of parallelism that can be achieved while requiring a significant fraction of the available control resources to be devoted to T-state production. By separating magic-state generation from logical storage, the proposed architecture avoids replicating these expensive resources across every module, enabling the system to scale by adding low-cost storage modules rather than fully featured compute nodes.

In what follows we will assume logical information to be encoded using surface-code patches [30], with logical operations implemented transversally [39] and via magic state injection. This implementation is well suited to reconfigurable neutral-atom arrays [6] and trapped-ion devices [37], where logical patches can be moved and coupled efficiently. However, the extensive encoding overhead from surface codes lead to different challenges in trapping and routing logical qubits across long distances; in this case, having smaller nodes and distributed logic comes again as an advantage.

For the quantitative evaluation in this work, we consider a neutral-atom implementation and adopt standard experimental device parameters from the literature [53], shown in Tab. I. These values are chosen to reflect current experimental capabilities rather than to optimize performance for the proposed architecture.

Parameter	Value
T-gate generation speed per factory	$10\,\mathrm{ms}$
T-gates per rotation	30
Bell pair generation speed	$1\,\mathrm{ms}$
Bell pair parallelism	$O(d)$
Bell pair cache size	$n = d^2$
Neutral atom movement acceleration	$5500\,\mathrm{m\,s^{-2}}$

TABLE I: **Parameters used in the architecture model.**

III. Related work

In fault-tolerant architecture design, the idea of specialized (heterogeneous) sections of a device is not new. These sections may use different QEC codes for different purposes (cold storage or using magic) [16], [42] or use identical QEC codes but with dedicated sections for minimizing contention and routing costs [27]. The idea of distributing devices and designing around local and non-local connections is also not new, or even recent [9], [18], [46]. Further, many fault-tolerant architectures involve some form of parallel processing with multiple sources of magic [47], [53].

However, prior work generally considers homogeneous distribution, where the same device is replicated multiple times and interconnected with a simple [14], [21], [36], [46] or complex [24], [51] network. The distinction of our work is that our architecture is heterogeneous but also distributed and extensible, and we show that our design performs better than a simple homogeneous distributed device for matched T factory count. The fact that our device is distributed means that it can be arbitrarily extended to match the size of the workload executed. The heterogeneity and simple network design ensures the extensibility is low-cost: Less-complex storage nodes are simply added to the chain. A recent work has started exploring heterogeneous distributed architectures [35]. But their proposal is less specialized than our architecture and relies on a centralized quantum bus to provide all-to-all logical connectivity.

Prior work has also considered using reconfigurable architectures and dynamic encoding changes to perform fermionic simulation efficiently [1], [28], [32]. These approaches are very promising for monolithic devices and discuss compatibility with fault-tolerant execution, but the dynamic encoding changes come with costs. The low-depth approach in [32] uses $O(N)$ ancilla qubits, the strongest results in [1] apply to geometrically local or bounded-range fermionic systems, and the arbitrary fermionic permutations in [28] require $O(\sqrt{N})$ depth and $O(N\sqrt{N})$ gates. However, none of these explicitly model a complete and distributed fault-tolerant architecture with magic-state production, modular communication, and contention. Furthermore, the required global reconfiguration or routing is natural within a monolithic device, but would require inter-module communication if the data were distributed across separate modules, a cost that these works do not consider. Our work instead takes the compiled Pauli strings as given and considers how to execute their local and non-local supports on an extensible distributed architecture. Explicit architectural studies of dynamic encodings on distributed devices are a promising direction for future work.

IV. Evaluation Methodology

A. Trotterization

After the fermion-to-qubit transformation, each Hamiltonian is expressed as a sum of Pauli operators,

$$H = \sum_{j=1}^{M} h_j P_j, \tag{1}$$

where h_j are real coefficients and P_j are tensor products of Pauli operators, also referred to as Pauli-strings. M is the total number of Pauli strings per Trotter step. The time-evolution operator for evolution time t is $U(t) = e^{-iHt}$.

We approximate the evolution using a first-order Trotter-Suzuki decomposition [44]. Dividing the total evolution time

into $r = 10$ equal steps of duration $\Delta t = t/r$ gives

$$U(t) = e^{-iHt} \approx \left[\prod_{j=1}^{M} e^{-ih_j P_j \Delta t} \right]^r + O(t^2/r). \quad (2)$$

Every exponential term $e^{-ih_j P_j \Delta t}$ corresponds to a Pauli rotation and is implemented using the standard compute-rotate-uncompute procedure. Pauli strings produced after encoding fermions into qubits are a mix of Pauli operators on different qubits, so single-qubit Clifford gates are first used to put each qubit in the support of P_j into the Z basis and then the parity of these qubits is accumulated onto an ancilla qubit using a sequence of CNOT gates. A single-qubit rotation with angle $\theta_j = 2h_j \Delta t$ is then applied to the ancilla, after which the parity computation is reversed (uncompute).

We assume that every single-qubit rotation needs 30 T-gates, regardless of architecture. Based on recent work [7], this is sufficient for a rotation error of 10^{-6} per rotation and any angle, where 40 T-gates suffice for a rotation error of 10^{-8}. We ignore the cost of the single-qubit Clifford gates as they are negligible compared to the other costs in this paper.

All the results and reported execution times correspond to a evolution with a single Trotter step ($r = 1$ in Eq. 2). However, for calculating the distance estimates for the surface code, we consider the total evolution consisting of $r = 10$ Trotter steps (see Appendix X-C for details).

B. Architectures

We compare our heterogeneous lane architecture with two baseline architectures: a monolithic architecture (Fig. 2(a)) and a homogeneous distributed architecture (Fig. 2(b)). As mentioned earlier, we assume the use of neutral atom device(s) and use the parameters shown in Table I. For a movement of a code block with code distance d along dimensionless length ℓ, we assign a physical distance $L_{\text{phys}} = \ell d\,(12\,\mu\text{m})$ and movement time

$$\tau_{\text{move}} = 2\sqrt{\frac{L_{\text{phys}}}{a}}, \quad (3)$$

where, $a = 5500\,\text{m}\,\text{s}^{-2}$ is the acceleration and $12\,\mu$m comes from the atom spacing [53]. The same movement model is used for both the heterogeneous architecture and the homogeneous 2D baseline.

1) Heterogeneous lane architecture (1D): Logical data qubits are divided into lanes. For example, an architecture with 4 lanes and 100 data qubits would have qubit q_0 to q_{24} on the zeroth lane, q_{25} to q_{49} on the next lane and so on. Each storage lane is then partitioned into cold-storage modules. Within each lane, qubits fill consecutive storage modules up to the 20-qubit capacity. If the number assigned to a lane is not divisible by 20, its final module contains only the remaining qubits. Unless otherwise stated, in this work this architecture uses six storage lanes, six compute rows, and five T-factory columns ($6 \times 5 = 30$ T factories in total).

Internal CNOTs between two qubits in the same module are implemented by moving the control qubit to the target and back, giving a total travel distance equal to twice the Manhattan separation between the qubits. Logical operations involving qubits located in different storage modules require non-local CNOT operations, carried out using cached Bell pairs. If the two supports on a lane are separated by multiple modules, then entanglement swapping is used to create a Bell pair link via the intermediate modules.

For a Pauli string, all local CNOTs required to accumulate parity across its supports are performed simultaneously across the storage modules. Then non-local CNOTs are used to accumulate the parities across a lane and feed it into the Magic core.

Pauli rotations are implemented using a compute-rotate-uncompute procedure. For a Pauli string P, the support of P is first partitioned according to storage lane. Within each occupied lane, parity is accumulated locally and transferred to the magic core using a non-local logical CNOT operation. The magic core combines the parities arriving from all occupied lanes and performs the required Pauli rotation. The compute step is then applied again but in reverse, which forms the uncompute step. Consequently, both local and non-local Pauli strings are executed using the same architectural primitive: lane-local parity accumulation followed by centralized rotation.

2) Monolithic architecture: The monolithic baseline represents an idealized monolithic quantum computer. It is the heterogeneous lane (1D) architecture but without communication costs. This is implemented in the simulator by giving the 1D architecture instantaneous Bell pairs. Consequently, the monolithic architecture provides a lower-bound for execution time, isolating the effects of available parity-computation resources and T-state production from communication constraints.

3) Homogeneous distributed architecture (2D): The homogeneous distributed device consists of a nearest-neighbor two-dimensional grid of identical modules. Each module contains up to 20 logical data qubits and 2-4 Bell pair/link zones. The number of local T-state factories is varied in this work. For a problem containing Q logical data qubits, the minimum number of modules required by storage capacity is $m = \left\lceil \frac{Q}{20} \right\rceil$. We instantiate a complete near-square rectangular grid with

$$X = \left\lceil \sqrt{m} \right\rceil, \qquad Y = \left\lceil \frac{m}{X} \right\rceil, \qquad N_{\text{grid}} = XY.$$

Thus, N_{grid} may be slightly larger than m. The Q data qubits are distributed across all N_{grid} modules in snake order.

For the 2D device we use a snake schedule, qubits on modules that are the support of Pauli string are first combined locally in parallel. Non-local CNOTs across the supporting modules are then used to finish the compute step. We explored other routing methods but they increased the number of hops between modules, increasing wall-clock time due to communication cost.

The T-state operation is performed at the module of the ancilla qubit and the compute process is then repeated to uncompute the parity. All support and transit modules used by the network/bus remain reserved throughout this process.

Model	L	n_q	Groups	Strings	W_{BK}	W_{JW}	d_{BK}	d_{JW}
FH	7	98	6	539	3080	5964	12	12
FH	10	200	5	1100	6630	17040	12	13
FH	12	288	5	1584	9816	29184	13	13
FH	15	450	6	2475	16254	56460	13	14
SYK	-	98	7	129	2028	5394	10	11
SYK	-	200	9	294	5713	24366	11	13
SYK	-	288	11	419	8873	50335	12	13
SYK	-	450	9	651	15589	117871	12	14

TABLE II: **Quantum simulation workloads.** We report the system size L (for FH only), the number of logical qubits n_q, the number of commuting groups, the total number of Pauli strings, the sum of Pauli weights across all strings W_{BK}, W_{JW}, and the estimated code distances d_{BK}, d_{JW} used in the wall clock time simulations for the different mappings. Here, W is the total Pauli weight, i.e., the sum of support sizes over all non-identity Pauli terms.

Arch.	N	Q/Module	Data	Hot	Bell	T	Total
Mono. 20L/5T	1	68,970	54,450	2,420	0	12,100	68,970
1D 6L5T	1+24	5,082/2,511	54,450	726	6,534	3,630	65,340
2D 1T	25	2,807	54,450	3,025	9,680	3,025	70,180
2D 5T	25	3,291	54,450	3,025	9,680	15,125	82,280

TABLE III: **Physical-qubit footprint for 450-logical-qubit workloads for each architecture.** We fix $d = 11$, differences in encoding choice can lead to slightly different distances.

C. Non-local operations and Bell pair generation

Our simulations are specific to the case where communication between modules (non-local operations) is realized via Bell pairs. But the results and analysis are broadly relevant to architectures where communication between subsystems is slower than within, and especially where some form of caching can be used.

In the heterogeneous architecture, each communication link maintains independent Bell pair pools for the two possible communication directions. In the homogeneous architecture, 2-4 pools are maintained depending on whether a module is on an edge or in the bulk.

For a code-distance d simulation, each external logical operation consumes $n = d^2$ Bell pairs from the corresponding pairs of pools. This is because a non-local CNOT via transversal operations requires one Bell pair per physical CNOT, and a transversal operation uses n physical CNOTs between the data qubits of the two codeblocks [41]. Importantly, provided the Bell-pair fidelity is within an order of magnitude of the local physical error rate, the logical error rate of a non-local operation is exponentially suppressed with respect to code distance [41], making error suppression possible. Detailed consideration of fidelities and decoherence are outside of the scope of this work.

We assume that Bell pairs are generated continuously at a rate $\tau_B = d$ Bell pairs per ms per directed pool. The parallel generation of $O(d)$ atom-photon pairs has been demonstrated experimentally in [29]. We take parallel generation of Bell pairs as an architectural assumption. If parallel generation is infeasible, the results of our paper can be adapted to serial ebit generation by multiplying the generation time by $O(d)$. We choose $O(d)$ Bell pair generation times of 1ms and 10ms. The 1ms time is roughly equal to syndrome extraction cycle time on a neutral atom device [53] and is experimentally feasible [43] for non-local implementations. The 10 ms time represents a pessimistic scenario in which the Bell state generation time is far slower than the QEC cycle time.

D. Pauli rotation and T-state generation

As described above, each Pauli rotation is assigned a fixed non-Clifford cost of 30 T states [26]. T states are supplied by dedicated T-state factories located within each module of the homogeneous architecture, and within the magic core of the heterogeneous architecture. A single T-factory produces one T state every 10 ms. Consequently, a Pauli rotation requiring 30 T states has a nominal T-stage time $\tau_T = \frac{300}{n_T}\mathrm{ms}$, where n_T is the number of available T-factory columns. In general, implementing these rotations requires a string of T and Clifford gates. Again, we ignore the cost of the Clifford gates as this is specific to the method of rotation synthesis used and is relatively negligible compared to the other costs in this paper. In the homogeneous architecture, each module has the same factory count. In the heterogeneous architecture, unless otherwise stated, we use $n_T = 5$ columns of T factories. This means each storage lane has five T factories dedicated to it. This corresponds to a nominal T-stage duration of 60 ms for (up to) the number of Pauli rotations equal to the number of lanes. A Pauli-string task reserves a compute row for its entire lifecycle, including parity accumulation, T-state consumption, and parity uncomputation.

E. Grouping

Given the Hamiltonian expressed as a sum of Pauli operators, the time evolution operator is constructed by partitioning its Pauli terms into commuting groups [2], [3], [45], with the groups executed sequentially.

The Pauli operators are partitioned into commuting groups using Qiskit's [22]

```
group_commuting(qubit_wise=False)
```

routine. Operators belonging to the same group commute and, therefore, can be executed in any order, while operators with disjoint supports can additionally be executed simultaneously. Once every operator in a group has been executed, the next group can start.

Within a group, the cycle-level scheduler greedily starts all Pauli strings whose qubit supports and lane sets are disjoint from currently active strings. Thus the grouping determines a legal ordering of mutually commuting rotations, while the simulator still enforces architectural conflicts at the level of qubits, lanes, and storage modules.

F. Simulation workloads

We consider systems of sizes ranging from 98 to 450 logical qubits for both the Fermi-Hubbard and sparse SYK models. Table II summarizes the properties of the workloads considered in this work.

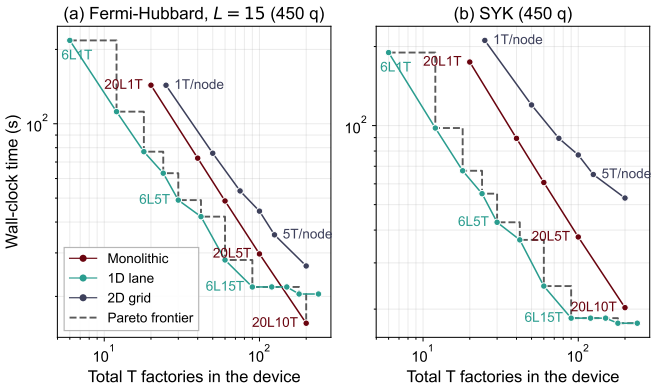

Fig. 3: **Wall-clock time against total T factory count Pareto plot. (a-b)** 450-qubit simulations at 1ms $O(d)$ Bell state generation rates. The data shown in these plots correspond to the Bravyi–Kitaev (BK) transformation.

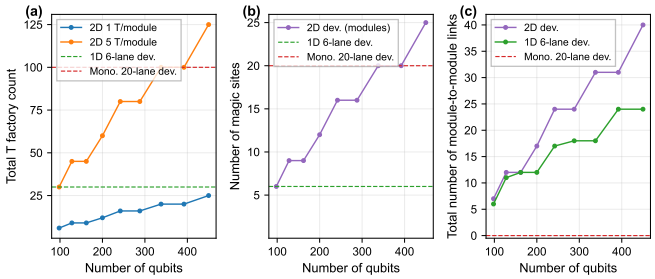

Fig. 4: **Number of T factories, magic sites and distributed connections. (a)** Total number of T factories across different distributed quantum computing architectures as problem size increases. **(b)** Total number of magic sites for each architecture. For the 1D and monolithic architectures, this remains fixed and is equal to the number of storage lanes. For the 2D architecture, it is equal to the number of modules and therefore increases with problem size. **(c)** Total number of distributed connections for different architectures against qubit count.

G. Scheduling model

The simulator advances all active objects in discrete cycles. At the beginning of each commuting group, all Pauli strings belonging to that group are placed in a scheduling queue. The scheduler greedily scans the queue and launches a Pauli-string task whenever its logical-qubit support and storage-lane set are disjoint from those of all currently active tasks. Strings that conflict with active tasks remain queued until the required resources become available.

Once a task has been launched, each occupied storage lane executes its parity-accumulation procedure independently. Consequently, the occupied lanes of a single Pauli string can operate in parallel, while multiple disjoint Pauli strings within the same commuting group are also executed concurrently.

A task completes only after all occupied lanes have finished uncomputation, at which point its logical qubits, storage lanes, communication resources, and assigned row of T factories are released.

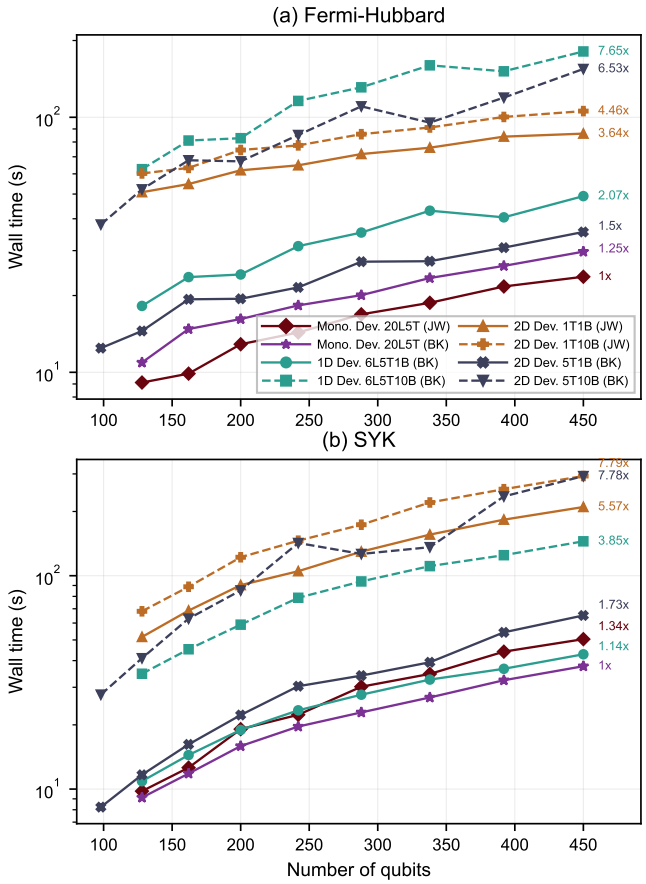

Fig. 5: **Performance of the architecture and its scaling against problem size. (a-b)** The performance of the device on an $L \times L$ grid Fermi-Hubbard simulation task with $N = 2L^2$ and a Sparse SYK simulation task with N complex fermionic modes. aL refers to a lanes, bT refers to b T-factory columns, and cB refers to $O(d)$ Bell pair generation time of c ms. End point ratios are against the fastest device.

V. System Performance

- **Observation 1: With substantially fewer magic sites and T-state factories, the heterogeneous (1D) architecture approaches monolithic and 2D device performance.**
- **Observation 2: For matched T factory count and under the same encoding, the 1D device is almost always faster than both the monolithic and 2D architectures.**

Figure 3 shows the performance–resource Pareto frontier for the architectures considered in this work as the total T factory count is varied. The six-lane heterogeneous (1D) architecture lies on the Pareto frontier across most resource budgets, only being overtaken by the 20 lane equivalent monolithic architecture once its 6 execution lanes become saturated. Importantly, the homogeneous distributed (2D) architecture never outperforms the 1D design for the same number of T factories. The performance gap is very significant for the SYK workload. While the 2D architecture for matched number of qubits, n_q, has more magic sites for parallel processing (see

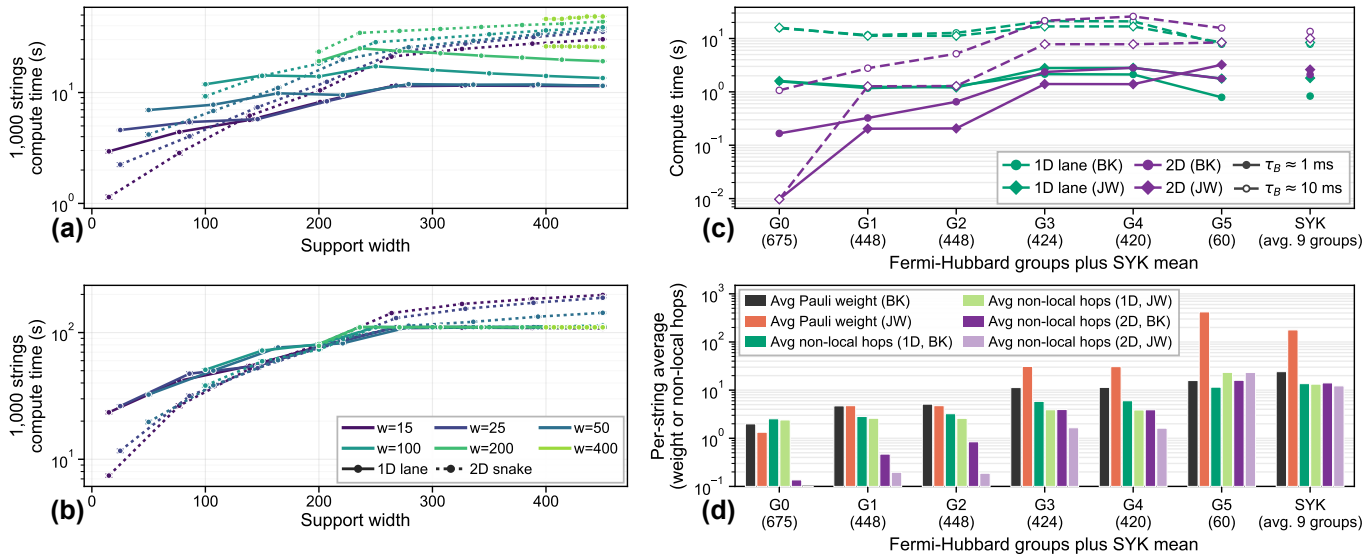

Fig. 6: **Compute time for random and work-load specific Pauli strings on different architectures. (a-b)** Random simulations. We fix $d = 11$ here. Each graph shows the total time taken to execute 1000 strings of weight w where for each Pauli string the supports are uniformly randomly distributed across a set of qubits $[c-W/2, c+W/2]$ of width W, where c is randomly chosen for each string. Error bars from re-running the experiment are smaller than the data points. **(a)** 1 ms $O(d)$ Bell pair generation speed. **(b)** $10\times$ slower $O(d)$ Bell pair generation time (10 ms). **(c)** Time to compute the different commuting groups of BK strings in a 450-qubit Fermi-Hubbard simulation for 1 ms and 10 ms $O(d)$ Bell pair generation times. We also report the results for a 450-qubit SYK simulation. This is averaged across the groups because the results were very similar for each group. **(d)** Average Pauli weight for BK and JW per Fermi-Hubbard commuting group (or SYK across all groups) and the average number of non-local hops in the different architectures and compilation methods.

Fig 4), the strings in this workload are difficult to execute in parallel due to their weight and disjointness (see Sec. VI and Sec. VIII for more details).

For the rest of this paper, we will focus on comparisons to the six lanes with five T-factories (6L5T) 1D architecture. This is because whilst its magic core has a larger physical qubit count (Tab. III) of ≈ 5100 compared to a 2D architecture module (≈ 3300) with five T-factories, the storage modules in the 1D device have on average ≈ 2500 physical qubits.

Figure 5(a) shows the performance of different architectures for Fermi-Hubbard simulations across a range of lattice sizes. Despite having only six parallel lanes that have five T factories each, the heterogeneous architecture is only about $2.1\times$ slower than the much larger monolithic baseline. This corresponds to executing 1 Trotter step in 49.1s. If execution scaled solely with the number of available lanes, one would expect a slowdown of approximately $3.3\times$. The smaller observed slowdown indicates that this simulation workload does not contain sufficient parallelism to fully utilize the larger monolithic device, even using the JW encoding.

With five T factories per node, the homogeneous architecture executes a Trotter step in $\approx$ 35.5s, $1.5\times$ slower than the monolithic device and $1.4\times$ faster than the 6L5T device. This is faster than the heterogeneous architecture. However, this performance comes from having approximately 125 T factories in total, which is more than the 100 T factory monolithic device, and $4\times$ more than the 6L5T heterogeneous architecture, which has only 30 T factories (Fig. 4(a)). When restricted to one T factory per node, corresponding to approximately 25 factories in total and therefore a resource budget comparable to that of the heterogeneous architecture (Fig. 4(a)), the homogeneous architecture is $\approx 2\times$ slower than the 1D device. Note that for the 2D 1T FH case, JW is used as it enables better use of the more limited magic compared to the 5T case, even if it induces a higher movement cost. In both cases, the homogeneous design also requires substantially greater communication connectivity and routing resources.

To evaluate sensitivity to communication performance, we reduce the generation speed of $O(d)$ Bell pairs from 1 ms to 10 ms. The execution time of the heterogeneous architecture nearly quadruples. However, the degradation in performance is substantially less than the $10\times$ drop that would be expected.

Figure 5(b) presents equivalent results for sparse SYK simulations. Here, as noted before, the 1D device performs consistently better, with the gap between it and the 2D device growing with qubit count. The gap widens for slower Bell state generation rate.

VI. Extensible storage performance

Observation 3: The 1D lane architecture becomes more competitive as Pauli-string support becomes more spatially non-local.

Observation 4: Slower Bell pair generation reduces sensitivity to Pauli weight because communication (rather than local parity accumulation) becomes the dominant cost.

Fig. 6 isolates the performance of the extensible storage system by excluding the Pauli rotation step and measuring only the time taken to compute Pauli strings. We find that only at the lowest weight and support width (high locality) does the 2D homogeneous architecture perform better than the 1D heterogeneous architecture. In this case, most Pauli strings are supported within a single module. We see that as the width of the distribution of supports goes up (increasing non-locality), the relative performance of the 1D system increases, becoming quickly better than the 2D device.

The 1D device's relative performance is best for the most distributed strings. This is because each of these strings are distributed across most or all of the 6 lanes of the device. This means they will be computed using all 6 lanes. Whereas for medium-distributed high-weight strings, it is difficult to continuously form commuting sets with disjoint supports that use all 6 lanes. This is why in Fig. 6(a) the runtime increases from low to medium support widths and then decreases again.

In Fig. 6(b) we see that decreasing the generation rate of Bell pairs by $10\times$ results in clustering behavior between Pauli strings of significantly different weights.

This is because the cost of internal movement becomes much lower so the difference between high weight and low weight strings that are both widely distributed (and therefore have similar external communication requirements) becomes minimal.

We also evaluate compute times, average Pauli weights and number of non-local hops under different encodings for the different commuting groups in Fermi-Hubbard simulation (and averaged across the entire simulation for SYK). First, we see how the behaviors identified in the random benchmark above continue in the real simulation: The 1D device is competitive for the more distributed commuting groups. That being said, the disjointness (and therefore parallelism) granted by the Jordan-Wigner transformation enables the 25 2D nodes for $n = 450$ qubits to compute with a high degree of parallelism, especially in the groups with lowest Pauli string weight. Importantly, we see how compute time is concentrated in the groups with the highest weight and most number of non local hops required. Therefore, even though small groups can be executed quickly and locally on the 2D device, the time to compute these is negligible compared to the larger groups where the 1D device is competitive. This shows how the 1D device is able to efficiently compute the strings, which contribute most to runtime, even given its lower physical costs. We analyze this further in the next section.

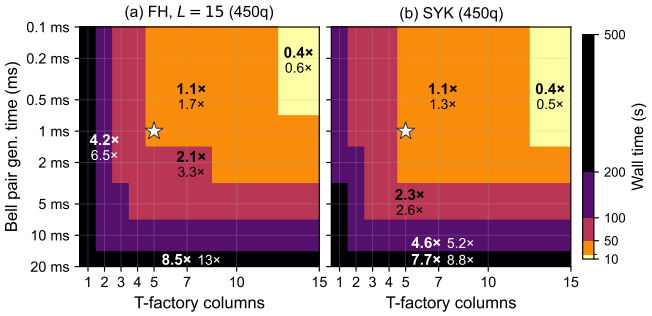

Fig. 7: **Communication time and T factory production** Heatmap shows wall clock time as a function of Bell generation rate and T-factory columns. Star indicates the 1D 6L5T device. Bold represents runtime-ratios against a monolithic version of this device. Non-bold represents runtime-ratios against the 20L5T monolithic device.

VII. Network and Magic Relationship

Observation 5: Additional T factories provide little benefit in a slow storage access regime.

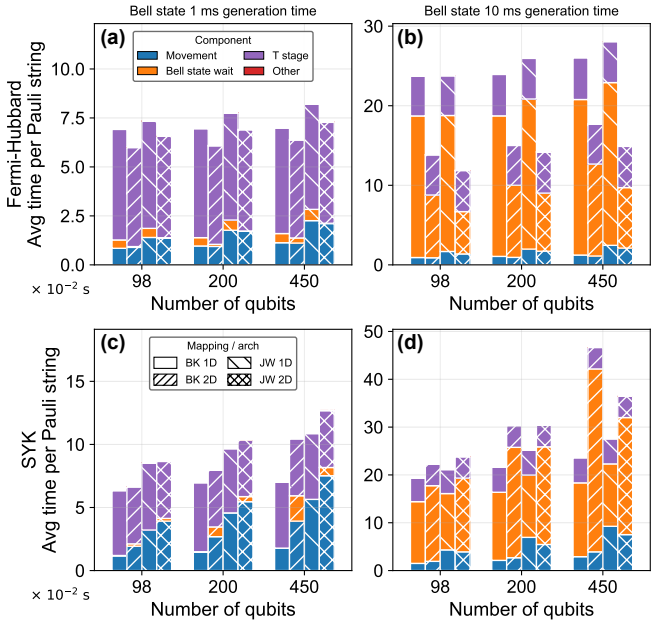

Fig. 8: **Breakdown of execution time per Pauli string. (a-d)** Average time per Pauli string with identified blockers for different workloads, Bell state speeds, number of qubits and encoding.

The heterogeneous architecture separates computation and storage into distinct subsystems. Consequently, overall performance is determined by the balance between the rate at which the magic core can supply non-Clifford resources and the rate at which the storage lanes can deliver parity information. If T states are produced rapidly, the storage system must be able to compute and communicate parities at a comparable rate in order to keep the compute node occupied. Conversely, if T-state production is slow, improvements in communication bandwidth provide little benefit because execution becomes magic-limited.

This behavior is illustrated in Fig.7(a-b), which shows execution time as a function of both Bell pair-generation rate and T-factory count. When Bell pair generation is fast enough, increasing the number of T-factories leads to approximately proportional improvements in execution time, indicating that the architecture is primarily compute-limited. However, as the Bell pair generation becomes slower, communication increasingly dominates execution. For example, at a Bell pair-

generation time of 10 ms, little performance benefit is obtained by increasing the number of T factories beyond three, as the storage system is no longer able to supply parity information quickly enough to utilize the additional non-Clifford resources.

To better understand these trends, Fig. 8(a-d) illustrates the average execution time of a Pauli string into its constituent components. At fast Bell-generation rates (1 ms), execution time is dominated by T-state generation, while communication overhead contributes only a small fraction of the total runtime. In contrast, at slower Bell state generation rates (10 ms), communication delays become the majority of the execution time. This effect is comparatively less pronounced for the homogeneous distributed architecture in the Fermi-Hubbard simulation, due to its ability to perform rotations locally for less distributed strings.

VIII. Parallelism

Observation 6: Quantum workload structure determines usable parallelism.

Observation 7: For overlapping non-local workloads, intra-string parallelism is more valuable than additional parallelism across Pauli-strings.

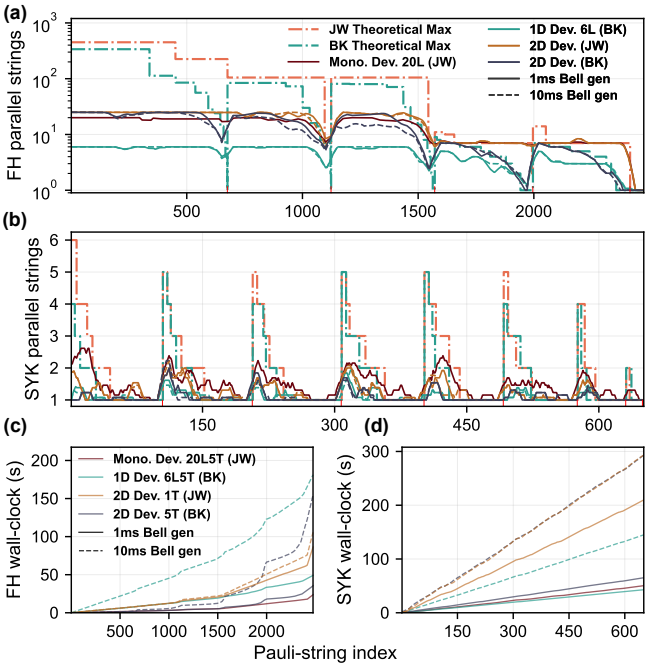

Fig. 9: **Parallelizable Pauli strings and wall clock time. (a, b)** Graphs of the theoretical max for different encodings and actually achieved for combinations of encodings, architectures, and workloads (**a**: Fermi-Hubbard **b**: SYK), number of simultaneously executable Pauli strings throughout execution of a single Trotter step. **(c, d)** Wall clock time against number of executed Pauli strings for different architectures and workloads. Dashed lines indicate 10ms Bell pair generation speed rather than 1ms.

Fig. 9(a-b) depicts the degree of available and achieved parallelism for the different workloads, architectures and encodings we study. Focusing on (a), we know that at the start of the simulation the average Pauli string weight of commuting groups is small (see Fig.6(c)). This means the theoretical maximum attainable parallelism, limited by commutativity and disjointness of the strings, is high. As the average Pauli weight increases, this decreases. We see that JW consistently has a higher degree of parallelism than BK, as expected. However, this metric becomes matched with BK later in the simulation. At earlier points, a hugely parallel device would be needed to take advantage of the theoretical limit.

Fig. 9(c) shows that at the start of the simulation the homogeneous architecture and monolithic baseline both initially (up to about the 1750^{th} Pauli string) have an increase in wall clock time with a lower gradient compared to the 1D device. The larger number of magic sites in these architectures therefore translates directly to an increased performance in this regime compared to the heterogeneous architecture.

However, beyond this their advantage largely disappears. The remaining operators are dominated by the larger Pauli strings, whose supports can overlap extensively. As shown by the decreasing theoretical parallelism curves in Fig. 9(a), only a small number of these strings can be executed simultaneously regardless of the amount of available hardware. The homogeneous architecture therefore loses most of its string-level parallelism advantage.

In this regime, the heterogeneous architecture benefits from a different form of parallelism. Rather than executing more Pauli strings simultaneously, it accelerates the execution of individual high-weight strings. The parity-computation circuit is distributed across multiple storage lanes, allowing the compute and uncompute stages to proceed concurrently. Consequently, even when only a single Pauli string is being executed, the parity accumulation process can be parallelized across up to six lanes. The homogeneous architecture lacks an equivalent mechanism and therefore performs the parity computation largely serially along its routing path. As a result, the largest Pauli strings, which dominate the tail of the execution profile, are completed significantly faster on the heterogeneous architecture.

This effect is accentuated in the SYK simulation because its Pauli strings are both non-local and strongly overlapping. The theoretical and achieved string-level parallelism remain close to one for much of the workload, leaving the additional execution sites of the homogeneous architecture underutilized. Performance is therefore determined primarily by the time to compute each individual parity, favoring the lane architecture's intra-string parallelism.

IX. Discussion and Future Work

The main concept introduced and evaluated in this work is a quantum storage system with specific access patterns and simple design. We have shown that this system can access arbitrary highly non-local Pauli strings with varying weights and non-locality faster than a homogeneous architecture (2D) that has higher connectivity. The storage lane concept enables inter- and intra-string parallelism, giving it flexibility across

different workloads. The storage system is also extensible: It can adapt to workload requirements and be extended at relatively low cost, especially compared to extending a device by adding fully capable modules. As a result, the proposed architecture is competitive with homogeneous architectures, despite using significantly fewer magic factories and sites.

The random parity access capability of the storage system enables saving in quantum simulation by making otherwise difficult, highly non-local fermionic encodings practical. This capability may benefit other quantum algorithms as well. There is scope to explore new or dynamic fermionic encodings that exploit the low cost access to non-local Pauli strings provided with this architecture. The storage system can also be modified, at some additional cost, to support tree-like or other hierarchical structures.

Further research should consider the use of qLDPC codes in place of the surface codes used in this work. Their higher encoding rate could increase storage density. Direct Pauli-based computation approaches, such as those proposed in [25], [50], could also be promising within a distributed storage system such as the one introduced here. Improvements to T-state production and application, including analog rotations on transversal [20] and extractor architectures [38], are expected to further increase the relative importance of storage access and routing overheads. Overall, we hope our results encourage further research into fault tolerant, extensible quantum storage systems.

X. Appendix

Here we provide details about the fermionic systems and mapping discussed in this work.

A. Fermionic models

We consider two representative fermionic models: the Fermi–Hubbard model [19] and the sparse Sachdev–Ye–Kitaev [49] (SYK) model.

The Fermi–Hubbard Hamiltonian on an $L \times L$ lattice with periodic boundary conditions is

$$H = -t \sum_{\sigma \in \{\uparrow,\downarrow\}} \sum_{\langle i,j \rangle} \left(c_{i\sigma}^\dagger c_{j\sigma} + c_{j\sigma}^\dagger c_{i\sigma} \right) + U \sum_i n_{i\uparrow} n_{i\downarrow}, \quad (4)$$

where $c_{i\sigma}^\dagger$ and $c_{i\sigma}$ are, respectively, the fermionic creation and annihilation operators for spin σ at lattice site $i < L^2$ with $N = 2L^2$, and $n_{i\sigma} = c_{i\sigma}^\dagger c_{i\sigma}$ is the corresponding number operator. The parameter t denotes the nearest-neighbor hopping strength, U is the on-site interaction strength, and $\langle i, j \rangle$ denotes pairs of neighboring lattice sites.

We use Qiskit [22] Nature's `FermiHubbardModel` to generate the Hamiltonian, with nearest-neighbor hopping amplitudes $t_{ij} = 1$ and on-site interaction strength $U = 8$. Each lattice site contains two fermionic modes, corresponding to spin up and spin down, and is therefore represented by two qubits after fermion-to-qubit mapping. An $L \times L$ lattice consequently requires $2L^2$ qubits.

We consider a sparse SYK Hamiltonian as defined in [32]. It is defined on N complex fermionic modes, or equivalently $2N$ Majorana modes. To match the Fermi–Hubbard qubit counts, the main sparse-SYK calculations use $N = 2L^2$ for the same nominal lattice parameter L. Quartets of Majorana modes are included independently with probability $p = d/\binom{2N-1}{3}$, where $d = 3$ is the mean-degree parameter. For each included quartet $I = (i, j, k, l)$, the coupling is sampled from a unit normal distribution and normalized by $1/\sqrt{2dN}$, giving

$$H_{\mathrm{SYK}} = \frac{1}{\sqrt{2dN}} \sum_I X_I J_I C_I, \quad (5)$$

where X_I is a Bernoulli inclusion variable, J_I is a Gaussian coupling coefficient, C_I is the appropriately phased product of the four Majorana operators associated with quartet I.

B. Fermion-to-qubit transformations

We use the Jordan–Wigner (JW) and Bravyi–Kitaev (BK) transformations [8], [23] to map the fermionic creation and annihilation operators to qubit operators. In both cases, a fermionic Hamiltonian is transformed into a weighted sum of Pauli strings.

Under the Jordan–Wigner transformation, the creation and annihilation operators for fermionic mode i are mapped as

$$c_i^\dagger \mapsto \frac{(X_i - iY_i)}{2} \Pi_{j=0}^{i-1} Z_j, \text{and } c_i \mapsto \frac{(X_i + iY_i)}{2} \Pi_{j=0}^{i-1} Z_j. \quad (6)$$

The string of Z operators records the parity of all fermionic modes preceding mode i, ensuring that the mapped operators satisfy the fermionic anticommutation relations. Although this mapping produces a highly regular support structure, the length of the parity string grows linearly with the mode index. The maximum Pauli weight therefore scales as $O(N)$ for a system containing N fermionic modes.

In the Bravyi–Kitaev transformation, occupation and parity information are distributed across qubits using a binary-tree encoding. Following the convention implemented by Qiskit Nature, each fermionic mode i is associated with an update set $U(i)$, a parity set $P(i)$, and a flip set $F(i)$. We additionally define the remainder set $R(i) = P(i) \setminus F(i)$.

The two Pauli components used to represent the creation and annihilation operators are

$$A_i = X_i \left(\Pi_{k \in U(i)} X_k \right) \left(\Pi_{k \in P(i)} Z_k \right), \quad (7)$$

and

$$B_i = Y_i \left(\Pi_{k \in U(i)} X_k \right) \left(\Pi_{k \in R(i)} Z_k \right). \quad (8)$$

The fermionic operators are then mapped according to

$$c_i^\dagger \mapsto \frac{1}{2} \left(A_i - iB_i \right), \qquad c_i \mapsto \frac{1}{2} \left(A_i + iB_i \right). \quad (9)$$

The update, parity, and flip sets are obtained from the binary-tree structure and contain at most $O(\log N)$ qubits. Consequently, the resulting Pauli operators have weight $O(\log N)$ rather than the $O(N)$ weight produced by the Jordan–Wigner transformation. This lower weight is obtained at the cost of a less geometrically local and less regular support structure, because the qubits associated with a single fermionic mode may be distributed across several parts of the encoding tree.

C. Code distance and physical-resource estimates

Code distances are estimated from the total number of logical operations required by the Pauli-string workload. Let W be the sum of all Pauli-string weights, M the number of non-identity Pauli strings, and N_{step} the assumed number of Trotter steps. The total logical-operation count is approximated as $N_{\mathrm{op}} = (2W + 30M)\, N_{\mathrm{step}}$ where the factor $2W$ accounts for parity computation and uncomputation, and the term $30M$ accounts for the fixed T-state cost assigned to each Pauli rotation. Unless otherwise stated, we assume $N_{\mathrm{step}} = 10$.

To ensure that the total probability of success with respect to physical error over the full simulation remains above a reasonable 90%, we assign a target logical error rate per operation of $p_L = 1 - 0.9^{1/N_{\mathrm{op}}}$.

The required code distance is then estimated using the phenomenological fit $d = \left\lceil d_0 + \frac{2\log(p_L/p_0)}{\log(r)} \right\rceil$ with $d_0 = 11$, $p_0 = 5 \times 10^{-7}$, the physical two-qubit error rate 10^{-3}, threshold 7×10^{-3}, and $r = 10^{-3}/(7 \times 10^{-3})$.

Acknowledgments

We thank Yifan (Frank) Zhang for helpful comments on the manuscript. J.S. and S.L. were summer interns at Lawrence Berkeley National Laboratory during the initiation of this project. J.S. and F.M. acknowledge support from NSF grants OSI-2531350, OSI-2410675, PHY-2325080 and OMA-2120757 as well as DOE grant DE-SC0025384. S.L. acknowledges support from the National Science Foundation STAQ project under Grant No. PHY-2325080. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC award ASCR-ERCAP0037552. Y.L. acknowledges support from DOE grant DE-SC0025384. This work was also supported in part by the U.S. Department of Energy, Office of Science, under Award No. DE-SCL0000039 to Lawrence Berkeley National Laboratory (PI: E.S.).

AI Use

Claude Opus 4.8 and ChatGPT 5.5 were used to extend and run the simulation code in this paper. They were also used to generate code to produce the figures in this paper from the simulator data. The diagrams (Fig.1 and Fig.2) were made by a human.

References

[1] G. Aigner, B. Klaver, M. Lanthaler, and W. Lechner, "Fermion lattices can be simulated by same-size qubit lattices with $\mathcal{O}(1)$ interaction overhead," 2026. [Online]. Available: https://arxiv.org/abs/2605.12600

[2] A. Anand and K. R. Brown, "Hamiltonian-based graph-state ansatz for variational quantum algorithms," *Physical Review A*, vol. 111, no. 1, p. 012437, 2025.

[3] A. Anand and K. R. Brown, "Leveraging commuting groups for an efficient variational hamiltonian ansatz," *Quantum Science and Technology*, vol. 10, no. 4, p. 045009, 2025.

[4] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, "Simulated quantum computation of molecular energies," *Science*, vol. 309, no. 5741, pp. 1704–1707, 2005.

[5] M. E. Beverland, P. Murali, M. Troyer, K. M. Svore, T. Hoefler, V. Kliuchnikov, G. H. Low, M. Soeken, A. Sundaram, and A. Vaschillo, "Assessing requirements to scale to practical quantum advantage," 2022. [Online]. Available: https://arxiv.org/abs/2211.07629

[6] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. Sales Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin, "Logical quantum processor based on reconfigurable atom arrays," *Nature*, vol. 626, no. 7997, pp. 58–65, Dec. 2024. [Online]. Available: http://dx.doi.org/10.1038/s41586-023-06927-3

[7] M. Bothe, C. Sünderhauf, M. J. Witham, E. T. Campbell, and N. S. Blunt, "More efficient clifford+t synthesis for small-angle rotations and application to trotterization," 2026. [Online]. Available: https://arxiv.org/abs/2605.31544

[8] S. B. Bravyi and A. Y. Kitaev, "Fermionic quantum computation," *Annals of Physics*, vol. 298, no. 1, pp. 210–226, 2002. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0003491602962548

[9] J. I. Cirac, A. K. Ekert, S. F. Huelga, and C. Macchiavello, "Distributed quantum computation over noisy channels," *Physical Review A*, vol. 59, no. 6, pp. 4249–4254, Jun 1999. [Online]. Available: http://dx.doi.org/10.1103/PhysRevA.59.4249

[10] A. M. Dalzell, S. McArdle, M. Berta, P. Bienias, C.-F. Chen, A. Gilyén, C. T. Hann, M. J. Kastoryano, E. T. Khabiboulline, A. Kubica, G. Salton, S. Wang, and F. G. S. L. Brandão, *Quantum Algorithms: A Survey of Applications and End-to-end Complexities*. Cambridge University Press, Apr. 2025. [Online]. Available: http://dx.doi.org/10.1017/9781009639651

[11] B. Eastin and E. Knill, "Restrictions on transversal encoded quantum gate sets," *Phys. Rev. Lett.*, vol. 102, p. 110502, Mar 2009. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.102.110502

[12] B. Fauseweh, "Quantum many-body simulations on digital quantum computers: State-of-the-art and future challenges," *Nature Communications*, vol. 15, no. 1, p. 2123, 2024. [Online]. Available: https://doi.org/10.1038/s41467-024-46402-9

[13] P. M. Fenwick, "A new data structure for cumulative frequency tables," *Software: Practice and Experience*, vol. 24, no. 3, pp. 327–336, 1994. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/spe.4380240306

[14] D. Filippov, P. Yang, and P. Murali, "Architecting distributed quantum computers: Design insights from resource estimation," 2026. [Online]. Available: https://arxiv.org/abs/2508.19160

[15] I. Georgescu, S. Ashhab, and F. Nori, "Quantum simulation," *Reviews of Modern Physics*, vol. 86, no. 1, p. 153–185, Mar. 2014. [Online]. Available: http://dx.doi.org/10.1103/RevModPhys.86.153

[16] C. Gidney, "How to factor 2048 bit rsa integers with less than a million noisy qubits," 2025. [Online]. Available: https://arxiv.org/abs/2505.15917

[17] D. Gottesman, "Opportunities and challenges in fault-tolerant quantum computation," *arXiv preprint arXiv:2210.15844*, 2022.

[18] L. K. Grover, "Quantum telecomputation," 1997. [Online]. Available: https://arxiv.org/abs/quant-ph/9704012

[19] J. Hubbard, "Electron correlations in narrow energy bands," *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, vol. 276, no. 1365, pp. 238–257, Nov 1963. [Online]. Available: https://doi.org/10.1098/rspa.1963.0204

[20] R. Ismail, I.-C. Chen, C. Zhao, R. Weiss, F. Liu, H. Zhou, S.-T. Wang, A. Sornborger, and M. Kornjača, "Transversal architecture for megaquop-scale quantum simulation with neutral atoms," *PRX Quantum*, vol. 7, p. 020343, May 2026. [Online]. Available: https://link.aps.org/doi/10.1103/j2fw-ccmy

[21] H. Jacinto, E. Gouzien, and N. Sangouard, "Network requirements for distributed quantum computation," *Phys. Rev. Res.*, vol. 8, p. 013205, Feb 2026. [Online]. Available: https://link.aps.org/doi/10.1103/v9ln-c4v2

[22] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, B. R. Johnson, and J. M. Gambetta, "Quantum computing with Qiskit," 2024.

[23] P. Jordan and E. Wigner, "Über das paulische äquivalenzverbot," *Zeitschrift für Physik*, vol. 47, no. 9, pp. 631–651, 1928.

[24] E. Kaur, S. Pouryousef, N. K. Chandra, H. Shapourian, J. Zhao, R. Kompella, and R. Nejabati, "Impact of network constraints on fault-tolerant distributed quantum computing," 2026. [Online]. Available: https://arxiv.org/abs/2606.17495

[25] S. Khan, S. Sethi, K. Sahay, Y. Lin, J. Alnas, S. Kurapati, A. Anand, J. M. Baker, and K. R. Brown, "Architecting early fault tolerant neutral atoms systems with quantum advantage," *arXiv preprint arXiv:2604.19735*, 2026.

[26] V. Kliuchnikov, K. Lauter, R. Minko, A. Paetznick, and C. Petit, "Shorter quantum circuits via single-qubit gate approximation," *Quantum*, vol. 7, p. 1208, Dec. 2023. [Online]. Available: http://dx.doi.org/10.22331/q-2023-12-18-1208

[27] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, "Even more efficient quantum computations of chemistry through tensor hypercontraction," *PRX Quantum*, vol. 2, p. 030305, Jul 2021. [Online]. Available: https://link.aps.org/doi/10.1103/PRXQuantum.2.030305

[28] D. Li, S. Xu, and Y. Ding, "Asymptotically optimal depth fermionic permutation on 2d grid quantum architecture without ancillas," 2026. [Online]. Available: https://arxiv.org/abs/2605.26041

[29] L. Li, X. Hu, Z. Jia, W. Huie, W. K. C. Sun, Aakash, Y. Dong, N. Hiri-O-Tuppa, and J. P. Covey, "Parallelized telecom quantum networking with an ytterbium-171 atom array," *Nature Physics*, vol. 21, no. 11, pp. 1826–1833, Sep 2025. [Online]. Available: http://dx.doi.org/10.1038/s41567-025-03022-4

[30] D. Litinski, "A game of surface codes: Large-scale quantum computing with lattice surgery," *Quantum*, vol. 3, p. 128, 2019.

[31] S. Lloyd, "Universal quantum simulators," *Science*, vol. 273, no. 5278, pp. 1073–1078, 1996. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.273.5278.1073

[32] N. Maskara, M. Kalinowski, D. Gonzalez-Cuadra, and M. D. Lukin, "Fast simulation of fermions with reconfigurable qubits," 2025. [Online]. Available: https://arxiv.org/abs/2509.08898

[33] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, "Quantum computational chemistry," *Rev. Mod. Phys.*, vol. 92, p. 015003, Mar 2020. [Online]. Available: https://link.aps.org/doi/10.1103/RevModPhys.92.015003

[34] C. Monroe and J. Kim, "Scaling the ion trap quantum processor," *Science*, vol. 339, no. 6124, pp. 1164–1169, 2013. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.1231298

[35] P. S. Mundada, A. Khindanov, Y. Wang, C. L. Edmunds, P. Coote, M. J. Biercuk, Y. Baum, and M. Hush, "Heterogeneous architectures enable a 138x reduction in physical qubit requirements for fault-tolerant quantum computing under detailed accounting," 2026. [Online]. Available: https://arxiv.org/abs/2604.06319

[36] J. Necaise, N. Anand, G. Gyawali, K. G. Johnson, J. D. Whitfield, and M. Mohseni, "Distribution complexity of electronic structure simulations on quantum supercomputers," 2026. [Online]. Available: https://arxiv.org/abs/2606.20805

[37] C. Ryan-Anderson, N. C. Brown, C. H. Baldwin, J. M. Dreiling, C. Foltz, J. P. Gaebler, T. M. Gatterman, N. Hewitt, C. Holliman, C. V. Horst, J. Johansen, D. Lucchetti, T. Mengle, M. Matheny, Y. Matsuoka, K. Mayer, M. Mills, S. A. Moses, B. Neyenhuis, J. Pino, P. Siegfried, R. P. Stutz, J. Walker, and D. Hayes, "High-fidelity teleportation of a logical qubit using transversal gates and lattice surgery," *Science*, vol. 385, no. 6715, pp. 1327–1331, 2024. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.adp6016

[38] S. Sethi, S. Khan, A. Awasthi, A. Anand, and J. M. Baker, "Injeqt: Improved magic-state injection protocol for fault-tolerant quantum extractor architectures," *arXiv preprint arXiv:2604.25094*, 2026.

[39] P. W. Shor, "Fault-tolerant quantum computation," in *Proceedings of 37th conference on foundations of computer science*. IEEE, 1996, pp. 56–65.

[40] J. Stack, Y. Liu, A. Ruocco, G. Zervas, and A. Beghelli, "Impact of a free space optics switch on ebit fidelity for distributed quantum computing," in *Quantum Computing, Communication, and Simulation VI*, P. R. Hemmer, A. L. Migdall, and I. A. Burenkov, Eds., vol. 13919, International Society for Optics and Photonics. SPIE, 2026, p. 139190Y. [Online]. Available: https://doi.org/10.1117/12.3073698

[41] J. Stack, M. Wang, and F. Mueller, "Transversal fault tolerant distributed quantum computing operations," *Nature Communications*, Jul. 2026. [Online]. Available: http://dx.doi.org/10.1038/s41467-026-75693-3

[42] S. Stein, S. Xu, A. W. Cross, T. J. Yoder, A. Javadi-Abhari, C. Liu, K. Liu, Z. Zhou, C. Guinn, Y. Ding, Y. Ding, and A. Li, "Hetec: Architectures for heterogeneous quantum error correction codes," in *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, ser. ASPLOS '25. New York, NY, USA: Association for Computing Machinery, 2025, pp. 515–528. [Online]. Available: https://doi.org/10.1145/3676641.3716001

[43] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, "High-rate, high-fidelity entanglement of qubits across an elementary quantum network," *Phys. Rev. Lett.*, vol. 124, p. 110501, Mar 2020. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.124.110501

[44] M. Suzuki, "Generalized trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems," *Communications in Mathematical Physics*, vol. 51, no. 2, pp. 183–190, 1976.

[45] E. Van Den Berg and K. Temme, "Circuit optimization of hamiltonian simulation by simultaneous diagonalization of pauli clusters," *Quantum*, vol. 4, p. 322, 2020.

[46] R. Van Meter, K. Nemoto, and W. Munro, "Communication links for distributed quantum computation," *IEEE Transactions on Computers*, vol. 56, no. 12, pp. 1643–1653, Dec 2007. [Online]. Available: http://dx.doi.org/10.1109/TC.2007.70775

[47] P. Webster, L. Berent, O. Chandra, E. T. Hockings, N. Baspin, F. Thomsen, S. C. Smith, and L. Z. Cohen, "The pinnacle architecture: Reducing the cost of breaking rsa-2048 to 100 000 physical qubits using quantum ldpc codes," 2026. [Online]. Available: https://arxiv.org/abs/2602.11457

[48] S. Williams, A. Waterman, and D. Patterson, "Roofline: an insightful visual performance model for multicore architectures," *Communications of the ACM*, vol. 52, no. 4, pp. 65–76, 2009.

[49] S. Xu, L. Susskind, Y. Su, and B. Swingle, "A sparse model of quantum holography," 2020. [Online]. Available: https://arxiv.org/abs/2008.02303

[50] T. J. Yoder, E. Schoute, P. Rall, E. Pritchett, J. M. Gambetta, A. W. Cross, M. Carroll, and M. E. Beverland, "Tour de gross: A modular quantum computer based on bivariate bicycle codes," Jun 2025, arXiv:2506.03094 [quant-ph]. [Online]. Available: http://arxiv.org/abs/2506.03094

[51] H. Zhang, Y. Xu, H. Hu, K. Yin, H. Shapourian, J. Zhao, R. R. Kompella, R. Nejabati, and Y. Ding, "Switchqnet: Optimizing distributed quantum computing for quantum data centers with switch networks," in *Proceedings of the 52nd Annual International Symposium on Computer Architecture*, ser. ISCA '25. New York, NY, USA: Association for Computing Machinery, 2025, pp. 1449–1463. [Online]. Available: https://doi.org/10.1145/3695053.3731046

[52] J. Zhao, S. Vinet, A. Minoofar, M. Kilzer, L. Wang, G. Moody, V. Pandey, R. Kompella, and R. Nejabati, "A universal quantum information preserving photonic switch for scalable quantum networks," 2026. [Online]. Available: https://arxiv.org/abs/2604.21902

[53] H. Zhou, C. Duckering, C. Zhao, D. Bluvstein, M. Cain, A. Kubica, S.-T. Wang, and M. D. Lukin, "Resource analysis of low-overhead transversal architectures for reconfigurable atom arrays," in *Proceedings of the 52nd Annual International Symposium on Computer Architecture*, ser. SIGARCH '25. ACM, Jun. 2025, pp. 1432–1448. [Online]. Available: http://dx.doi.org/10.1145/3695053.3731039