

## Fault-tolerant optical interconnects for neutral-atom arrays

Josiah Sinclair<sup>1</sup>, Joshua Ramette,<sup>1</sup> Brandon Grinkemeyer,<sup>2</sup> Dolev Bluvstein<sup>1</sup>,<sup>2</sup> Mikhail D. Lukin,<sup>2</sup> and Vladan Vuletić<sup>1</sup>

<sup>1</sup>*Department of Physics, MIT-Harvard Center for Ultracold Atoms and Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

<sup>2</sup>*Department of Physics, Harvard University, Cambridge, Massachusetts 02139, USA*



(Received 28 August 2024; accepted 28 January 2025; published 25 March 2025)

We analyze the use of photonic links to enable large-scale fault-tolerant connectivity of locally error-corrected modules based on neutral atom arrays. Our approach makes use of recent theoretical results showing the robustness of surface codes to boundary noise and combines recent experimental advances in atom-array quantum computing with logical qubits with optical quantum networking techniques. We find the conditions for fault tolerance can be achieved with local two-qubit Rydberg gate and nonlocal Bell-pair errors below 1% and 10%, respectively, without requiring distillation or space-time overheads. Realizing the interconnects with a lens, a single optical cavity, or an array of cavities enables—with sufficient multiplexing—a Bell-pair generation rate in the 1–50 MHz range. When directly interfacing logical qubits, this rate translates to error-correction cycles in the 25–2000 kHz range, satisfying all requirements for fault tolerance and in the upper range fast enough for 100 kHz logical clock cycles.

DOI: 10.1103/PhysRevResearch.7.013313

Quantum computers made from noisy components require quantum error correction (QEC) to scale [1]. However, building error-corrected devices is a challenging task, as it requires connecting and controlling a large number of physical qubits with high fidelity: physical components and operations must introduce noise sufficiently below a code-dependent threshold [2]. Recently, programmable atom arrays with hundreds of qubits, long coherence times, and high-fidelity Rydberg gates [3–5] have emerged as a leading platform with experiments demonstrating continuous operation [6,7], control of up to 48 logical qubits, and reduction of the error rate with increasing code distance [8]. In these experiments, optical multiplexing enables efficient control over logical rather than physical qubits, greatly easing experimental complexity and paving the way for intermediate-scale error-corrected devices containing hundreds of logical qubits.

Further scaling will greatly benefit from modular architectures, where modules containing many physical qubits and supporting high-fidelity and error-corrected “local” quantum operations are connected via long-range entanglement, mediated by photons [9–11]. Modular architectures have the appeal of transforming the challenge of scaling into the task of realizing a unit module with a sufficiently high rate and fidelity quantum input-output interface for fault-tolerant communication. Once this is accomplished, the system can then be scaled up arbitrarily by adding more modules, eliminating the need to solve new scientific and engineering challenges at each new larger size.

In this paper, we propose and analyze such a unit module for fault-tolerantly linking Rydberg arrays via photonic interconnects. We first consider the noise requirements for fault-tolerant quantum communication. While multiple recent works have targeted network noise below 1% to achieve modular fault tolerance [12–15], we show the network noise can under certain conditions exceed 10%. We apply to neutral atoms the recent results of Ref. [16], which showed that distinct surface-code patches can be connected in a fault-tolerant manner even in the presence of substantial noise along their connecting interface. Translated to a Rydberg qubit error model, we find this corresponds to a local Rydberg gate and nonlocal Bell-pair (photonic interconnect) threshold of 1% and 10%, respectively, representing a tenfold increase in the Bell-pair threshold and comparing favorably to the current state of the art [3,17–28]. This high tolerance for communication errors brings the fault-tolerant connection of error-corrected modules within reach of existing atom array technology, with the primary remaining challenge being the development of sufficiently fast and efficient photonic links.

Turning to this challenge, we next explore several module designs depicted in Fig. 1 that aim to achieve fast quantum communication. Our module designs utilize a zoned architecture [8], leveraging the recently demonstrated capability of fast coherent transport [22,30,31]. Bell pairs are generated in a “remote entanglement generation zone” and then transported to a quiescent “computation zone” where logical qubits are maintained and operated. In the first design, a large numerical aperture (NA) lens and detector array are used to entangle atoms remotely. In this approach, the rate of entanglement generation,  $\sim 200$  Hz/atom, is limited by the lens’s collection efficiency, but very high rates can be reached with sufficient multiplexing. Contrasting with this slow but highly parallel approach, we consider a design employing a single optical cavity using the realistic parameters from Young *et al.* [32].

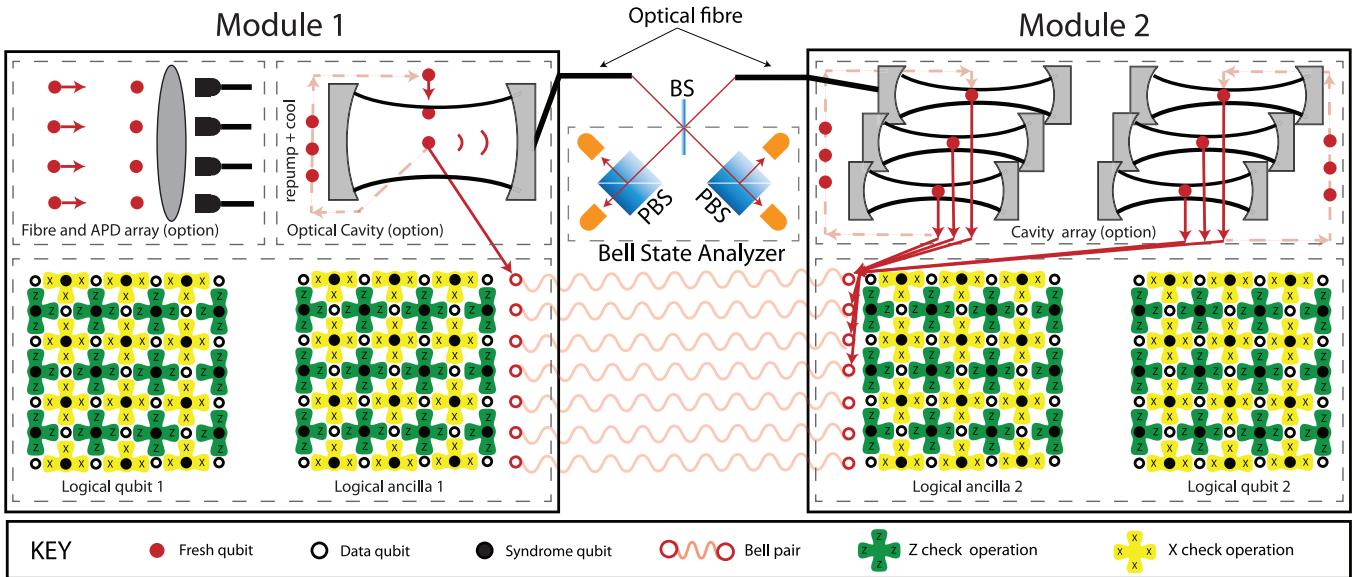


FIG. 1. Two modules equipped with sufficient quantum input/output for fault-tolerant communication. Surface-code patches in each module are realized using an array of atoms and connected using teleported gates. The Bell pairs are generated using either a lens, an optical cavity, or an array of optical cavities. Once a Bell pair is created, it is transported (red arrows) from the remote entanglement zone to the computation zone, placed on a logical qubit boundary, and used to enact a parity check between modules. Interfacing logical qubits in this way allows logical CNOT gates between modules, for example via lattice surgery CNOT gate [29].

Parallel control of multiple itinerant qubits combined with temporal multiplexing can enable Bell-pair generation rates reaching  $\sim 1$  MHz and error correction cycles of 25 kHz. Finally, we seek to combine the parallelism of the free-space approach with the speed of an optical cavity by introducing a third module design that utilizes an array of optical micro-cavities. For an array of 30 cavities, we estimate a maximum Bell-pair generation rate of approximately 50 MHz, capable of 2 MHz error correction cycles. Estimating the number of qubits needed for high-speed fault-tolerant communication, we discuss tradeoffs between communication speed and qubit overheads in each approach.

## I. COMMUNICATION FIDELITY REQUIREMENTS FOR FAULT-TOLERANT SCALABILITY

A significant challenge for all quantum error-corrected modular architectures is noise in the quantum communication between modules. Typically, architectures based on the surface code target both local gate and network noise levels of 0.1% [12–15] to ensure all operations occur around 10 $\times$  below the usual surface-code threshold of 1%. While Rydberg gates have recently surpassed the 1% threshold, the lowest Bell-pair infidelity achieved to date in atomic qubits lags behind at 6–10% [25,33] (see Ref. [34] for a review). Additionally, teleported gates require several local operations to be executed on the Bell pair, further increasing network noise. To reduce network noise, entanglement distillation has been proposed as a method to convert many low-fidelity Bell pairs into a smaller number of higher-fidelity Bell pairs [35]. However, even simple protocols involve many local gates and, therefore, can only distill down to an error rate approximately ten times larger than the local gate errors [36,37]. More sophisticated protocols are necessary to reduce the noise

further, but the added complexity lowers the success rate, reducing the achievable error-correction cycle rate and increasing the memory errors per cycle [9,11,38–41].

As an alternative to distillation, it has been previously shown [16,42,43] that the surface code is intrinsically robust to higher levels of noise along a lower-dimensional boundary. As shown in Fig. 1, this boundary is where the physical operations required to fault-tolerantly connect code patches occur, for example, using lattice surgery to perform a logical CNOT gate [29]. For noise confined to the boundary, the boundary can be thought of as acting like an embedded lower-dimensional code with correspondingly less stringent thresholds [16,42–44]. Regardless of the degree of noise on the boundary and the bulk, most logical failures are isolated to either the lower-dimensional boundary or the bulk. Consequently, the thresholds for boundary and bulk gates are almost entirely independent. One can then approximate the total logical error rate as arising from logical errors occurring on the bulk and on the boundary independently:

$$P_{\text{fail}}(p_{\text{bulk}}, p_{\text{bound}}) \sim \left( \frac{p_{\text{bulk}}}{p_{\text{th}}^{\text{bulk}}} \right)^{L/2} + \left( \frac{p_{\text{bound}}}{p_{\text{th}}^{\text{bound}}} \right)^{L/2}, \quad (1)$$

where errors not confined to either the bulk or the boundary are negligible, and the thresholds obey  $p_{\text{th}}^{\text{bound}} > p_{\text{th}}^{\text{bulk}}$ . From Eq. (1), we can see that errors on the boundary will not dominate the total error until  $p_{\text{bound}}/p_{\text{th}}^{\text{bound}}$  is comparable to  $p_{\text{bulk}}/p_{\text{th}}^{\text{bulk}}$ .

Given the long coherence time of neutral atoms, errors occur primarily due to imperfect local Rydberg gates (with probability  $\epsilon_{\text{Ryd}}$ ), local fluorescence readout ( $\epsilon_M$ ), and nonlocal Bell pairs ( $\epsilon_{\text{Bell}}$ ). We model this noise by counting up the probabilities  $p$  and  $q$  per error-correction cycle for the sum of these processes to create a bit-flip error on data and

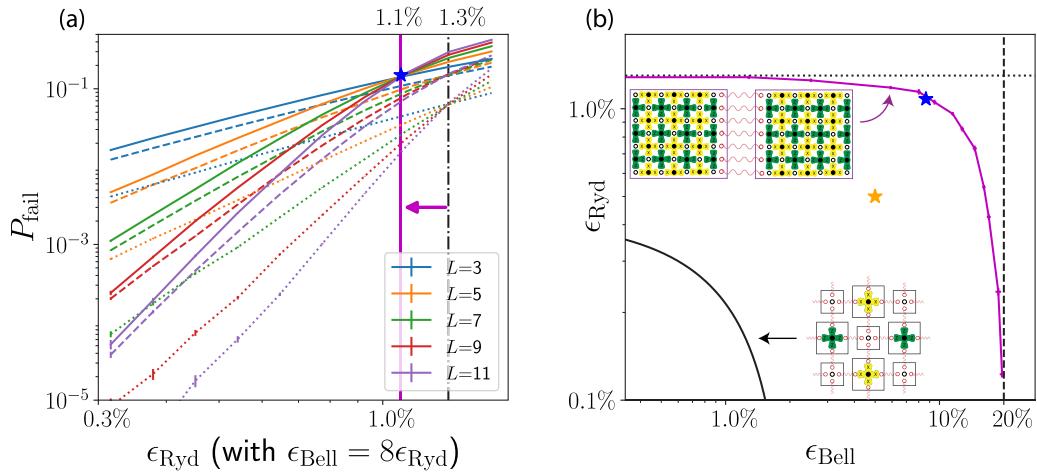


FIG. 2. (a) Dotted:  $P_{\text{fail}}$  for a bulk surface-code patch with *no boundary* errors (in terms of  $\epsilon_{\text{Ryd}}$ ). Dashed: plotting *only the boundary* logical errors in terms of  $\epsilon_{\text{Ryd}}$ , fixing  $\epsilon_{\text{Bell}} = 8\epsilon_{\text{Ryd}}$ ,  $\epsilon_M = \epsilon_{\text{Ryd}}$  to shift the boundary threshold to equal the bulk threshold (black dot-dashed vertical line). Solid:  $P_{\text{fail}}$  with *both bulk and boundary* errors simultaneously. (b) Magenta: the threshold for  $P_{\text{fail}}$  plotted in the  $(\epsilon_{\text{Ryd}}, \epsilon_{\text{Bell}})$  space. Points below and left of the curves are below threshold. Blue star indicates threshold in (a), orange star indicates best experimental Rydberg gate and nonlocal Bell-pair infidelity achieved separately to date [3,25]. Black: for reference, the curve  $3\% = 1.5\epsilon_{\text{Bell}} + 7\epsilon_{\text{Ryd}}$ , approximates the expected threshold if all parity checks in the bulk were carried out with teleported gates.

syndrome qubits, respectively, both in the bulk and on the boundary (see Appendixes for more information about the error model), with the results summarized in Table I. While noise in the bulk comes solely from local operations necessary to measure the local check operators, Bell-pair noise and additional local operations for enacting the nonlocal teleported gate circuit contribute to elevated noise along the boundary.

Using this Rydberg error model, we perform Monte Carlo simulations of surface codes with phenomenological errors in the bulk and on the boundary [45], and plot the logical failure rate  $P_{\text{fail}}$  in Fig. 2(a) in terms of  $\epsilon_{\text{Ryd}}$ , choosing  $\epsilon_{\text{Ryd}} = \epsilon_M$  and having fixed the ratio  $\epsilon_{\text{Bell}} = 8\epsilon_{\text{Ryd}}$ . The cases of only boundary noise (dashed), only bulk noise (dotted), and both boundary and bulk noise occurring simultaneously (solid) are each shown. The elevated noise on the boundary, due to both the additional local operations necessary to operate the teleported gate as well as the noise from the Bell pairs, only modestly shifts the  $\epsilon_{\text{Ryd}}$  threshold. Figure 2(b) shows how the same threshold depends on the choice of parameters  $\epsilon_{\text{Ryd}}$ ,  $\epsilon_{\text{Bell}}$ , with the blue star identifying the point  $\epsilon_{\text{Ryd}} = 1.1\%$ ,  $\epsilon_{\text{Bell}} = 8.8\%$  corresponding to the threshold in Fig. 2(a). To highlight how favorable these threshold requirements are, the orange star identifies the best remote atomic Bell pair and Rydberg gate fidelities experimentally achieved so far in separate experiments [3,25], while, for comparison, the black line shows the expected thresholds if each module contains only a couple of qubits and all parity checks in the

bulk are carried out with teleported gates (see Appendixes for a detailed description of the error model).

## II. COMMUNICATION SPEED REQUIREMENTS FOR FAULT-TOLERANT SCALABILITY

Having shown a high degree of robustness to noise along the interface between modules, successful error correction now only requires executing error correction cycles sufficiently fast. To ensure that background decoherence errors remain negligible compared to other noise sources, we require a cycle time  $T$  satisfying  $T/\tau_{\text{dec}} \ll p_{\text{bulk}}$ , where a decoherence time  $\tau_{\text{dec}} \sim 2$  s has been demonstrated [22,46]. Beyond decoherence, faster error correction cycles are desirable for improving logical clock speeds. The rate of generating the  $2L$  Bell pairs required per cycle, where  $L$  is the surface-code distance, is expected to bottleneck  $T$ . (The fastest experimentally achieved remote entanglement generation rate in neutral atoms or trapped ions is currently 250 Hz [47].) However, through either multiplexing entanglement generation in free space, or through the use of a single optical cavity we show that it should be possible to generate the  $2L$  Bell pairs per error correction cycle much faster than necessary to interface logical qubits of size and quality sufficient for large-scale applications. We also show that even higher rates are possible if entanglement generation can be further parallelized using an array of microcavities and compare rates and qubit overhead for each approach.

For generating nonlocal Bell pairs, we consider remote entanglement protocols where atom-photon entanglement is swapped to atom-atom entanglement [25,47–49]. Atom-atom entanglement is probabilistic but heralded by the successful detection of two photons. For concreteness, we consider the Rubidium atom-atom entanglement protocol from Refs. [32,50,51], operating on the D2  $F = 1$  to  $F' = 0$  transition.

TABLE I. Bit-flip error rates associated with bulk and boundary qubits under a generic, unbiased noise model.

	Bulk	Boundary
$p$	$2\epsilon_{\text{Ryd}}$	$0.5\epsilon_{\text{Bell}} + 2.5\epsilon_{\text{Ryd}} + \epsilon_M$
$q$	$2\epsilon_{\text{Ryd}} + \epsilon_M$	$0.5\epsilon_{\text{Bell}} + 2.5\epsilon_{\text{Ryd}} + 2\epsilon_M$

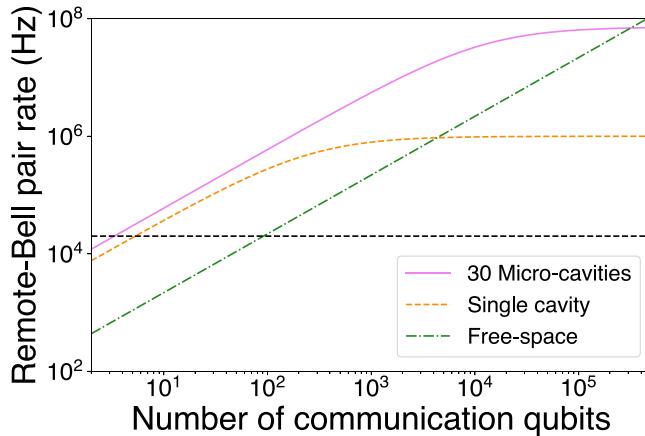


FIG. 3. The average Bell-pair rate depends on the number of qubits available for communication and the interconnect design. The dashed black line indicates the rate required to execute error-correction cycles every 2 ms for one distance  $L = 20$  logical qubit. For a single cavity (orange dashed line), above 250 qubits the rate becomes limited by addressing beam switching time. Higher rates (purple solid line) are possible using an array of microcavities. In this case, the rate is eventually limited by the switching time and by the number of microcavities, which larger microscope objectives could further increase. In free space (green dot-dashed line), lower collection efficiency means more qubits are necessary to achieve the same Bell pair rate.

#### A. Free-space collection

Considering first just a single atom, the single-photon emission success probability is given by the mode overlap of the dipole electric field of the atom and a Gaussian mode. Assuming a numerical aperture of 0.7, the collection efficiency  $\eta^{\text{lens}} = 0.12$  [32]. The atom-atom entanglement generation probability per attempt  $P_{\text{aa}}$  is then given by

$$P_{\text{aa}}^{\text{lens}} = \frac{1}{2}(\eta^{\text{lens}}\eta_{\text{det}})^2, \quad (2)$$

where  $\eta_{\text{det}}$  is the subsequent detection efficiency. We assume  $\eta_{\text{det}} = 0.7$ , making the atom-atom entanglement generation success probability per attempt  $P_{\text{aa}}^{\text{lens}} = 0.0035$ .

With a single qubit, the repetition rate is limited by the time required to reset the qubit after each failed attempt (about 6  $\mu\text{s}$ ) and the time required to occasionally recool the qubits (averaging to 10  $\mu\text{s}$  per attempt) [32]. The resulting mean time to successfully prepare a Bell pair is then 16  $\mu\text{s}/0.0035 = 4.6$  ms. Faster rates are possible using a workflow where atom reinitialization and cooling are carried out elsewhere to allow entanglement attempts to be repeated more frequently. In addition, entanglement can be attempted on multiple atoms simultaneously using a detector array (see Fig. 3). In this case, the Bell-pair generation rate is limited only by the number of atoms which are available and which can be simultaneously controlled and/or the number of detectors.

As summarized in Table II, the cycle time  $T$ , dominated by the time to distribute Bell pairs, will then be  $T = 2L\tau_{\text{Bell}}$ , the time to create the  $2L$  Bell pairs along the interface per cycle.  $T$  grows with  $L$ , but as a reference point, with local gates and boundary noise operating 10 $\times$  below the threshold, as in Eq. (1), logical qubits with distance  $L = 20$  have logical errors sufficiently low for large-scale applications [52]. In

TABLE II. Quantities relevant for estimating QEC cycle rates for a code patch with distance  $L = 20$  when using a lens, a single cavity, or a microcavity array with the parameters from Table III. For free space, times are reported per communication qubit. Where an optical cavity is used, the times reported are the asymptotic values from Fig. 3 limited by the number (1 or 30) of optical modes. Bell pairs are denoted BP.

		Free space	1 cavity	30 cavities
Code distance	$L$	20	20	20
Qubit # $N$	$\sim 4L^2$	1521	1521	1521
BP attempt time	$\tau$	16 $\mu\text{s}/\text{qubit}$	0.10 $\mu\text{s}$	3.3 ns
BP time $\tau_{\text{Bell}}$	$\tau/P_{\text{aa}}^{\text{cav}}$	4.6 ms/qubit	1.0 $\mu\text{s}$	14 ns
Cycle time $T$	$2L\tau_{\text{Bell}}$	180 ms/qubit	40 $\mu\text{s}$	560 ns
Log. gate time	$2L^2\tau_{\text{Bell}}$	3.7 s/qubit	0.8 ms	11.2 $\mu\text{s}$
Decoh. time	$\tau_{\text{dec}}$	2.0 s	2.0 s	2.0 s

this case,  $\tau_{\text{Bell}} = 4.6$  ms/atom, so control of  $\sim 100$  qubits is sufficient to produce 40 Bell pairs in 2 ms, fast enough so that  $T/\tau_{\text{dec}} = 10^{-3}$ . Implementing a logical gate requires  $L$  code cycles, giving a logical clock speed of 25 Hz.

#### B. Single large-volume cavity

By enhancing collection efficiency, an optical cavity can create Bell pairs faster than in free space, enabling faster logical clock speeds with fewer communication qubits. We adopt the “medium near-concentric” cavity design proposed in Ref. [32], where to be compatible with Rydberg excitation the cavity length is conservatively chosen to be 4 mm with a 5  $\mu\text{m}$  waist. The remaining experimental parameters are summarized in Table III. The cavity enhances collection efficiency from the atom, achieving  $\eta^{\text{cav}} = 0.66$ . Assuming the same detection efficiency of 0.7, the atom-atom entanglement generation probability per attempt is  $P_{\text{aa}}^{\text{cav}} = 0.1$ .

If only one qubit is available for entanglement generation, Ref. [32] previously found that entanglement rates of about 5800 Hz are possible, limited by the 16  $\mu\text{s}$  required to reset and occasionally recool qubits between subsequent attempts. To achieve faster rates of remote entanglement, we envision a

TABLE III. Cavity properties: for the single cavity, we adopt the “medium near-concentric” entry from Table II of Ref. [32]. For the microcavity array, we envision a cavity with a radius of curvature of  $\sim 100$   $\mu\text{m}$ , allowing for a large cavity bandwidth and small mode volume. The small size of the mode allows for these cavities to be densely packed.

	Single cavity	Microcavity array
# of cavities	1	30
Length	4 mm	90 $\mu\text{m}$
Waist	5 $\mu\text{m}$	2.5 $\mu\text{m}$
$(g, \kappa, \Gamma)$	(17, 28, 6 MHz)	(184, 404, 6 MHz)
Cooperativity $C$	3.2	56
$\eta^{\text{cav}}$	0.66	0.98
$\eta_{\text{det}}$	0.7	0.7
$P_{\text{aa}}^{\text{cav}}$	0.1	0.24
$\tau_{\text{cav}}$	5.7 ns	0.39 ns

module design that utilizes the coherent transport capabilities of neutral atoms [22,30,31] to temporally multiplex remote entanglement generation [14,53]. In each module, depicted in Fig. 1, fresh qubits are initialized and transported through the cavity, where attempts are made to establish remote entanglement. If successful, the qubits comprising the Bell pair are transported to the edge of the code patch (red lines in Fig. 1) where they will be used to enact syndrome checks across the boundary.

The speed at which atoms can be transported through the cavity mode is limited by trap depth (available laser power) and heating to approximately  $10 \mu\text{m}/\mu\text{s}$  [22]. With a  $5 \mu\text{m}$  waist, atoms will spend much longer in the cavity mode than is required to attempt remote entanglement. To increase the attempt rate, off-resonant addressing beams focused to  $1 \mu\text{m}$  can be used to Stark shift atoms as they exit and enter the central region of the cavity mode, allowing for a more closely spaced stream of qubits and an attempt every  $100 \text{ ns}$ . An additional possibility proposed in Ref. [12] is to transport qubits in batches into the cavity and attempt remote entanglement sequentially (relying again on individually addressed Stark shifts and now limited by  $100 \text{ ns}$  switching time). Once a qubit has been successfully entangled, it must be transported to the computation zone and back before it can be used again. This is the slowest step, taking approximately  $100 \mu\text{s}$ , but it is only required whenever entanglement generation succeeds.

Applying these methods with a  $P_{aa}^{\text{cav}}$  success probability of 0.1, gives an attempt time per qubit of  $26 \mu\text{s}$ . Therefore, about 260 communication qubits (see Fig. 3) should make it possible to reach an attempt rate limited by the switching time of the addressing beams, enabling remote entanglement attempts approximately every  $100 \text{ ns}$  and a maximum mean Bell-pair time of  $\tau_{\text{Bell}} = 1 \mu\text{s}$ . The time to produce the 40 Bell pairs required to implement an error-correction cycle is then  $T = 40 \mu\text{s}$ , resulting in a  $25 \text{ kHz}$  code cycle rate and a logical clock speed of  $1.25 \text{ kHz}$ .

### C. Microcavity array

With one cavity, the availability of a single optical mode presents a fundamental limit on the Bell-pair rate, as entanglement attempts must be made sequentially. This limit can be increased through interfacing multiple cavities to the same array [54–56]. However, field-of-view constraints require a different approach to cavity design. The small size of microcavities allows for them to be densely packed into a microcavity array. The main considerations for the density are the clipping loss of the cavity waist on the microcavity mirror and the total length of the cavity including the substrate.

We propose microcavities requiring a  $50 \mu\text{m}$  mirror diameter and a substrate less than  $100 \mu\text{m}$  thick. With a  $1.5 \text{ mm}$  field-of-view microscope objective, about 30 cavities can be fit side by side. Stabilization can then be performed through a combination of microelectromechanical systems (MEMS) and thermal tuning to allow for independent locking of each high-finesse cavity [55,57]. The resulting cavities are  $90 \mu\text{m}$  in length, and are tightly focused with a waist of  $2.5 \mu\text{m}$  enabling a large cooperativity of  $C = 56$ . To ensure compatibility with Rydberg excitation, doped silicon should be used for the substrate, reducing charge noise by making the bulk of

the structure conductive. Additionally, if necessary, Rydberg gates can be performed in the computation zone away from the cavities. For a summary of their other properties, see Table III.

The parallelization afforded by this array of 30 microcavities (as opposed to the single cavity from the previous section) directly improves the maximum achievable entanglement attempt rate by a factor of 30 as long as sufficient communication qubits ( $\sim 7500$ ) are available so that transport times do not bottleneck the process, see Fig. 3. Additionally, because of the higher cooperativity, the success probability also increases by approximately a factor of 2, overall resulting in 60 times higher Bell-pair rates. Per Table II, the average time to create a Bell pair is then  $\tau_{\text{Bell}} = 14 \text{ ns}$ , resulting in a  $2 \text{ MHz}$  QEC cycle rate and a logical clock speed of  $100 \text{ kHz}$ .

### D. Comparison of rates and qubit resources

In Fig. 3, we compare the qubit resources required to achieve a certain remote entanglement rate for our three approaches. We plot the remote Bell-pair rate versus the number of communication qubits for a lens (dot-dashed line), a single large-volume cavity (dashed line), and a microcavity array (solid line). For the comparison, the cavity parameters used are given in Table III (see Appendixes for a detailed description of the rates shown). The corners at around 260 communication qubits for the single large-volume cavity and around 7500 communication qubits for the microcavity array show the point at which the rate becomes limited by the number of optical modes. The dashed black line indicates a Bell-pair rate sufficiently fast to execute error correction cycles every  $2 \text{ ms}$  for one distance  $L = 20 \text{ logical qubit}$ , satisfying the requirement that  $T/\tau_{\text{dec}} \leq 10^{-3}$ .

## III. SUPERMODULES

To reduce the number of total modules required, supermodules can be constructed comprising multiple arrays in one vacuum chamber connected by optical lattice conveyor belts. Each array (or “submodule”) would be separately controlled and under its own microscope, with the crucial challenge being fast and high-fidelity communication between submodules via atom transport. For the former, many Bell pairs would be created deterministically using parallelized local Rydberg gates. Then, one qubit out of each pair would be accelerated to match the speed of the optical lattice conveyor belt, placed in the lattice, and transported to a distant array.

Because Bell pairs can be generated using Rydberg gates, they could be initially prepared with fidelity exceeding 99%. Therefore, the predominant source of Bell pair infidelity would be decoherence during transport. Transport times are set by the conveyor-belt speed, which is limited by the maximum velocity that qubits can reach without excessive heating ( $\sim 1 \mu\text{m}/\mu\text{s}$ ). Travel times between arrays separated by  $10 \text{ cm}$  (limited by space required for microscope objectives) would therefore be approximately  $100 \text{ ms}$ , resulting in decoherence below the 10% Bell-pair threshold from Fig. 2(b).

## IV. CONCLUSION

Combining unique atom-array capabilities with the recently reported robustness of surface codes to boundary

noise, we have proposed a path to fault-tolerant scaling of error-corrected modules requiring only current-to-near-term technology. Rydberg gate errors below 1% have been achieved [3–5] as well as remote entanglement of atomic qubits with infidelity below 10% [25,47]. Additionally, theoretical analysis of the expected fidelity for remote Bell pairs indicates further improvements can be expected [12,32]. Generating Bell pairs sufficiently fast remains a challenge [12,32,34], but we have described several promising approaches for realizing sufficient Bell-pair rates using free-space collection, a large-volume optical cavity, or a microcavity array. Injecting physical Bell pairs into logical qubits [14] and running more sophisticated distillation schemes [58,59] could require fewer physical Bell pairs per logical gate, but would incur additional space-time overheads, and further work is warranted to investigate these tradeoffs. Further gains may also be possible in certain regimes by using transversal gates between modules [60] and exploiting algorithmic fault tolerance [61,62]. By lowering the bar for communication fidelity and outlining multiple paths towards achieving sufficiently fast photonic interconnects, the present analysis motivates the realization and exploration of networked logical quantum processors that satisfy the requirements for scalable fault tolerance in the near term.

#### ACKNOWLEDGMENTS

This project was funded in part by DARPA under the ONISQ program (Grant No. W911NF2010021), the US Department of Energy (Quantum Systems Accelerator, Grant No. DE-AC02-05CH11231), the Center for Ultracold Atoms (an NSF Physics Frontiers Center, Grant No. PHY-1734011), QuSeC-TAQS from NSF (Grant No. 2326787), (Grant No. W911NF2320219), the Army Research Office (Grant No. W911NF2010082), IARPA under the ELQ program (Grant No. W911NF2320219), and QuEra Computing. D.B. acknowledges support from the NSF Graduate Research Fellowship Program (Grant No. DGE1745303) and The Fannie and John Hertz Foundation.

#### APPENDIX A: TELEPORTED GATE

When the interface is realized via distributed entanglement, that entanglement serves as a resource for enacting nonlocal, teleported gates between qubits in distinct modules. Figure 4 shows how bit- and phase-flip noise on the distributed Bell pair propagates to the control and target qubits in the distinct modules that the (pink) teleported gate acts on. The propagation

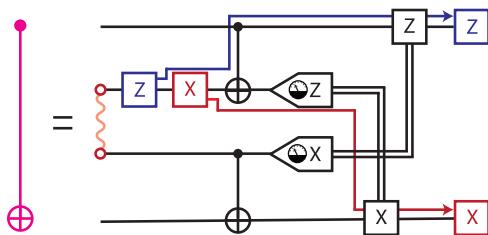


FIG. 4.  $X, Z$  errors on the Bell pair used in teleported gate propagate to the two qubits it operates on.

is identical for errors occurring on either of the Bell-pair qubits, as must be the case since the Bell pair is invariant under application of  $XX$  and  $ZZ$ .

#### APPENDIX B: MICROSCOPIC SEAM ERROR MODEL

For Bell pairs and CNOT gates, we model each as a perfect operation followed by all combinations of  $X, Y, Z$  errors on the two qubits, each with error  $\epsilon_B/15$  or  $\epsilon_{CX}/15$ :  $IX, IY, IZ, XI, XX, XY, XZ, YI, YX, YY, YZ, ZI, ZX, ZY, ZZ$  [42]. As the Bell pair is stabilized under the application of  $XX, ZZ$ , all 15 types of errors are equivalent to  $II$  with  $\frac{1}{5}\epsilon_B$  and  $IX, IY$ , or  $IZ$  with probability  $\frac{4}{15}\epsilon_B$ . Similarly, as  $Y \propto XZ$ , we can count up the total probability a given qubit gets hit by  $X$  and by  $Z$  errors following a CNOT:  $\frac{8}{15}\epsilon_{CX}$ .

Using this error model, we first consider a standard surface-code patch. In this case there is no seam and errors only arise from CNOT gate errors and readout errors. As each data qubit is subject to four parity-check operations per code cycle, the bit-flip probability for data qubits is  $p_{\text{bulk}} = 4 \times \frac{8}{15}\epsilon_{CX} \approx 2\epsilon_{CX}$ . Similarly for syndrome qubits, but with an additional error associated with syndrome qubit readout,  $q_{\text{bulk}} = 4 \times \frac{8}{15}\epsilon_{CX} + \epsilon_M \approx 2\epsilon_{CX} + \epsilon_M$ . To determine the dependence of the logical error probability on the code distance and local gate errors, we use a Monte Carlo simulation of errors ( $p_{\text{bulk}}$  and  $q_{\text{bulk}}$ ) and a local minimum weight perfect matching decoder [45], exhibiting a threshold of 1.3% (see main text Fig. 2).

Next, we consider the situation depicted in main text Fig. 1 where two surface-code patches are being merged via teleported gates spanning the seam. As shown, stretching the parity-check operators across the seam involves one teleported gate per (pink) data/syndrome qubit pair on the seam. Figure 4 shows how a teleported gate propagates bit and phase ( $X$  and  $Z$ ) errors occurring on a Bell pair. Evidently, bit-flip errors on the Bell pair propagate exclusively to the target qubit. Similarly, phase-flip errors on the Bell pair propagate exclusively to the control qubit. In total, the bit-flip probability on the control qubit is  $\frac{8}{15}\epsilon_{CX}$ , whereas the bit-flip probability on the target is  $\frac{8}{15}\epsilon_B + 2 \times \frac{8}{15}\epsilon_{CX} + \epsilon_M$ . Similarly, the phase-flip probability on the control qubit is  $\frac{8}{15}\epsilon_B + 2 \times \frac{8}{15}\epsilon_{CX} + \epsilon_M$ , whereas the phase-flip probability on the target is  $\frac{8}{15}\epsilon_{CX}$ . Per code cycle, each (pink) qubit along the seam experiences three local CNOTs followed by one teleported gate. As the seam shown in main text Fig. 1 is along the  $X_L$  direction, it is most susceptible to logical bit-flip errors  $X_L$  arising from bit flips along the length of the seam, specifically on the pink qubits in code patch 2. In contrast, although phase-flip errors are also occurring with elevated probability along the length of the vertical seam (specifically on the pink qubits in code patch 1), they contribute little to logical phase errors, which correspond to horizontal strings.

For qubits on the seam (pink), our phenomenological weighted error model is as follows. A plaquette-syndrome qubit on code patch (CP) 2 is the target of three local CNOTs ( $3 \times \frac{8}{15}\epsilon_{CX}$ ), as well as the target of one teleported gate ( $\frac{8}{15}\epsilon_B + 2 \times \frac{8}{15}\epsilon_{CX} + \epsilon_M$ ) and one final readout ( $\epsilon_M$ ). The resulting bit-flip probability for a plaquette-syndrome qubit in code patch 2 is  $q_{\text{seam}} = 3 \times \frac{8}{15}\epsilon_{CX} + (\frac{8}{15}\epsilon_B + 2 \times$

TABLE IV. Phenomenological bit-flip error probabilities per code cycle  $p$  and  $q$  on data and syndrome qubits. Entries describe how during a given code cycle, local operations and Bell pairs to the total phenomenological error probability. Phase-flip error rates are identical. “Bulk” and “Seam” columns correspond to regions depicted in main text Fig. 1, and for comparison, the “Small modules” column shows the case where all gates in a surface code are done with teleported gates.

	Bulk	Seam	Small modules
$p$	$2\epsilon_{\text{CX}}$	$0.5\epsilon_B + 2.5\epsilon_{\text{CX}} + \epsilon_M$	$\epsilon_B + 3\epsilon_{\text{CX}} + 2\epsilon_M$
$q$	$2\epsilon_{\text{CX}} + \epsilon_M$	$0.5\epsilon_B + 2.5p_{\text{CX}} + 2\epsilon_M$	$2\epsilon_B + 4\epsilon_{\text{CX}} + 5\epsilon_M$

$\frac{8}{15}\epsilon_{\text{CX}} + \epsilon_M) + \epsilon_M \approx 0.5\epsilon_B + 2.5\epsilon_{\text{CX}} + 2\epsilon_M$ . A data qubit on the seam of code patch 2 experiences identical bit-flip probability to a plaquette-syndrome qubit in the same patch, less one final readout,  $p_{\text{seam}} = 3 \times \frac{8}{15}\epsilon_{\text{CX}} + (\frac{8}{15}\epsilon_B + 2 \times \frac{8}{15}\epsilon_{\text{CX}} + \epsilon_M) + \epsilon_M \approx 0.5\epsilon_B + 2.5\epsilon_{\text{CX}} + \epsilon_M$ . Additionally, data qubits in code patch 2 experience phase-flip errors at the same rate as bulk qubits. Correspondingly, similar considerations apply to data and star syndrome qubits in code patch 1, which experience elevated levels of phase-flip errors. Table I in the main text summarizes the noise.

### APPENDIX C: SMALL MODULES ERROR MODEL

For comparison, in minimal-size module configuration, where each module contains a single data or syndrome qubit and just one communication qubit, teleported gates must be used for each two-qubit operation in all the check operators, leading to substantially worse performance. This configuration is illustrated in the cartoon in main text Fig. 2(b).

Each code cycle, a given data qubit is the target of 2 teleported CNOTs and control of 2 teleported CNOTs, giving a bit-flip probability per code cycle of  $p_{\text{small}} = 2 \times (8\epsilon_B/15 + 2 \times 8\epsilon_{\text{CX}}/15 + \epsilon_M) + 2 \times (8\epsilon_{\text{CX}}/15) \approx \epsilon_B + 3\epsilon_{\text{CX}} + 2\epsilon_M$ .

TABLE V. Entanglement generation success probability and attempt rate for each approach described in the main text. Multiplying  $P_{\text{aa}}$  by  $R_{\text{attempt}}(N)$  yields the curves shown in main text Fig. 3.

	$P_{\text{aa}}$	$R_{\text{attempt}}(N) [\mu\text{s}^{-1}]$
Free space	0.0035	$1/(16/N + P_{\text{aa}} \cdot 100/N)$
Single cavity	0.1	$1/(0.1 + 16/N + P_{\text{aa}} \cdot 100/N)$
30 microcavities	0.24	$1/(0.003 + 16/N + P_{\text{aa}} \cdot 100/N)$

Similarly, consider a plaquette-syndrome qubit, which is the target of four teleported CNOTs, then followed with additional readout error  $\epsilon_M$ . This plaquette has an error probability of  $q_{\text{small}} = 4 \times (8\epsilon_B/15 + 2 \times 8\epsilon_{\text{CX}}/15 + \epsilon_M) + \epsilon_M \approx 2\epsilon_B + 4\epsilon_{\text{CX}} + 5\epsilon_M$ . The result is the same for star-operator-syndrome qubits.

Also, since the threshold would occur at  $p = 3\%$  for  $q = p$ , expressing the threshold in terms of  $\epsilon_B$  and  $\epsilon_{\text{CX}}$  amounts to determining to what extent each of these contributes to  $p$ . In fact, taking for simplicity  $\epsilon_{\text{CX}} = \epsilon_M$  and averaging the small modules expressions for  $p$  and  $q$  from Table IV gives a reasonable approximation for the threshold relationship  $3\% = 1.5\epsilon_B + 7\epsilon_{\text{CX}}$ . An exact simulation with the small modules weights for  $p$  and  $q$  gives results quite close to this.

### APPENDIX D: RATE COMPARISON

In main text Fig. 3, the remote Bell-pair rate for all approaches is the entanglement attempt rate  $R_{\text{attempt}}$  multiplied by atom-atom entanglement generation probability per attempt  $P_{\text{aa}}$ . With only a few communication qubits, the attempt rate is limited by the time to repump, recool, and transport atoms. The rate increases with the number of qubits through a combination of temporal and spatial multiplexing, until it is limited, when applicable, by the number of optical modes or, for free space, by the number of detectors. The success probabilities and the attempt rates for each approach are summarized in Table V.

- [1] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, *Phys. Rev. A* **52**, R2493(R) (1995).
- [2] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 10th ed. (Cambridge University Press, Cambridge, 2011).
- [3] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, H. Levine, G. Semeghini, M. Greiner, V. Vuletić, and M. D. Lukin, High-fidelity parallel entangling gates on a neutral-atom quantum computer, *Nature (London)* **622**, 268 (2023).
- [4] M. Peper, Y. Li, D. Y. Knapp, M. Bileksa, S. Ma, G. Liu, P. Peng, B. Zhang, S. P. Horvath, A. P. Burgers, and J. D. Thompson, Spectroscopy and modeling of Rydberg states for high-fidelity two-qubit gates, *Phys. Rev. X* **15**, 011009 (2025).
- [5] R. B.-S. Tsai, X. Sun, A. L. Shaw, R. Finkelstein, and M. Endres, Benchmarking and fidelity response theory of high-fidelity rydberg entangling gates, *PRX Quantum* **6**, 010331 (2025).
- [6] F. Gyger, M. Ammenwerth, R. Tao, H. Timme, S. Snigirev, I. Bloch, and J. Zeiher, Continuous operation of large-scale atom arrays in optical lattices, [arXiv:2402.04994](https://arxiv.org/abs/2402.04994).
- [7] M. A. Norcia, H. Kim, W. B. Cairncross, M. Stone, A. Ryou, M. Jaffe, M. O. Brown, K. Barnes, P. Battaglino, T. C. Bohdanowicz, A. Brown, K. Cassella, C. A. Chen, R. Coxe, D. Crow, J. Epstein, C. Griger, E. Halperin, F. Hummel, A. M. W. Jones *et al.*, Iterative assembly of  $^{171}\text{Yb}$  atom arrays with cavity-enhanced optical lattices, [arXiv:2401.16177](https://arxiv.org/abs/2401.16177).
- [8] 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 (London)* **626**, 58 (2024).

[9] N. H. Nickerson, Y. Li, and S. C. Benjamin, Topological quantum computing with a very noisy network and local error rates approaching one percent, *Nat. Commun.* **4**, 1756 (2013).

[10] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim, Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects, *Phys. Rev. A* **89**, 022317 (2014).

[11] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links, *Phys. Rev. X* **4**, 041041 (2014).

[12] Y. Li and J. Thompson, High-rate and high-fidelity modular interconnects between neutral atom quantum processors, [arXiv:2401.04075](https://arxiv.org/abs/2401.04075).

[13] H. Leone, S. Srikara, P. P. Rohde, and S. Devitt, Upper bounds for the clock speeds of fault-tolerant distributed quantum computation using satellites to supply entangled photon pairs, *Phys. Rev. Res.* **5**, 043302 (2023).

[14] S. Sunami, S. Tamiya, R. Inoue, H. Yamasaki, and A. Goban, Scalable networking of neutral-atom qubits: Nanofiber-based approach for multiprocessor fault-tolerant quantum computers, *PRX Quantum* **6**, 010101 (2025).

[15] S. Singh, F. Gu, S. de Bone, E. Villaseñor, D. Elkouss, and J. Borregaard, Modular architectures and entanglement schemes for error-corrected distributed quantum computation, [arXiv:2408.02837](https://arxiv.org/abs/2408.02837).

[16] J. Ramette, J. Sinclair, N. P. Breuckmann, and V. Vuletić, Fault-tolerant connection of error-corrected qubits with noisy links, *npj Quantum Inf.* **10**, 58 (2024).

[17] T. Wilk, A. Gaëtan, C. Evellin, J. Wolters, Y. Miroshnychenko, P. Grangier, and A. Browaeys, Entanglement of two individual neutral atoms using Rydberg blockade, *Phys. Rev. Lett.* **104**, 010502 (2010).

[18] L. Isenhower, E. Urban, X. L. Zhang, A. T. Gill, T. Henage, T. A. Johnson, T. G. Walker, and M. Saffman, Demonstration of a neutral atom controlled-not quantum gate, *Phys. Rev. Lett.* **104**, 010503 (2010).

[19] Y.-Y. Jau, A. M. Hankin, T. Keating, I. H. Deutsch, and G. W. Biedermann, Entangling atomic spins with a Rydberg-dressed spin-flip blockade, *Nat. Phys.* **12**, 71 (2016).

[20] T. M. Graham, M. Kwon, B. Grinkemeyer, Z. Marra, X. Jiang, M. T. Lichtman, Y. Sun, M. Ebert, and M. Saffman, Rydberg-mediated entanglement in a two-dimensional neutral atom qubit array, *Phys. Rev. Lett.* **123**, 230501 (2019).

[21] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin, Parallel implementation of high-fidelity multiqubit gates with neutral atoms, *Phys. Rev. Lett.* **123**, 170503 (2019).

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

[23] I. S. Madjarov, J. P. Covey, A. L. Shaw, J. Choi, A. Kale, A. Cooper, H. Pichler, V. Schkolnik, J. R. Williams, and M. Endres, High-fidelity entanglement and detection of alkaline-earth Rydberg atoms, *Nat. Phys.* **16**, 857 (2020).

[24] P. Maunz, S. Olmschenk, D. Hayes, D. N. Matsukevich, L.-M. Duan, and C. Monroe, Heralded quantum gate between remote quantum memories, *Phys. Rev. Lett.* **102**, 250502 (2009).

[25] 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.* **124**, 110501 (2020).

[26] V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D. A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B. P. Lanyon, and T. E. Northup, Entanglement of trapped-ion qubits separated by 230 meters, *Phys. Rev. Lett.* **130**, 050803 (2023).

[27] C. Nölleke, A. Neuzner, A. Reiserer, C. Hahn, G. Rempe, and S. Ritter, Efficient teleportation between remote single-atom quantum memories, *Phys. Rev. Lett.* **110**, 140403 (2013).

[28] S. Langenfeld, S. Welte, L. Hartung, S. Daiss, P. Thomas, O. Morin, E. Distante, and G. Rempe, Quantum teleportation between remote qubit memories with only a single photon as a resource, *Phys. Rev. Lett.* **126**, 130502 (2021).

[29] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, Surface code quantum computing by lattice surgery, *New J. Phys.* **14**, 123011 (2012).

[30] A. Lengwenus, J. Kruse, M. Schlosser, S. Tichelmann, and G. Birkl, Coherent transport of atomic quantum states in a scalable shift register, *Phys. Rev. Lett.* **105**, 170502 (2010).

[31] M. Schlosser, S. Tichelmann, J. Kruse, and G. Birkl, Scalable architecture for quantum information processing with atoms in optical micro-structures, *Quantum Info. Proc.* **10**, 907 (2011).

[32] C. B. Young, A. Safari, P. Huft, J. Zhang, E. Oh, R. Chinnarasu, and M. Saffman, An architecture for quantum networking of neutral atom processors *Appl. Phys. B* **128**, 151 (2022).

[33] L. Hartung, M. Seubert, S. Welte, E. Distante, and G. Rempe, A quantum-network register assembled with optical tweezers in an optical cavity, *Science* **385**, 179 (2024).

[34] J. P. Covey, H. Weinfurter, and H. Bernien, Quantum networks with neutral atom processing nodes *npj Quantum Inf.* **9**, 90 (2023).

[35] W. Dür and H.-J. Briegel, Entanglement purification for quantum computation, *Phys. Rev. Lett.* **90**, 067901 (2003).

[36] E. T. Campbell, Distributed quantum-information processing with minimal local resources, *Phys. Rev. A* **76**, 040302(R) (2007).

[37] S. Krastanov, V. V. Albert, and L. Jiang, Optimized entanglement purification, *Quantum* **3**, 123 (2019).

[38] L. Jiang, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, Distributed quantum computation based on small quantum registers, *Phys. Rev. A* **76**, 062323 (2007).

[39] Y. Li and S. C. Benjamin, High threshold distributed quantum computing with three-qubit nodes, *New J. Phys.* **14**, 093008 (2012).

[40] K. Fujii, T. Yamamoto, M. Koashi, and N. Imoto, A distributed architecture for scalable quantum computation with realistically noisy devices, [arXiv:1202.6588](https://arxiv.org/abs/1202.6588).

[41] R. Nigmatullin, C. J. Ballance, N. de Beaudrap, and S. C. Benjamin, Minimally complex ion traps as modules for quantum communication and computing, *New J. Phys.* **18**, 103028 (2016).

[42] A. G. Fowler, D. S. Wang, C. D. Hill, T. D. Ladd, R. Van Meter, and L. C. L. Hollenberg, Surface code quantum communication, *Phys. Rev. Lett.* **104**, 180503 (2010).

[43] S. J. S. Tan, C. A. Patterson, M. McEwen, and J. Preskill, Resilience of the surface code to error bursts, [arXiv:2406.18897](https://arxiv.org/abs/2406.18897).

[44] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, Topological quantum memory, *J. Math. Phys.* **43**, 4452 (2002).

[45] O. Higgott, Pymatching: A python package for decoding quantum codes with minimum-weight perfect matching (2021).

[46] S. Ma, A. P. Burgers, G. Liu, J. Wilson, B. Zhang, and J. D. Thompson, Universal gate operations on nuclear spin qubits in an optical tweezer array of  $^{171}\text{Yb}$  atoms, *Phys. Rev. X* **12**, 021028 (2022).

[47] J. O'Reilly, G. Toh, I. Goetting, S. Saha, M. Shalaev, A. Carter, A. Risinger, A. Kalakuntla, T. Li, A. Verma, and C. Monroe, Fast photon-mediated entanglement of continuously-cooled trapped ions for quantum networking, [arXiv:2404.16167](https://arxiv.org/abs/2404.16167).

[48] J.-W. Pan, D. Bouwmeester, H. Weinfurter, and A. Zeilinger, Experimental entanglement swapping: Entangling photons that never interacted, *Phys. Rev. Lett.* **80**, 3891 (1998).

[49] J. Schupp, V. Krermarsky, V. Krutyanskiy, M. Meraner, T. E. Northup, and B. P. Lanyon, Interface between trapped-ion qubits and traveling photons with close-to-optimal efficiency, *PRX Quantum* **2**, 020331 (2021).

[50] J. Hofmann, M. Krug, N. Ortegel, L. Gérard, M. Weber, W. Rosenfeld, and H. Weinfurter, Heralded entanglement between widely separated atoms, *Science* **337**, 72 (2012).

[51] W. Rosenfeld, D. Burchardt, R. Garthoff, K. Redeker, N. Ortegel, M. Rau, and H. Weinfurter, Event-ready bell test using entangled atoms simultaneously closing detection and locality loopholes, *Phys. Rev. Lett.* **119**, 010402 (2017).

[52] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).

[53] V. Krutyanskiy, M. Canteri, M. Meraner, V. Krermarsky, and B. P. Lanyon, Multimode ion-photon entanglement over 101 kilometers of optical fiber, [arXiv:2308.08891](https://arxiv.org/abs/2308.08891).

[54] M. Trupke, J. Goldwin, B. Darquié, G. Dutier, S. Eriksson, J. Ashmore, and E. A. Hinds, Atom detection and photon production in a scalable, open, optical microcavity, *Phys. Rev. Lett.* **99**, 063601 (2007).

[55] C. Derntl, M. Schneider, J. Schalko, A. Bittner, J. Schmiedmayer, U. Schmid, and M. Trupke, Arrays of open, independently tunable microcavities, *Opt. Express* **22**, 22111 (2014).

[56] G. Wachter, S. Kuhn, S. Minniberger, C. Salter, P. Asenbaum, J. Millen, M. Schneider, J. Schalko, U. Schmid, A. Felgner, D. Hüser, M. Arndt, and M. Trupke, Silicon microcavity arrays with open access and a finesse of half a million, *Light Sci. Appl.* **8**, 37 (2019).

[57] J. Gallego, S. Ghosh, S. K. Alavi, W. Alt, M. Martinez-Dorantes, D. Meschede, and L. Ratschbacher, High-finesse fiber fabry-perot cavities: Stabilization and mode matching analysis, *Appl. Phys. B* **122**, 47 (2016).

[58] C. Gidney, Tetrahedrally compact entanglement purification, [arXiv:2311.10971](https://arxiv.org/abs/2311.10971).

[59] C. A. Patterson, G. Baranes, J. P. B. Ataides, M. D. Lukin, and H. Zhou, Fast quantum interconnects via constant-rate entanglement distillation, [arXiv:2408.15936](https://arxiv.org/abs/2408.15936).

[60] See Supplementary Fig. 3 of Ref. [29].

[61] M. Cain, C. Zhao, H. Zhou, N. Meister, J. P. B. Ataides, A. Jaffe, D. Bluvstein, and M. D. Lukin, Correlated decoding of logical algorithms with transversal gates, [arXiv:2403.03272](https://arxiv.org/abs/2403.03272).

[62] H. Zhou, C. Zhao, M. Cain, D. Bluvstein, C. Duckering, H.-Y. Hu, S.-T. Wang, A. Kubica, and M. D. Lukin, Algorithmic fault tolerance for fast quantum computing, [arXiv:2406.17653](https://arxiv.org/abs/2406.17653).