Scalable Fault-Tolerant Quantum Technologies with Silicon Colour Centres

Stephanie Simmons on behalf of Photonic, Inc. (Dated: November 10, 2023)

The scaling barriers currently faced by both quantum networking and quantum computing technologies ultimately amount to the same core challenge of distributing high-quality entanglement at scale. In this Perspective, a novel quantum information processing architecture based on optically active spins in silicon is proposed that offers a combined single technological platform for scalable fault-tolerant quantum computing and networking. The architecture is optimized for overall entanglement distribution and leverages colour centre spins in silicon (T centres) for their manufacturability, photonic interface, and high fidelity information processing properties. Silicon nanophotonic optical circuits allow for photonic links between T centres, which are networked via telecom-band optical photons in a highly-connected graph. This high connectivity unlocks the use of low-overhead quantum error correction codes, significantly accelerating the timeline for modular, scalable faulttolerant quantum repeaters and quantum processors.

I. CONTEXT

Quantum information processing unlocks novel technological capabilities that cannot be achieved classically. To this end, there is a long and growing list of quantum algorithms [1], to be executed on future quantum computers, which offer super-polynomial speedups over their known classical counterparts. Some of these algorithms are critical for useful tasks including quantum chemistry and the design of novel materials as they enable precise simulations of chemical processes in important technological applications ranging from catalysts [2] to batteries [3] and pharmaceuticals [4]. Other quantum algorithms will impact our cybersecurity standards. For instance, once implemented on a large-scale fault-tolerant quantum computer, Shor's algorithm will allow efficient decryption of any data protected by RSA cryptography and its variants, which currently underpins >90% of all financial and internet traffic [5]. Interestingly, numerous so-called "quantum-safe" RSA-replacement encryption algorithms have already fallen to classical or quantum attacks, including Diffie-Hellman [6], Soliloquy [7], Elliptic Curve [8, 9], Rainbow [10], SIKE [11], and (through a side-channel) CRYSTALS-Kyber [12]. Some PQC algorithms under consideration have held up against attack so far, and have the potential for strong long-lasting security in practice, but similar to RSA, rest upon intrinsically unprovable computational security [13].

In parallel to quantum computing, quantum networks provide greater connectivity and security for quantum devices, and enable applications beyond the power of standalone quantum devices. These include secure communications [14] that even future quantum computers cannot break, "blind" computing of both classical and quantum algorithms [15], and modular quantum computing [16], timing [17], and sensing [18]. Analogously to today's high performance modular classical supercomputers, modular quantum supercomputers over quantum networks will enable astonishing computational capabilities by enabling true horizontal scaling of quantum resources, realizing practical quantum advantage for the aforementioned applications and giving rise to quantum applications yet to be imagined. Importantly, modular quantum processors connected by a quantum network directly offer all capabilities required of quantum repeaters, which are necessary to scale quantum networks' topologies and ultralong distance quantum networking applications.

However, today's noisy intermediate-scale quantum (NISQ) devices cannot be trusted to execute practicalscale networking and computational instructions beyond small instances of a few specific problems [19–21]. For each of today's monolithic quantum architectures it has been a challenge to deliver high-fidelity operations at a scale of even a few hundred physical qubits, and presently it seems that more (and perhaps substantially more) high-fidelity qubits will be necessary for commercial advantage. Similarly, the scaling of quantum networks is currently blocked by the unavailability of reliable quantum repeaters. Therefore, it remains an open question if any commercial value can be realized with these NISQ devices [22].

Delivering the full potential of quantum information processing requires the construction of scalable faulttolerant quantum (SFTQ) computing and networking technologies. Although the global race to deliver SFTQ technology is already considerable and accelerating, and the performance of some small-scale quantum systems are approaching the levels needed to operate individual logical qubits [23], scaling such systems seems to be a considerable challenge, and hence the advent of largescale SFTQ networking and computing technologies is believed by some to be a decade or more away [24] based upon the quantum architectures known today.

In this Perspective the argument is made that quantum networks and quantum information processors will both achieve ultimate scale by combining them into the same core entanglement distribution technology. The networking of quantum computers will allow for true horizontal scaling of quantum resources through modularity; and the introduction of quantum repeaters (small quantum computers) will scale telecom quantum networks' topologies and total distance. By way of illustration, a quantum architecture emphasizing entanglement distribution is proposed (Figure 1), which could be reasonably ex-



FIG. 1. Photonic's scalable quantum technology architecture. A quantum chip is cooled in a 1 K cryostat. This chip hosts integrated silicon T centres within optical cavities, photonic switches, and single photon detectors. Optical input-output (IO) ports via telecom fibre connects to a room temperature photonic switch network and control electronics. This naturally allows a highly-connected architecture with non-local connectivity even as the system scales in size. Telecom fibre also enables horizontal system scaling by connecting multiple cryostats together through their optical IO. This enables both expansion of computing power and long-distance quantum networks.

pected to deliver SFTQ technologies and the corresponding transformational real-world use cases far earlier than current predictions.

This technology is being built based upon networked spins in silicon, specifically the T centre spin-photon interface [25]. The architecture itself is broadly applicable to all long-lived spin-photon interfaces such as solid state colour centres [26] as well as some optically active atoms in vacuum [27, 28]. This architecture exploits the connectivity and modularity of quantum networks to scale the power of fault-tolerant quantum processors, and the error-corrected memory of distributed processors to scale the reach of quantum networks. This document is an overview of that vision.

II. CONNECTIVITY

The stringent thresholds for fault-tolerant quantum error correction imply that qubits must operate in strictly controlled environments, including (for example) low temperatures, ultra-high vacuum, electromagnetic shielding, with high-purity materials, and more. Each of these physical constraints result in a certain system size beyond which the marginal difficulty of adding the next qubit gets harder, not easier. Constraints such as these imply that, for each quantum computing platform, there exists a natural system size past which it would be far easier to link multiple computer modules instead of building ever-larger monolithic quantum supercomputers. This shift into modular quantum processing has multiple advantages: it can directly unlock quantum repeaters and scalable quantum networks if telecom photons are used to link the modules. For some quantum computing architectures the maximum number of future qubits in any one module could in principle be quite large, and the requirement to shift into horizontally-scalable modular quantum computing may be positioned in the distant future. However, for competitive (rapid) scaling of quantum resources, it is suggested in this Perspective that modular quantum processing will be an ultimate long-term goal for all quantum architectures on a practical basis.

If one assumes modularity will unleash limitless horizontal scaling of fault-tolerant quantum networks and quantum supercomputers, as it has for classical networking and supercomputers, the role of connectivity between modules deserves specific attention. For reference, modern classical high-performance computing hinges entirely upon parallelization across interconnected computing modules [29, 30]. In the quantum case even higher degrees of inter-module connectivity will be critical. Even with quantum algorithms that minimize the number of logical operations across modules [31–33], entanglement will need to be distributed efficiently. This is deeply connected to how quantum error correction is implemented in physical systems as is described next.

The scaling of connections between modules is a critical design parameter, as module size sets a lower bound on the number of entangled pairs that must be shared to apply an arbitrary operation on a joint two module system [34]. More importantly, connectivity remains significant in the specific case where each module encodes its own set of logical qubits. To implement logical twoqubit gates between separate modules, many physical qubits must be entangled between the individual modules. By way of example, for a CSS code, a transversal logical CNOT gate between logical qubits from two code blocks in separate modules requires transversal (pairwise) physical CNOT gates between every single one of their constituent physical qubits. Without a high degree of inter-module connectivity, the entanglement distribution between the two modules becomes a bottleneck in the performance of the distributed system [35, 36]. As systems scale, this bottleneck can be avoided only if the connectivity of qubits between modules scales with the error-detecting capacity of the code. In codes where this error-detection capacity scales proportionally with module size, this poses a stringent mandate: an interconnect for each physical qubit in the ideal case [37]. In largescale quantum systems the interconnectivity of modules may in fact dominate the temporal and spatial quantum resource estimates of quantum networking and computing applications—and yet connectivity between and within modules is almost never accounted for, and typically not even mentioned, in current quantum resource estimation literature. Simply put, these one-to-one physical entangling operations between the modules must be parallelizable for good overall system performance. This observation implies certain physical qubit capabilities in the ideal case; it is argued below that maximal parallel entanglement distribution in a modular architecture preferably implies direct telecom optical access to each physical qubit.

With the assumption of modular quantum technologies and high connectivity between modules for the purpose of high-performance parallel physical operations across modules, the natural question becomes: how should all of the qubits be connected to each other? The key to fast logical gates across modules is efficient nonlocal entanglement distribution. If indeed each qubit is designed to be easily entangled with other physical qubits from other modules, presumably that qubit supports a physical process which allows for nonlocal, as opposed to proximity-based, entanglement generation. Under such assumptions, connectivity both within and across modules are not necessarily constrained to proximity-based qubit topologies and can be connected in whatever way offers the best total system performance.

Most of today's quantum architectures are built on the paradigm of planar proximity-based entanglement generation, and aim to leverage a variant of the well-studied low-connectivity quantum error correction (QEC) code known as the surface code [38, 39] to achieve fault tolerance. Even these near-optimal codes for planar architectures require tremendous development and resources. For example, under optimistic settings for the factoring algorithm in Ref. [40], the surface code would require upwards of 3000 physical qubits for every fault-tolerant logical qubit. Moreover, in many planar architectures, the



FIG. 2. Temporal overheads associated with inter-module connectivity. An illustrative example of two 7-qubit CSS codeblocks with a) limited and b) ideal inter-module connectivity. c) Maximal inter-module connectivity where each physical qubit can be directly entangled with its partner qubit from the second module in parallel results in a single-timestep transversal CNOT gate, which implements the logical CNOT gate between these two modules. In d), the limited intermodule connectivity of this example results in a serial slow-down for the two-qubit logical CNOT operation with an illustrative, but sub-optimal, circuit implementation. Substantial additional temporal overheads are incurred (not shown here) if all-to-all connectivity within the module is not available as is assumed for this example.

vast majority of logical instruction cycles require swapping qubits in order to implement proximity-based multiqubit operations [41].

Rejecting the assumption of planar connectivity (Figure 3) for high-connectivity SFTQ architectures can give rise to a significant reduction of time and physical resource requirements [42]. Avoiding these and similar overheads by allowing for higher connectivity can significantly reduce the number of operations required [9, 43]. Even the distillation of magic states, which is largely the preferred method for providing the missing element of universality needed for surface code architectures, would be $2 \times$ more resource-efficient using high-connectivity distillation algorithms [44], and near-term codes could have $10 \times$ lower overhead [45]. Unfortunately, neither a moderate amount of long-range connectivity [46], nor small modules of qubits with substantial internal connectivity, offer a significant improvement over simple planar connectivity. The connectivity must keep pace with the system size to realize substantial improvements [47].



FIG. 3. Comparison between two-qubit operations on devices with (a) planar, nearest neighbour connectivity graphs and (b) non-local connectivity. A CNOT gate between two distant qubits on the planar graph requires sequential swap operations. Each consecutive operation accumulates errors and the overall fidelity drops precipitously, even over the small distances shown. With non-local connectivity, two qubit operations are equivalent across the graph and the device can scale with greater fidelity.

Quantum Low Density Parity Check (QLDPC) codes, quantum analogues of the classical LDPC codes underpinning 5G telecommunications networks, are a class of codes characterized by low-weight stabilizer checks. When these codes are unshackled from the constraint of local connectivity, they can achieve truly astounding performance. A recent flurry of work [48–53] culminated in the construction of QLDPC code families with asymptotically optimal code parameters, closing this long-open question after decades of relatively stagnant progress. In conjunction with previous work [54], these code families established that physical overheads could be driven significantly lower than those observed for planar code architectures. The inherent properties of QLDPC codes unlock additional performance perks including single-shot operation [55–58], which removes the requirement to repeat QEC cycles to separate measurement errors from qubit errors [59, 60]. The work to implement logical operations has only just begun, but early signs point to methods on par to those of planar codes [61, 62]. Surprisingly, no-go theorems implying the need for costly methods to achieve universality may not even apply to certain QLDPC codes [63].

Even more encouraging than these theoretical results on high-connectivity QLDPC codes has been the assessment of their practical performance, which appears to match or outperform the planar code in essentially every way. They offer similar error thresholds [46, 64], similar check weights [45], and proven efficient decoding algorithms for large finite codes [65]. QLDPC codes that encode 1 logical qubit per every 3 physical qubits have been demonstrated to saturate the quantum hashing bound under depolarizing noise (one of the biggest points of pride of the planar surface code) using decoders already in practical use today [66]. Even with rather small system sizes (where the asymptotic behavior of QLDPC codes need not be representative of actual performance) and under nuanced noise models, QLDPC codes have been demonstrated to lower resource costs by over a factor of 10 in comparison to planar codes [67] even when extending to physical resource estimation [68]. Certain families of QLDPC codes offer constant overheads as the code scales in size. For instance, explicit codes of around 1,000 physical qubits support well over 100 logical qubits [68]. With constant overhead QLDPC codes, the encoding rate could remain the same as the system scales horizontally. Unfortunately for purely planar architectures, there is no substitute for high connectivity – it is a fundamental requirement for any system hoping to capitalize on this tantalizing potential [69, 70]. While full spacetime logical circuit overhead estimates are currently scarce, existing methods [45, 71] already upper bound this overhead at O(n), and there is optimism that QLDPC logical circuit compiling can match or beat state-of-the-art methods for surface codes [72] in practice.

Even in advance of SFTQ, connectivity is also proving critical in NISQ applications. For example, to build a quantum repeater, connectivity allows for substantially more efficient entanglement distribution across many users [73, 74] and higher entanglement throughput via multi-path routing [75]. Other examples where connectivity improves near-term applications include mapping circuits to hardware [76], and quantum image processing [77].

Taken together, the rate and quality of nonlocal entanglement generation should be seen as the key characteristic dictating the ultimate performance and resource requirements of applications using large-scale modular fault-tolerant quantum systems. This is true for both large-scale networks leveraging quantum repeaters as well as modular and horizontally scalable quantum supercomputers. In this vision, quantum supercomputers and quantum networks are both ultimately large-scale entanglement distribution systems. Although many physical qubit types can in principle support nonlocal entanglement generation, and the architecture proposed here could be applied to each, practical details which impact the rate and quality of a qubit's nonlocal entanglement generation capabilities can have truly dramatic consequences on the final system performance, and any given physical qubit type should be selected for further development with this long-term requirement in mind.

III. BUILDING BLOCKS: SILICON, TELECOM, MEMORY, ENTANGLEMENT

A. Silicon

Silicon is a pinnacle material for both quantum and classical applications. It is the industrial standard for high-performance integrated electronics, as well as for low-loss, high-density, photonic integrated circuits both active and passive. Spin qubits within silicon have also proven to be exceptional quantum memories—they have set performance records for fidelity [78] and lifetimes [79]. The industrial dominance and extensive development of silicon offers such incomparable competitive advantages that, historically, if a solution is found using silicon, the silicon solution usually wins.

B. Telecom photons

Telecommunications-band (telecom) photonic flying qubits will be the backbone of any highly connected global quantum network and the backbone of modular quantum computers. Telecom photons can be flexibly routed with arbitrary connectivity to connect matter qubits both locally and remotely, with low loss in cryogenic-compatible waveguides and at room temperature using modern telecommunications infrastructure. Although a number of research efforts are now advancing the transduction of other qubits into telecom photons [80–84], the overhead of transduction processes can be avoided entirely by working with quantum systems that interact with telecom photons directly—such as the silicon T centre. In this Perspective, the claim is made that telecom photons are essential for high-connectivity quantum technologies at scale but are likely not practically sufficient on their own without a quantum memory.

C. Quantum memory

The central challenge with photonic qubits of all wavelengths, across quantum computing and quantum networking, is that they suffer from unavoidable loss. Each photonic process over optical links succeeds only a fraction of the time, and this is true even for telecom photons which have the lowest loss photonic components available by far (e.g. switches) [85]. In large-scale, highconnectivity settings, each switch layer introduces yet further loss. Quantum memories offer a straightforward path to protect against photon loss for both high-fidelity computing and networking applications.

Quantum memories are necessary for the construction of quantum repeaters with the highest functionalities [86]. In repeaters, ultra-long quantum memory lifetimes are essential for high-fidelity qubit storage and the coordination of photon signals with long time-of-flight (long distance) connections, as well as in high-loss environments where many entanglement-generation attempts are necessary on average before photons successfully arrive. Solid-state colour centre qubits with access to a spin degree of freedom [26, 87], as well as some optically active atoms in vacuum [27, 28, 88], offer all of the aforementioned ideal quantum functionality. Each of these physical qubit types offer a direct high quality photon interface into at least one long-lived [79], high-fidelity [89], universally-controllable spin qubit [90]. Every physical qubit in this category has a direct and dedicated photonic interface allowing for parallel entanglement generation.

For high-connectivity quantum computing, a similar architectural advantage is proposed here where the photons distribute entanglement but do not process quantum information directly: the processing is done within the spin qubits. This design is inherently tolerant to photon loss as probabilistic entanglement-generation attempts can be repeated until success, without losing the spin qubit state [91]. Essentially, the high connectivity between matter qubits – which could be physically arranged on planar chips – is provided by entanglement carried by telecom photons and the arbitrary connectivity that these photons provide.

Maximally entangled Bell pairs (BPs - pairs of qubits in one of the four Bell states) of long-lived spins can be produced via a variety of photonic methods (see Figure 6 below for a specific implementation [92]). Once entanglement in the form of BPs is delivered to the spins, this entanglement can be used as a resource to construct cluster states to be consumed for computation within the measurement-based computing paradigm [93, 94] or teleported gates in the traditional gate model of quantum computing [95, 96]. Both of the cluster state, and teleported gate models of quantum computation allow for blind-computing applications over a network [15, 97]. Below the focus is on teleported gates as an example implementation.

For this kind of quantum technology, quantum networks' repeaters and quantum computing modules will be almost identical in their core construction, which reflects the fact that the key to efficient, high-performance large-scale quantum networking and quantum computing is high-bandwidth high-connectivity entanglement distribution. In the long term, it is suggested that this unified technology will outperform architectures with disparate networking and computing cores that require extra layers of quantum interfaces which will inevitably compromise entanglement distribution rates and fidelities.

D. Computing and networking with T centres

Even in a loss-tolerant design, achieving efficient distribution of entanglement at scale is a critical performance metric for high-connectivity SFTQ. This further motivates the adoption of solid-state spin-photon interfaces, ideally telecom colour centres, which can be directly placed into low-loss integrated photonic circuits in silicon. Direct integration not only maximizes photon collection efficiency from every single qubit, but also enhances the photon emission rate and quality using routine photonic engineering, while the spins, particularly long-lived nuclear spins, can perform at the high levels necessary for low-overhead QEC.

The above considerations foretell a future where scalable fault-tolerant quantum networks and scalable fault tolerant quantum computers will be the same core technology, with long-lived spins embedded into silicon integrated photonics, directly interfacing with telecom photons. In particular, silicon T centre [25, 98, 99] is proposed as the exemplary foundational quantum unit: in addition to direct telecom access, it contains one unpaired electron spin and up to three spin-1/2 nuclear spins (one hydrogen and two carbon atoms, see Figure 1), each of which serves as a high fidelity and long-lived qubit with good performance comfortably above 1 K [25] where substantial cooling power is available for large-scale modular processors.

Solid-state spin performance is largely determined by the host crystal environment. High fidelity performance and ultra-long coherence times [79, 100] are common to many silicon spin centres including T centres [25]. In particular, the capability to isotopically purify silicon to the nuclear spin-free ²⁸Si isotope has shown nuclear spin coherence times of 3 hours [79], with fidelities above 99.9% [78]. T centres in bulk ²⁸Si samples show spin coherence times (T_2 Hahn echo) of 2.1 ms for the electron and 1.1 s for the hydrogen nuclear spin; both spins have T_1 lifetimes far longer than a second [25].

The T centre spins almost uniquely [101] possess a direct telecom photonic interface in silicon. It can efficiently interact with a pump laser pulse and emit a spin-entangled O-band photon. A photonic cavity can change the photonic environment around the T and enhance the emission of the photon into desired optical modes through the Purcell effect [102]. These modes are coupled to optical waveguides with well-defined modes, and therefore spin-entangled photons can be accurately (and with low loss) routed to their destination either through integrated photonic waveguides or by coupling into optical fibre. Optical fibres can connect T centres across multiple chips, enabling a naturally modular

and horizontally-scalable architecture. Modular scaling can relieve IO density challenges and environmental constraints that emerge as systems grow in size.

IV. QUANTUM OPERATIONS

The silicon T centre, only recently rediscovered for quantum applications in 2020 [25], offers all the ideal characteristics listed above to enable high-performance SFTQ in the near and long term. It merges the advantages of silicon fabrication and scalability, telecom emission, and long-lived spin memories into a quantum system uniquely adapted to a modular, high-connectivity architecture. Below it is described how the qubits of the T centre perform high-fidelity operations according to this vision.

A. State preparation and measurement

For readout and initialisation, the electron spin can be projectively measured and initialized to high fidelity using spin-dependent optical excitation (Figure 4a) [25]. The initialization and readout of the full T centre spin



FIG. 4. (a) Energy levels of a T centre. Transition C is optically pumped to generate entanglement and for readout. h - excited state hole spin; e - ground state electron spin. (b) Optical spectrum of the T centre's electron spin, overlaid with optical cavity to enhance emission of the C transition.

register can be realized via successive rounds of high fidelity state swaps between the nuclear spins and the electron spin. Additionally, quantum non-demolition (QND) readout of the nuclear spins can be achieved by executing two-qubit CNOT gates with the electron spin as the target, optically reading out the electron spin, and then repeating as necessary [103].

The four spin qubits in a T centre's local register admit independent and high-fidelity single-qubit and multiqubit gates through standard electron and nuclear magnetic resonance techniques [25]. Additionally, the low spectral diffusion of the spin transitions in ²⁸Si allows for full access to state-selective transitions, and high fidelity CNOT gates between the electron and its coupled nuclear spins are directly available [25].

B. Bell pair entanglement

Two-qubit gates rely on proximity-based or nonlocal entangling operations between qubits. Remote T centres interact through photon-mediated entanglement. Each T centre can be optically triggered to emit single telecom O-band photons [99], optionally entangled with the electron spin. To generate entanglement between spins, these photons must be indistinguishable. This means matching every degree of freedom between two photons: wavelength, linewidth, timing, polarization, and so on. Importantly, the wavelength of these photons can be tuned with either strain [104] or electric fields [105, 106] to compensate for the variations in T centres' local environments. Other degrees of freedom can be engineered to match using integrated photonics, and timing can be synchronized by calibrating control signals against optical path lengths.

The indistinguishability of spin-entangled photons dictates the quality of the generated Bell pairs, and hence the maximum fidelity of the multi-qubit operations. Highly indistinguishable photons can be produced on demand when the emitter's optical transition linewidth approaches its lifetime-limited value, rather than being broadened by environmental noise. This can be dramatically assisted by photonic engineering, specifically by spin-selective Purcell enhancement of the optical transition when on resonance (Figure 4b) with a high-Q photonic cavity [107]. This offers faster emission rates, leading to larger lifetime contributions to the linewidth, and therefore more indistinguishable photons, as well as higher cyclicity (defined as the conservation of the spin state through an optical cycle).

Integration into photonic cavities has already led to demonstrations of $> 20 \times$ reductions in single T centre excited state lifetimes compared to the 940 ns unenhanced value [108], enhancing indistinguishability and increasing maximum photon rates. Instantaneous linewidths for single T centres only $5 \times$ larger than the lifetime-limited value have also been measured [108]. Additional improvements towards high indistinguishability include enhanced



FIG. 5. (a) Two T centres emit synchronized photons. Detection times will vary due to the finite time width of the photon wavepackets. (b) The detection time difference dt can be used to herald only high-fidelity Bell pairs. As the heralding threshold grows more stringent, the Bell pair generation rate decreases. The model curves above were calculated using reasonable projections of system parameters from current performance.

fabrication methods to reduce the optical spectral diffusion.

The triggered emission of two indistinguishable highpurity spin-entangled photons can be used to project the participating electron spins into a maximally entangled Bell Pair spin state [109, 110]. An illustrative example of the BP generation protocols that exist, the Barrett-Kok [92] protocol family is notable in that it is whichway symmetric, does not require interferometric stability of the setup, and heralds success with minimal false positives.

The Barrett-Kok protocol proceeds as follows (Figure 6): the electron spins of two separated T centres are prepared in a superposition of spin-up and spin-down, and triggered to emit a photon resonant with the spinup transition. The photonic modes are interfered on a beamsplitter and detected. Then the spin states are inverted, and the emission and detection is repeated. Using exactly one photon detection in each of the first and second optical cycles as a herald signal, the spins are projected into a maximally-entangled BP.

Although photon loss can cause an individual attempt to fail, success is heralded and the entire process can be repeated until success, with small degradation in the local nuclear spin coherence for each attempt (see below). Moreover, for this protocol, higher fidelity BPs can be



FIG. 6. Two-qubit gate mechanism. The electron spin is initialized in a superposition state and optically pumped to emit a spin-entangled photon. Repeating this process and heralding on photon detector clicks produces photon-mediated electron spin-spin entanglement, which is consumed in remote T centres to effect a CNOT gate between two nuclear spins.

obtained at the expense of lower entanglement rates by shifting the heralding threshold, the time difference between detector clicks, e.g. using time-bin filtering for time-frequency mismatch (Figure 5) [111].

C. Teleported gates

The teleported-gate model has a number of advantages at scale. Teleportation delivers two-qubit gates on demand, after coincident, identical photons are heralded. This inherently eliminates accidental multi-qubit operations and cross-talk that can occur when two-qubit interactions are "always-on". Furthermore, qubits coupled to common local environments are susceptible to correlated errors that can dramatically degrade QEC performance [112–114]. Since the teleported gate is a remote operation, each T centre can occupy a unique local environment, thus minimizing correlations in noise.

Once the BP entanglement is heralded and delivered to the two T centres' electron spins, that BP can be consumed in a teleported gate sequence to apply a nonlocal multi-qubit gate between the two T centres' nuclear spins using only local operations and classical feedforward operations (see Figure 6) [96, 115, 116]. Alternatively, a series of successful BPs can be distilled through local measurements and gates onto a remote memory qubit BP [117, 118], with fidelity beyond the photonindistinguishability, before this distilled pair is consumed to deliver higher fidelity teleported gates. A register of multiple spin memory qubits such as the T centre's internal nuclear spins allows a tiered distillation protocol [119]. Distillation can prioritize the dominant error pathway of the teleported gate, yielding, for example, teleported CNOT process fidelity exceeding the heralded or distilled BP fidelity [110].

The approach outlined here uses the electron as a spin-photon interface to generate and distribute entanglement, and nuclear spins for processing and memory. A configuration must be chosen that protects the memory qubits during entanglement distribution attempts. Repeated attempts can, in general, degrade the information stored in memory by perturbing the nuclear spin states [87].

Protection techniques include working in a decoherence free subspace [120], aligning an external magnetic field along one of the principal axes of the hyperfine tensor [121], and reducing the optical excited state lifetime [122, 123]. With sufficient protection, BP generation does not materially affect data stored in the nuclear spins and so the teleported gate operations are made photon-loss tolerant. Moreover, because of the long memory, this system works well for distillation and creating large entangled states, even after minimizing the required number of attempts by minimizing optical losses.

V. SCALABLE QUANTUM TECHNOLOGY

Scaling the number of interacting qubits in this architecture requires mediating entanglement across a growing network of qubits. Enabling multi-qubit operations between any two T centres amounts to routing their indistinguishable photons to the same beamsplitter and detector module. Here, the silicon-on-insulator (SOI) platform can be leveraged to drive on-chip device integration to a stage where photons never need to leave the photonic integrated circuit, maximizing photon collection efficiency and entanglement heralding rates. Furthermore, optical interconnects off-chip allow connections to outside networks and modular scaling of computing power.

A visual summary of this physical architecture is shown in Figure 1. T centres are placed into photonic cavities which are directly coupled into low-loss photonic waveguides. Emitted photons are routed on-chip through switches and, optionally, through high-efficiency optical IO ports into one of the many optical fibres. An optical switch fabric governed by run-time electronics guides T centre emission to single-photon detectors and Bell-state measurement modules within the cryostat.

The ability to then link distinct photonic chips through off-chip interconnects enables modularizing and distributing IO across multiple chips to horizontally scale the system size both within one and amongst many cryostats. Thanks to the low loss of telecom photons in optical fibre, the cryostats can be metres to kilometres apart.

A. On-chip scalability

T centres are atomic in scale and can be localized within small photonic components without direct optical crosstalk, making it possible to pattern millions of individually addressable qubits per chip [99]; however, the ultimate qubit count per chip will be governed by the larger footprint of active photonic components, control electronics, and signal IO.

The low thermal conductivity of optical fibre enables incredibly high count optical IO into cryogenic environments (e.g. 37,000 optical fibres in Ref. [124]). Highdensity optical connections to integrated photonic devices have also been demonstrated in cryogenic environments [125]. Furthermore, the T centre achieves operating performance (electron spin lifetime, optical linewidth) at only 1-2 K [25]. This relatively high operation temperature circumvents the need for dilution refrigeration technologies, and permits increased cooling power to counteract the heat load generated by the signal IO. As high density photonic IO continues to develop, it is expected that each chip could be able to support upwards of 4,000 fibres with 1,000+ qubits per system and beyond.

T centre qubits can be controlled using electrical signals segmented into global fields and local fields. Global fields include a static magnetic field B_0 to define the T centre spin Hamiltonian and a spin control field B_1 that can address qubits individually using global microwave control techniques [126–128]. Local control fields include electric field E, a trimming magnetic field δB_0 , and optical switch and detector signals. Multiplexing strategies reduce the electrical line count for all signals, particularly the local control [129] and detector [130] signals.

B. Local quantum networks

As mentioned in Section III, quantum computing and quantum networking are essentially the same in the architecture presented here. The challenge that is overcome through this approach to SFTQ technologies – highbandwidth, high-connectivity entanglement distribution over lossy optical networks – is the same challenge faced in scaling up quantum communication networks.

For this high-connectivity architecture, horizontal scaling of quantum resources is achieved by racking multiple units adjacent to each other into a modular local quantum network of qubits (Figure 1). Where additional network topology, inter-chip links, or connections between distant T centres are needed, optical IO can route photons to switch networks. Because optical photons can noiselessly traverse from cryogenics to room temperature without transduction, the switching network can be placed either within the cryostat or at room temperature.

This design allows for arbitrary and flexible connectivity using switchable telecommunications hardware at room temperature with simple connections to cryogen9



FIG. 7. The non-local connectivity of the system works the same way when entangling qubits across the chips in the same cryostat, 10 m away, or 100 km away.

ics via optical fibre. Additionally, on-chip switching and multiplexing can then be used to scale the qubit count within each chip beyond limitations imposed by the IO fibre count.

For the architecture described in this paper, T centres from distinct cryogenic units can be entangled in the same way as adjacent T centres (Figure 7). Put another way, the multi-qubit operations between different cryostats can be similar in performance to the multiqubit operations within a chip. This design maximizes the performance of important distributed logical quantum algorithms [131]. The physical and logical qubits will have effectively all-to-all connectivity either by routing photons to common beamsplitters, by teleporting qubit states to new neighbourhoods of connectivity, or through entanglement swapping protocols [16].

C. Quantum repeaters

Broad adoption of a quantum internet [132] is hampered by issues that arise when quantum communication channels are extended beyond local (≈ 10 km) distances: the communication rate of point-to-point terrestrial fibre links drops exponentially with distance due to signal attenuation intrinsic in even the best available optical fibres. Similarly, beam divergence, atmospheric absorption, and turbulence impose fundamental losses and noise on free-space telescopic links. And even though telecom networks have achieved point-to-point quantum communications over distances as large as 1000 km [133], optical quantum repeaters remain the key to unlocking global quantum networking applications [86].

Fibre-coupled, T centre processor nodes contain telecom optical quantum memories that are capable of supporting long-distance communication over low-loss fibre links, and are able to generate, store and process optically-entangled spin qubits. Hence, they can be naturally operated as a telecom-wavelength quantum repeater to realize a long-distance, scalable multi-user quantum internet, with the potential to securely distribute entanglement or connect quantum processors between any two users on the global network (Figure 1).

As an initial implementation, the nested quantum repeater protocol [134] creates long distance entanglement by heralded entanglement generation over shorter links and entanglement swapping between adjacent nodes followed by entanglement distillation. The steps are repeated between successive, now twice as long, entangled links until the end nodes are entangled. The individual entanglement generation steps taken during this procedure are exactly the same as local, distributed SFTO computing using the proposed architecture, allowing implementation by T centre processors. This repeater protocol belongs to the so called first generation of quantum repeaters [135]. Here the loss errors are addressed by heralded entanglement generation and operational errors are addressed by entanglement distillation, which is a specific type of error correction.

Looking to the second generation of quantum repeaters, one can leverage the ability of T centre processors to perform multi-qubit encoding and logical operations between nodes. As before, loss errors on individual photons are addressed by heralded entanglement generation, but the operational errors are addressed by logical encoding at each node, namely with codes permitting transversal two-qubit operations [136], i.e. the teleported CNOT implemented in Fig. 6. Each node possesses one or more logical qubits, which become logically entangled to neighbouring logical qubits through entanglement distribution. Subsequent logical entanglement swapping allows the end-node logical qubits to be entangled, similar to the end-node physical qubits of the first generation of repeaters. This second generation uses the scalability and connectivity of T centres in individual nodes to distribute error-corrected logical qubits over long distance.

For these reasons, the proposed T centre based architecture is not only best suited to tackle the challenges of SFTQ computing, but also offers a very attractive platform for implementing large scale quantum networks and their various applications.

VI. QUANTUM NETWORKING APPLICATIONS

A truly scalable network, extendable both in distance and number of users, unlocks powerful applications [137, 138], such as quantum key distribution (QKD), distributed quantum computing, blind quantum computing [15], and enhanced sensing. Here, the focus is on two applications, blind quantum computing and quantum key distribution, and discuss how these could be naturally implemented in the proposed architecture.



FIG. 8. (a) A protocol for loading of a time bin photonic qubit state into the spin memory of a T centre. (b) A quantum circuit that utilizes the protocol in (a) to generate a cryoptographic key for two users (laser clients, LCs) connected to the same hub (c). Loading photons from two users and then performing a Bell-state measurement on them computes the parity of their states, allowing secret key generation in MA-MDI QKD.

A. Memory-assisted, measurement-device independent quantum key distribution

Quantum key distribution is the quantum networking application nearest to widespread adoption. A first implementation of a QKD network using T centres consists of users possessing laser clients (LCs) that are optically connected to a quantum processor containing T centres (see Figure 8), possibly in a hub-spoke configuration. Users can encode quantum information onto weak coherent pulses (for example, using time-bin encoding) and send this to the processor hub. There, using the T centre's spin-photon entangled state and quantum teleportation, the encoded quantum state is loaded to a local spin state [139] for storage and further processing. Multiple hubs could be interconnected to extend range and capacity (Figure 9).

The proposed networking configuration realizes memory-assisted measurement-device independent (MA-MDI) QKD [140]. Within this protocol a secret key between two users is established by first loading each user's states onto the hub's qubits followed by a Bell-state measurement between these two qubits. The result of that measurement is announced to the users, who use this information to extract a shared secret key.

The MA-MDI QKD protocol eliminates detector-side channel attacks [141], allowing the hub to operate as an untrusted node. Entanglement-based linking of nodes can also be untrusted, providing significantly more security than existing trusted-node networks [142]. At



FIG. 9. (a) Quantum circuit for loading qubit states from distant laser clients to two different quantum hubs and performing a nonlocal Bell state measurement between the two T centre spins. The teleported CNOT (middle) is mediated by entanglement distribution between the two hubs. (b) Because the CNOT between the two T centres is teleported, this circuit could be used to generate cryptographic keys between two users connected to different cryostats that are linked by telecom fibre.

the same time, the memory-assistance provides improved keyrate scaling [140] due to the ability to store the clients' quantum information in the T centre spins. This means the client photons do not have to arrive simultaneously; their loading to the spins is heralded, and only once both are loaded is the Bell-state measurement performed between spins.

Within this architecture, each hub will support thousands of users in a hub-spoke model, well exceeding the usual point-to-point connections of currently commercially available QKD systems. Performing entanglement distribution between hubs as in Fig. 9 allows implementing the same MA-MDI-QKD protocol in a distributed setting. The shared entanglement can be generated independently of the qubit loading events, and then consumed to implement the nonlocal Bell-state measurement.

Whether connected to a single hub or a repeater chain, each user requires only a simple cost-effective source: a room-temperature laser attenuated below the singlephoton level. It is equipped with a modulator able to produce time-bin encoded photonic qubits, which are spectrally matched to T centres in the hub. Such a device is very similar to a modern datacenter's telecom transceivers, and the photonic qubits produced by these sources can be similarly routed through existing datacentre or metropolitan telecom fibre to the hub.

B. Blind quantum computing

Although quantum repeaters allow for long-distance quantum computation and telecom-wavelength photons allow for modular quantum design, quantum computers may remain constrained by complex hardware required to protect and manipulate quantum states. This restriction suggests that quantum supercomputers are expected to start, and perhaps remain, as network-accessible largescale devices. Cloud access over the classical internet may

not be able to provide the necessary privacy for all potential end-user applications. Additionally, the operators of the quantum computers would have access to both the data and the algorithm run on it. Similar concerns for classical cloud computing galvanized research in the field of homomorphic encryption [143], where a computer performs logic directly on encrypted data without ever gaining the capability to decrypt and learn the underlying information. Blind quantum computing [15], which may be viewed as a quantum analogue to homomorphic encryption, allows users to perform arbitrary computations using remote quantum processing resources, while restricting the computer from having meaningful access to both the data and the algorithm. In this highly-connected architecture, the same process of qubit loading employed for MA-MDI QKD can also be employed to realize remote and blind quantum computation, without a change of user hardware. This is due to the intrinsic spin-photon interface of the T centre, which allows loading arbitrary external user qubits into the computation in a heralded way. Loaded qubits are largely used to direct the computation remotely, while some are reserved as checks to confirm the operation of the computer [144].

VII. CONCLUSION

The highest-value quantum applications known today require fault-tolerant capabilities at scale. Modular quantum systems, potentially distributed over global distances, will be the ultimate version of horizontallyscaled quantum information processing and networking. Underpinning this view is the observation that combining quantum computing and quantum networking technology removes the fundamental obstacles to scale that each of these technologies are facing in isolation, ie truly scalable architectures are horizontally scalable. Namely, to unlock quantum networks at scale one needs to develop quantum repeaters which in the high-performance limit are essentially fault-tolerant quantum computing modules, and to unlock truly scalable quantum computing one needs to leverage the entanglement distribution capabilities of quantum networking to link quantum computing modules into quantum supercomputers.

Given that high-bandwidth, high-quality entanglement distribution ultimately sets the performance of both scalable (modular) fault-tolerant quantum computing and networking, quantum systems should be engineered to optimize entanglement distribution. Under this model, quantum computing and quantum networking are (in the ideal case) ultimately the same fundamental technology.

This Perspective proposes a scalable quantum (networking *and* computing) architecture with this end goal in mind. It suggests a specific implementation using telecom colour centres in silicon, namely the T centre, but this architectural model is broadly applicable to many qubit systems. Because of the high connectivity offered by the spin-photon interface, this architecture can take advantage of fixed- and low-overhead QLDPC codes to deliver fault tolerance. Thanks to the integrated silicon photonics platform, thousands of qubits can be fabricated and addressed on a single chip with optical and electronic control and routing, and modules can be connected together across existing global telecommunications infrastructure without any transduction losses. Using a T cen-

- S. Jordan, Quantum Algorithm Zoo, https: //quantumalgorithmzoo.org/.
- [2] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, Quantum computing enhanced computational catalysis, Physical Review Research 3, 033055 (2021), 2007.14460.
- [3] I. H. Kim, Y.-H. Liu, S. Pallister, W. Pol, S. Roberts, and E. Lee, Fault-tolerant resource estimate for quantum chemical simulations: Case study on Li-ion battery electrolyte molecules, Physical Review Research 4, 023019 (2022).
- [4] J. J. Goings, A. White, J. Lee, C. S. Tautermann, M. Degroote, C. Gidney, T. Shiozaki, R. Babbush, and N. C. Rubin, Reliably assessing the electronic structure of cytochrome P450 on today's classical computers and tomorrow's quantum computers, Proceedings of the National Academy of Sciences **119**, e2203533119 (2022), 2202.01244.
- [5] L. Kee, RSA Is Dead We Just Haven't Accepted It Yet, https://www.forbes. com/sites/forbestechcouncil/2021/05/06/ rsa-is-dead---we-just-haventaccepted-ityet/.
- [6] A. Raya and K. Mariyappn, Diffie-Hellman Instantiations in Pre- and Post- Quantum World: A Review Paper, in 2020 Fifth International Conference on Research in Computational Intelligence and Communication Networks (ICRCICN) (2020) pp. 130–136.
- [7] P. Campbell, M. G. Shepherd, and D. Shepherd, Soliloquy: A cautionary tale, Tech. Rep.
- [8] M. Webber, V. Elfving, S. Weidt, and W. K. Hensinger, The impact of hardware specifications on reaching quantum advantage in the fault tolerant regime, AVS Quantum Science 4, 013801 (2022).
- [9] E. Gouzien, D. Ruiz, F.-M. L. Regent, J. Guillaud, and N. Sangouard, Computing 256-bit Elliptic Curve Logarithm in 9 Hours with 126133 Cat Qubits, arXiv:2302.06639.
- [10] W. Beullens, Breaking rainbow takes a weekend on a laptop, Cryptology ePrint Archive, Paper 2022/214 (2022), https://eprint.iacr.org/2022/214.
- [11] W. Castryck and T. Decru, An efficient key recovery attack on sidh (preliminary version), Cryptology ePrint Archive, Paper 2022/975 (2022), https: //eprint.iacr.org/2022/975.
- [12] E. Dubrova, K. Ngo, and J. Gärtner, Breaking a fifthorder masked implementation of crystals-kyber by copypaste, Cryptology ePrint Archive, Paper 2022/1713 (2022), https://eprint.iacr.org/2022/1713.
- [13] R. Bavdekar, E. Jayant Chopde, A. Agrawal, A. Bhatia, and K. Tiwari, Post quantum cryptography: A review of techniques, challenges and standardizations, in 2023 International Conference on Information Network-

tre network to distribute verified quantum entanglement allows for device-independent networking protocols, providing the ultimate protection against eavesdropper attacks, as well as other high-value applications leveraging entanglement distribution such as blind computing. Taken together, the future for truly scalable quantum technology is bright.

ing (ICOIN) (2023) pp. 146–151.

- [14] C. H. Bennett and G. Brassard, Quantum cryptography: Public key distribution and coin tossing, in *Proceedings* of the International Conference on Computers, Systems & Signal Processing (1984) pp. 175–179.
- [15] A. Broadbent, J. Fitzsimons, and E. Kashefi, Universal Blind Quantum Computation, in 2009 50th Annual IEEE Symposium on Foundations of Computer Science (2009) pp. 517–526, 0807.4154.
- [16] M. Caleffi, M. Amoretti, D. Ferrari, D. Cuomo, J. Illiano, A. Manzalini, and A. S. Cacciapuoti, Distributed Quantum Computing: a Survey, arXiv:2212.10609.
- [17] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, A quantum network of clocks, Nature Physics 10, 582 (2014), 1310.6045.
- [18] Z. Zhang and Q. Zhuang, Distributed quantum sensing, Quantum Science and Technology 6, 043001 (2021).
- [19] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'brien, A variational eigenvalue solver on a photonic quantum processor, Nature communications 5, 4213 (2014).
- [20] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, arXiv:1411.4028 ().
- [21] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, arXiv:quant-ph/0001106 ().
- [22] J. W. Z. Lau, K. H. Lim, H. Shrotriya, and L. C. Kwek, NISQ computing: where are we and where do we go?, AAPPS Bulletin **32**, 27 (2022).
- [23] Y. Wang, S. Simsek, T. M. Gatterman, J. A. Gerber, K. Gilmore, D. Gresh, N. Hewitt, C. V. Horst, M. Matheny, T. Mengle, B. Neyenhuis, and B. Criger, Faulttolerant one-bit addition with the smallest interesting colour code, arXiv:2309.09893.
- [24] M. Mosca and M. Piani, 2022 Quantum Threat Timeline Report, Tech. Rep.
- [25] L. Bergeron, C. Chartrand, A. T. K. Kurkjian, K. J. Morse, H. Riemann, N. V. Abrosimov, P. Becker, H. J. Pohl, M. L. W. Thewalt, and S. Simmons, Silicon-Integrated Telecommunications Photon-Spin Interface, PRX Quantum 1, 020301 (2020).
- [26] D. D., Awschalom, R. Hanson, J. Wrachtrup, and B. B. Zhou, Quantum technologies with optically interfaced solid-state spins, Nature Photonics 12, 516 (2018).
- [27] B. B. Blinov, D. L. Moehring, L.-M. Duan, and C. Monroe, Observation of entanglement between a single trapped atom and a single photon, Nature 428, 153 (2004).
- [28] T. Wilk, S. C. Webster, A. Kuhn, and G. Rempe, Single-Atom Single-Photon Quantum Interface, Science 317, 488 (2007).
- [29] J. L. Gustafson, Reevaluating amdahl's law, Commun.

ACM 31, 532–533 (1988).

- [30] K. Asanović, R. Bodik, B. C. Catanzaro, J. J. Gebis, P. Husbands, K. Keutzer, D. A. Patterson, W. L. Plishker, J. Shalf, S. W. Williams, and K. A. Yelick, *The Landscape of Parallel Computing Research: A View* from Berkeley, Tech. Rep. UCB/EECS-2006-183 (EECS Department, University of California, Berkeley, 2006).
- [31] R. Van Meter and S. J. Devitt, The path to scalable distributed quantum computing, Computer 49, 31 (2016).
- [32] D. Ferrari, A. S. Cacciapuoti, M. Amoretti, and M. Caleffi, Compiler design for distributed quantum computing, IEEE Transactions on Quantum Engineering 2, 1 (2021).
- [33] D. Cuomo, M. Caleffi, K. Krsulich, F. Tramonto, G. Agliardi, E. Prati, and A. S. Cacciapuoti, Optimized compiler for distributed quantum computing, ACM Transactions on Quantum Computing 4, 1 (2023).
- [34] D. Stahlke and R. B. Griffiths, Entanglement requirements for implementing bipartite unitary operations, Physical Review A 84, 032316 (2011).
- [35] R. Nigmatullin, C. J. Ballance, N. d. Beaudrap, and S. C. Benjamin, Minimally complex ion traps as modules for quantum communication and computing, New J. Phys. 18, 103028 (2016).
- [36] Y. Li and S. C. Benjamin, Hierarchical surface code for network quantum computing with modules of arbitrary size, Phys. Rev. A 94, 042303 (2016).
- [37] As a prototypical example, consider the case where each module accommodates a codeblock of a large n-qubit error correction code with distance d. Here it is argued that a fault-tolerant entangling logical operation between these modules will always require O(d) simultaneous physical entangling operations between the two modules. Performing entangling logical operations between two codeblocks necessarily involves at least O(d)physical entangling operations, as the support of a logical Pauli in each module must grow by O(d) in the other module, though these could be mostly within a codeblock in principle. If this operation is performed fault-tolerantly, any one single-qubit Pauli error cannot grow a support larger than O(1) in either module. Since a logical Pauli operation in one module consists of the tensor product of O(d) single qubit Pauli operations, O(d) single qubit Pauli operations in one module must each gain distinct support from one another in the other module. This implies that there needs to be O(d) entangling operations *across* the two modules. Since we would like the total circuit depth for the logical operation to be as close to O(1) as possible to maintain fault-tolerance, we should be capable of performing these O(d) moduleto-module entangling operations simultaneously. For example, this is how lattice surgery will be accomplished in planar QEC architectures. More importantly, if we consider an optimally scaling code where d = O(n), this already implies we need to be able to apply O(n)entangling operations in parallel between two n-qubit modules. A classic example of a logical operation with this requirement is the transversal CNOT operation in CSS codes.
- [38] S. B. Bravyi and A. Y. Kitaev, Quantum codes on a lattice with boundary, arXiv:quant-ph/9811052 (1998).
- [39] A. Kitaev, Fault-tolerant quantum computation by anyons, Annals of Physics 303, 2 (2003), quantph/9707021.

- [40] C. Gidney and M. Ekerå, How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits, Quantum 5, 433 (2021), 1905.09749.
- [41] M. McEwen, D. Bacon, and C. Gidney, Relaxing Hardware Requirements for Surface Code Circuits using Time-dynamics, arXiv:2302.02192.
- [42] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, Nature 549, 172 (2017), 1612.07330.
- [43] D. Litinski and N. Nickerson, Active volume: An architecture for efficient fault-tolerant quantum computers with limited non-local connections, arXiv:2211.15465 (2022).
- [44] S. Bravyi and J. Haah, Magic-state distillation with low overhead, Physical Review A 86, 052329 (2012), 1209.2426.
- [45] L. Z. Cohen, I. H. Kim, S. D. Bartlett, and B. J. Brown, Low-overhead fault-tolerant quantum computing using long-range connectivity, Science Advances 8, eabn1717 (2022), 2110.10794.
- [46] M. A. Tremblay, N. Delfosse, and M. E. Beverland, Constant-Overhead Quantum Error Correction with Thin Planar Connectivity, Physical Review Letters 129, 050504 (2022), 2109.14609.
- [47] N. Baspin and A. Krishna, Quantifying Nonlocality: How Outperforming Local Quantum Codes Is Expensive, Physical Review Letters **129**, 050505 (2022), 2109.10982.
- [48] J.-P. Tillich and G. Zémor, Quantum ldpc codes with positive rate and minimum distance proportional to the square root of the blocklength, IEEE Transactions on Information Theory 60, 1193 (2013).
- [49] M. B. Hastings, J. Haah, and R. O'Donnell, Fiber bundle codes: breaking the n 1/2 polylog (n) barrier for quantum ldpc codes, in *Proceedings of the 53rd An*nual ACM SIGACT Symposium on Theory of Computing (2021) pp. 1276–1288.
- [50] N. P. Breuckmann and J. N. Eberhardt, Balanced product quantum codes, IEEE Transactions on Information Theory 67, 6653 (2021).
- [51] P. Panteleev and G. Kalachev, Quantum ldpc codes with almost linear minimum distance, IEEE Transactions on Information Theory 68, 213 (2021).
- [52] A. Leverrier and G. Zémor, Quantum tanner codes, in 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS) (IEEE, 2022) pp. 872–883.
- [53] P. Panteleev and G. Kalachev, Asymptotically good quantum and locally testable classical ldpc codes, in *Proceedings of the 54th Annual ACM SIGACT Sympo*sium on Theory of Computing (2022) pp. 375–388.
- [54] D. Gottesman, Fault-tolerant quantum computation with constant overhead, arXiv:1310.2984.
- [55] N. P. Breuckmann and V. Londe, Single-Shot Decoding of Linear Rate LDPC Quantum Codes with High Performance, arXiv:2001.03568.
- [56] A. O. Quintavalle, M. Vasmer, J. Roffe, and E. T. Campbell, Single-shot error correction of threedimensional homological product codes, PRX Quantum 2, 020340 (2021).
- [57] O. Higgott and N. P. Breuckmann, Improved single-shot decoding of higher-dimensional hypergraph-product codes, PRX Quantum 4, 020332 (2023).
- [58] S. Gu, E. Tang, L. Caha, S. H. Choe, Z. He, and A. Kubica, Single-shot decoding of good quantum ldpc codes,

arXiv:2306.12470.

- [59] H. Bombín, Single-shot fault-tolerant quantum error correction, Physical Review X 5, 031043 (2015).
- [60] E. T. Campbell, A theory of single-shot error correction for adversarial noise, Quantum Science and Technology 4, 025006 (2019).
- [61] A. O. Quintavalle, P. Webster, and M. Vasmer, Partitioning qubits in hypergraph product codes to implement logical gates, arXiv:2204.10812.
- [62] A. Cowtan and S. Burton, Css code surgery as a universal construction, arXiv:2301.13738.
- [63] P. Webster, M. Vasmer, T. R. Scruby, and S. D. Bartlett, Universal fault-tolerant quantum computing with stabilizer codes, Physical Review Research 4, 013092 (2022), 2012.05260.
- [64] J. Roffe, D. R. White, S. Burton, and E. Campbell, Decoding across the quantum low-density parity-check code landscape, Phys. Rev. Res. 2, 043423 (2020).
- [65] P. Panteleev and G. Kalachev, Degenerate Quantum LDPC Codes With Good Finite Length Performance, Quantum 5, 585 (2021), 1904.02703.
- [66] S. Yang and R. Calderbank, Spatially-coupled qdlpc codes, arXiv:2305.00137.
- [67] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, arXiv:2308.07915.
- [68] Q. Xu, J. Ataides, C. A. Pattison, N. Raveendran, D. Bluvstein, J. Wurtz, B. Vasic, M. D. Lukin, L. Jiang, and H. Zhou, Constant-overhead fault-tolerant quantum computation with reconfigurable atom arrays, arXiv:2308.08648.
- [69] N. Baspin and A. Krishna, Connectivity constrains quantum codes, Quantum 6, 711 (2022).
- [70] N. Delfosse, M. E. Beverland, and M. A. Tremblay, Bounds on stabilizer measurement circuits and obstructions to local implementations of quantum ldpc codes, arXiv:2109.14599 (2021).
- [71] D. Gottesman, Fault-tolerant quantum computation with constant overhead, arXiv preprint arXiv:1310.2984 (2013).
- [72] M. Beverland, V. Kliuchnikov, and E. Schoute, Surface code compilation via edge-disjoint paths, PRX Quantum 3, 020342 (2022).
- [73] Q. Zhuang and B. Zhang, Quantum communication capacity transition of complex quantum networks, Physical Review A 104, 022608 (2021), 2011.07397.
- [74] Y. Lee, E. Bersin, A. Dahlberg, S. Wehner, and D. Englund, A quantum router architecture for high-fidelity entanglement flows in quantum networks, npj Quantum Information 8, 75 (2022), 2005.01852.
- [75] C. Harney and S. Pirandola, Analytical Methods for High-Rate Global Quantum Networks, PRX Quantum 3, 010349 (2022), 2104.10701.
- [76] T. Peham, L. Burgholzer, and R. Wille, On Optimal Subarchitectures for Quantum Circuit Mapping, arXiv:2210.09321.
- [77] A. Geng, A. Moghiseh, C. Redenbach, and K. Schladitz, Improved FRQI on superconducting processors and its restrictions in the NISQ era, Quantum Information Processing 22, 104 (2023).
- [78] J. T. Muhonen, A. Laucht, S. Simmons, J. P. Dehollain, R. Kalra, F. E. Hudson, S. Freer, K. M. Itoh, D. N. Jamieson, J. C. McCallum, A. S. Dzurak, and

A. Morello, Quantifying the quantum gate fidelity of single-atom spin qubits in silicon by randomized benchmarking, Journal of Physics: Condensed Matter **27**, 154205 (2015), 1410.2338.

- [79] K. Saeedi, S. Simmons, J. Z. Salvail, P. Dluhy, H. Riemann, N. V. Abrosimov, P. Becker, H.-J. Pohl, J. J. L. Morton, and M. L. W. Thewalt, Room-Temperature Quantum Bit Storage Exceeding 39 Minutes Using Ionized Donors in Silicon-28, Science **342**, 830 (2013).
- [80] M. Mirhosseini, A. Sipahigil, M. Kalaee, and O. Painter, Superconducting qubit to optical photon transduction, Nature 588, 599 (2020), 2004.04838.
- [81] G. Arnold, M. Wulf, S. Barzanjeh, E. S. Redchenko, A. Rueda, W. J. Hease, F. Hassani, and J. M. Fink, Converting microwave and telecom photons with a silicon photonic nanomechanical interface, Nature Communications 11, 4460 (2020), 2002.11628.
- [82] V. Krutyanskiy, M. Meraner, J. Schupp, V. Krcmarsky, H. Hainzer, and B. P. Lanyon, Light-matter entanglement over 50 km of optical fibre, npj Quantum Information 5, 72 (2019), 1901.06317.
- [83] A. Stolk, K. v. d. Enden, M.-C. Roehsner, A. Teepe, S. Faes, C. Bradley, S. Cadot, J. v. Rantwijk, I. t. Raa, R. Hagen, A. Verlaan, J. Biemond, A. Khorev, R. Vollmer, M. Markham, A. Edmonds, J. Morits, T. Taminiau, E. v. Zwet, and R. Hanson, Telecom-Band Quantum Interference of Frequency-Converted Photons from Remote Detuned NV Centers, PRX Quantum 3, 020359 (2022).
- [84] R. D. Delaney, M. D. Urmey, S. Mittal, B. M. Brubaker, J. M. Kindem, P. S. Burns, C. A. Regal, and K. W. Lehnert, Superconducting-qubit readout via lowbackaction electro-optic transduction, Nature 606, 489 (2022).
- [85] O. Alibart, V. D'Auria, M. De Micheli, F. Doutre, F. Kaiser, L. Labonté, T. Lunghi, É. Picholle, and S. Tanzilli, Quantum photonics at telecom wavelengths based on lithium niobate waveguides, Journal of Optics 18, 104001 (2016).
- [86] K. Azuma, S. E. Economou, D. Elkouss, P. Hilaire, L. Jiang, H.-K. Lo, and I. Tzitrin, Quantum repeaters: From quantum networks to the quantum internet, arXiv:2212.10820.
- [87] M. Pompili, S. L. N. Hermans, S. Baier, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. d. S. Martins, B. Dirkse, S. Wehner, and R. Hanson, Realization of a multinode quantum network of remote solid-state qubits, Science **372**, 259 (2021), 2102.04471.
- [88] D. P. Nadlinger, P. Drmota, B. C. Nichol, G. Araneda, D. Main, R. Srinivas, D. M. Lucas, C. J. Ballance, K. Ivanov, E. Y. Tan, P. Sekatski, R. L. Urbanke, R. Renner, N. Sangouard, and J. D. Bancal, Experimental quantum key distribution certified by bell's theorem, Nature 607, 682 (2022), original title: Device-Independent Quantum Key DistributionChanged in response to referees. See referee reports on file.
- [89] H. H. Vallabhapurapu, I. Hansen, C. Adambukulam, R. Stohr, A. Denisenko, C. H. Yang, and A. Laucht, High Fidelity Control of a Nitrogen-Vacancy Spin Qubit at Room Temperature using the SMART Protocol, arXiv:2208.14671.
- [90] G. Waldherr, Y. Wang, S. Zaiser, M. Jamali, T. Schulte-

Herbrüggen, H. Abe, T. Ohshima, J. Isoya, J. F. Du, P. Neumann, and J. Wrachtrup, Quantum error correction in a solid-state hybrid spin register, Nature **506**, 204 (2014), 1309.6424.

- [91] Alternatively, the photons could themselves be considered as qubits to directly mediate nonlocal interactions through strong coupling schemes [145], which could be made loss-tolerant under a broker-client [146] model of entanglement distribution.
- [92] S. D. Barrett and P. Kok, Efficient high-fidelity quantum computation using matter qubits and linear optics, Physical Review A 71, 060310 (2005), quantph/0408040.
- [93] R. Raussendorf, D. E. Browne, and H. J. Briegel, Measurement-based quantum computation on cluster states, Physical Review A 68, 022312 (2003), quantph/0301052.
- [94] R. Jozsa, An introduction to measurement based quantum computation, arXiv:quant-ph/0508124.
- [95] D. Gottesman and I. L. Chuang, Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations, Nature 402, 390 (1999), quant-ph/9908010.
- [96] J. Eisert, K. Jacobs, P. Papadopoulos, and M. B. Plenio, Optimal local implementation of nonlocal quantum gates, Physical Review A 62, 10.1103/physreva.62.052317 (2000).
- [97] C. A. Pérez-Delgado and J. F. Fitzsimons, Iterated Gate Teleportation and Blind Quantum Computation, Physical Review Letters 114, 220502 (2015), 1411.4777.
- [98] N. S. Minaev and A. V. Mudryi, Thermally-induced defects in silicon containing oxygen and carbon, physica status solidi (a) 68, 561 (1981).
- [99] D. B. Higginbottom, A. T. K. Kurkjian, C. Chartrand, M. Kazemi, N. A. Brunelle, E. R. MacQuarrie, J. R. Klein, N. R. Lee-Hone, J. Stacho, M. Ruether, C. Bowness, L. Bergeron, A. DeAbreu, S. R. Harrigan, J. Kanaganayagam, D. W. Marsden, T. S. Richards, L. A. Stott, S. Roorda, K. J. Morse, M. L. W. Thewalt, and S. Simmons, Optical observation of single spins in silicon, Nature **607**, 266 (2022).
- [100] A. M. Tyryshkin, S. Tojo, J. J. Morton, H. Riemann, N. V. Abrosimov, P. Becker, H.-J. Pohl, T. Schenkel, M. L. Thewalt, K. M. Itoh, *et al.*, Electron spin coherence exceeding seconds in high-purity silicon, Nature materials **11**, 143 (2012).
- [101] Y. Xiong, C. Bourgois, N. Sheremetyeva, W. Chen, D. Dahliah, H. Song, S. M. Griffin, A. Sipahigil, and G. Hautier, High-throughput identification of spinphoton interfaces in silicon (2023).
- [102] E. M. Purcell, H. C. Torrey, and R. V. Pound, Resonance absorption by nuclear magnetic moments in a solid, Phys. Rev. 69, 37 (1946).
- [103] P. Neumann, J. Beck, M. Steiner, F. Rempp, H. Fedder, P. R. Hemmer, J. Wrachtrup, and F. Jelezko, Singleshot readout of a single nuclear spin, Science **329**, 542 (2010).
- [104] A. N. Safanov, E. C. Lightowlers, and G. Davies, Carbon-hydrogen deep level luminescence centre in silicon responsible for the t-line, Materials Science Forum 196-201, 909 (1995).
- [105] T. Müller, I. Aharonovich, L. Lombez, Y. Alaverdyan, A. N. Vamivakas, S. Castelletto, F. Jelezko, J. Wrachtrup, S. Prawer, and M. Atatüre, Wide-

range electrical tunability of single-photon emission from chromium-based colour centres in diamond, New Journal of Physics **13**, 075001 (2011), 1101.4911.

- [106] Manuscript in preparation.
- [107] A. L. Crook, C. P. Anderson, K. C. Miao, A. Bourassa, H. Lee, S. L. Bayliss, D. O. Bracher, X. Zhang, H. Abe, T. Ohshima, E. L. Hu, and D. D. Awschalom, Purcell Enhancement of a Single Silicon Carbide Color Center with Coherent Spin Control, Nano Letters 20, 3427 (2020), 2003.00042.
- [108] Manuscript in preparation.
- [109] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, Entanglement of single-atom quantum bits at a distance, Nature 449, 68 (2007).
- [110] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, Heralded entanglement between solid-state qubits separated by three metres, Nature **497**, 86 (2013).
- [111] J. Metz and S. D. Barrett, Effect of frequencymismatched photons in quantum-information processing, Phys. Rev. A 77, 042323 (2008).
- [112] C. D. Wilen, S. Abdullah, N. A. Kurinsky, C. Stanford, L. Cardani, G. D'Imperio, C. Tomei, L. Faoro, L. B. Ioffe, C. H. Liu, A. Opremcak, B. G. Christensen, J. L. DuBois, and R. McDermott, Correlated charge noise and relaxation errors in superconducting qubits, Nature 594, 369 (2021).
- [113] J. Yoneda, J. S. Rojas-Arias, P. Stano, K. Takeda, A. Noiri, T. Nakajima, D. Loss, and S. Tarucha, Noisecorrelation spectrum for a pair of spin qubits in silicon (2022).
- [114] J. S. Rojas-Arias, A. Noiri, P. Stano, T. Nakajima, J. Yoneda, K. Takeda, T. Kobayashi, A. Sammak, G. Scappucci, D. Loss, and S. Tarucha, Spatial noise correlations beyond nearest-neighbor in ²⁸Si/SiGe spin qubits (2023).
- [115] D. I. Schuster, A. Wallraff, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. M. Girvin, and R. J. Schoelkopf, ac Stark Shift and Dephasing of a Superconducting Qubit Strongly Coupled to a Cavity Field, Physical Review Letters 94, 123602 (2005), cond-mat/0408367.
- [116] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis, Demonstration of the trapped-ion quantum CCD computer architecture, Nature **592**, 209 (2021), 2003.01293.
- [117] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels, Physical Review Letters 76, 722 (1996), quant-ph/9511027.
- [118] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. W. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, Entanglement distillation between solid-state quantum network nodes, Science **356**, 928 (2017), 1703.03244.
- [119] Y. Li and S. C. Benjamin, High threshold distributed quantum computing with three-qubit nodes, New Journal of Physics 14, 093008 (2012).
- [120] A. Reiserer, N. Kalb, M. S. Blok, K. J. van Bemmelen, T. H. Taminiau, R. Hanson, D. J. Twitchen, and

M. Markham, Robust quantum-network memory using decoherence-protected subspaces of nuclear spins, Physical Review X **6**, 10.1103/physrevx.6.021040 (2016).

- [121] Manuscript in preparation.
- [122] L. Jiang, M. V. G. Dutt, E. Togan, L. Childress, P. Cappellaro, J. M. Taylor, and M. D. Lukin, Coherence of an optically illuminated single nuclear spin qubit, Physical Review Letters **100**, 10.1103/physrevlett.100.073001 (2008).
- [123] M. S. Blok, N. Kalb, A. Reiserer, T. H. Taminiau, and R. Hanson, Towards quantum networks of single spins: analysis of a quantum memory with an optical interface in diamond, Faraday Discuss. 184, 173 (2015).
- [124] A. J. R. MacDonald, G. G. Popowich, B. D. Hauer, P. H. Kim, A. Fredrick, X. Rojas, P. Doolin, and J. P. Davis, Optical microscope and tapered fiber coupling apparatus for a dilution refrigerator, Review of Scientific Instruments 86, 013107 (2015).
- [125] M. Häußler, R. Terhaar, M. A. Wolff, H. Gehring, F. Beutel, W. Hartmann, N. Walter, M. Tillmann, M. Ahangarianabhari, M. Wahl, T. Röhlicke, H.-J. Rahn, W. H. P. Pernice, and C. Schuck, Scaling waveguide-integrated superconducting nanowire singlephoton detector solutions to large numbers of independent optical channels, arXiv:2207.12060.
- [126] I. Hansen, A. E. Seedhouse, A. Saraiva, A. Laucht, A. S. Dzurak, and C. H. Yang, Pulse engineering of a global field for robust and universal quantum computation, Physical Review A 104, 062415 (2021), 2108.00776.
- [127] E. Vahapoglu, J. P. Slack-Smith, R. C. C. Leon, W. H. Lim, F. E. Hudson, T. Day, J. D. Cifuentes, T. Tanttu, C. H. Yang, A. Saraiva, N. V. Abrosimov, H.-J. Pohl, M. L. W. Thewalt, A. Laucht, A. S. Dzurak, and J. J. Pla, Coherent control of electron spin qubits in silicon using a global field, npj Quantum Information 8, 126 (2022), 2107.14622.
- [128] B. E. Kane, A silicon-based nuclear spin quantum computer, Nature **393**, 133 (1998).
- [129] B. P. Wuetz, P. L. Bavdaz, L. A. Yeoh, R. Schouten, H. v. d. Does, M. Tiggelman, D. Sabbagh, A. Sammak, C. G. Almudever, F. Sebastiano, J. S. Clarke, M. Veldhorst, and G. Scappucci, Multiplexed quantum transport using commercial off-the-shelf CMOS at sub-kelvin temperatures, npj Quantum Information 6, 43 (2020), 1907.11816.
- [130] S. Doerner, A. Kuzmin, S. Wuensch, I. Charaev, F. Boes, T. Zwick, and M. Siegel, Frequencymultiplexed bias and readout of a 16-pixel superconducting nanowire single-photon detector array, Applied Physics Letters **111**, 032603 (2017), 1705.05345.
- [131] B. R. d. Supinski, M. Hall, T. Gamblin, T. Häner, D. S. Steiger, T. Hoefler, and M. Troyer, Distributed Quantum Computing with QMPI, SC21: International Conference for High Performance Computing, Networking, Storage and Analysis **00**, 1 (2021), 2105.01109.
- [132] H. J. Kimble, The quantum internet, Nature 453, 1023 (2008).

- [133] Y. Liu, W.-J. Zhang, C. Jiang, J.-P. Chen, C. Zhang, W.-X. Pan, D. Ma, H. Dong, J.-M. Xiong, C.-J. Zhang, *et al.*, Experimental twin-field quantum key distribution over 1000 km fiber distance, Physical Review Letters **130**, 210801 (2023).
- [134] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, Quantum repeaters: The role of imperfect local operations in quantum communication, Phys. Rev. Lett. 81, 5932 (1998).
- [135] S. Muralidharan, L. Li, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, Optimal architectures for long distance quantum communication, Scientific reports 6, 20463 (2016).
- [136] L. Jiang, J. M. Taylor, K. Nemoto, W. J. Munro, R. V. Meter, and M. D. Lukin, Quantum repeater with encoding, Physical Review A 79, 032325 (2009), pauli frame, 0809.3629.
- [137] S. Wehner, D. Elkouss, and R. Han-Quantum internet: Α vision for the son. ahead, Science **362**, eaam9288(2018)road https://www.science.org/doi/pdf/10.1126/science.aam9288.
- [138] C. Simon, Towards a global quantum network, Nature Photonics 11, 678 (2017).
- [139] W. Gao, P. Fallahi, E. Togan, A. Delteil, Y. Chin, J. Miguel-Sanchez, and A. Imamoğlu, Quantum teleportation from a propagating photon to a solid-state spin qubit, Nature Communications 4, 2744 (2013), 1307.1142.
- [140] C. Panayi, M. Razavi, X. Ma, and N. Lütkenhaus, Memory-assisted measurement-device-independent quantum key distribution, New Journal of Physics 16, 043005 (2014), 1309.3406.
- [141] H.-K. Lo, M. Curty, and B. Qi, Measurement-Device-Independent Quantum Key Distribution, Physical Review Letters 108, 130503 (2012), 1109.1473.
- [142] X. Liu, J. Liu, R. Xue, H. Wang, H. Li, X. Feng, F. Liu, K. Cui, Z. Wang, L. You, Y. Huang, and W. Zhang, 40user fully connected entanglement-based quantum key distribution network without trusted node, PhotoniX 3, 2 (2022).
- [143] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, A survey on homomorphic encryption schemes: Theory and implementation, ACM Computing Surveys (Csur) 51, 1 (2018).
- [144] T. Kapourniotis, E. Kashefi, D. Leichtle, L. Music, and H. Ollivier, A Framework for Verifiable Blind Quantum Computation, arXiv 10.48550/arxiv.2206.00631 (2022), 2206.00631.
- [145] C. M. Knaut, A. Suleymanzade, Y.-C. Wei, D. R. Assumpcao, P.-J. Stas, Y. Q. Huan, B. Machielse, E. N. Knall, M. Sutula, G. Baranes, *et al.*, Entanglement of nanophotonic quantum memory nodes in a telecommunication network, arXiv preprint arXiv:2310.01316 (2023).
- [146] S. C. Benjamin, D. E. Browne, J. Fitzsimons, and J. J. Morton, Brokered graph-state quantum computation, New Journal of Physics 8, 141 (2006).