WHITE PAPER

Dr. Stephanie Simmons
November 8, 2023

What could networks of quantum supercomputers look like?

Modern computers and communication networks based upon semiconductor physics have already given us so much, yet even more capabilities await if we can use the laws of quantum physics to process information. We already know there are a huge range of advancements which will be enabled by large-scale, reliable quantum computers and networks—it's only a matter of time.

Already today, the potential of quantum is clear in far-ranging areas: for understanding and designing chemicals, drugs and materials suitable for all sectors, for revolutionizing cybersecurity, for understanding the catalysts needed to mitigate climate change, and many more. And there's likely several unforeseen quantum applications which we will only discover once the underlying technology is built.

But, right now, quantum networks and quantum computers are small-scale and unreliable, and, because of that, are in the early stages of meeting their potential.

From a system engineering perspective, there are a few lessons we have learned about quantum technology design for the technology to successfully offer reliability at scale—and begin to unlock some of quantum's most beneficial impacts for people around the world.

Lesson 1

All platforms will evolve towards the construction of networked quantum modules.

All quantum systems accumulate errors from unintentionally interacting with the environment, as well as from imperfections in the intentional manipulations of those systems.

Ultimately, the technical perfection needed to make qubits reliable is an almost impossibly high bar at any scale. This high degree of perfection means all qubits (and there are many kinds under development) must operate in a strictly controlled environment – some kind of box. And most quantum-grade environments have a maximum-capacity problem in that only so many qubits can fit in each box, and this amounts to a scaling barrier at about 1,000 qubits or so for most qubit types. We need significantly more qubits than that to realize quantum's promise for humanity.

In some sense, this isn't a new problem. Today's computers and networks got around their scaling problems by scaling horizontally – by linking many cores or modules together with telecom-band light in fibre-optic cables.

Quantum technologies will ultimately follow that same path. Individual particles of telecom light – called "telecom photons"– are extremely reliable qubits. Indeed, they are the most reliable way to transmit quantum information at room temperature. The only "box" they need to perform well is a regular fibre-optic cable. Through the input/output (I/O) ports of quantum modules, telecom photons will be used to network quantum modules together. In this way, arbitrarily large quantum networks and quantum computers can be built. Networked modules bring scale and, thus, compute power!

All-to-all Connectivity (networked multi-module chips)

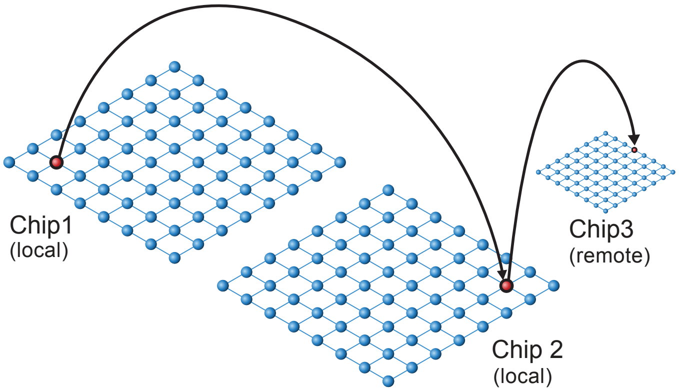

Lesson 2

I/O *quantity* matters, so, ideally, we want quantum I/O for each qubit.

One of the key differences between today's digital networked supercomputers and future networked quantum supercomputers is the enormous number of input/output (I/O) ports that will be required. For classical computers, we're used to thinking of network connectivity at the system level. For quantum computers, anything less than network connectivity at the qubit level will produce a bottleneck that throttles total system performance.

Quantum I/O is used to interact with the qubits for computing and networking and, critically, to transmit and receive quantum information stored in telecom photon qubits.

For the most exciting quantum applications we know of already, there's an enormous and unavoidable amount of I/O needed across modules. For example, for the most efficient kinds of reliable quantum information processing, almost every physical qubit in the entire module must be directly linked with a partner qubit in one of the other modules whenever any of their reliable qubits need to perform any quantum logic together.

This is massively more I/O than modern computing networks demand, the more quantum I/O the better – in the extreme case having an I/O for each qubit will still offer performance gains over anything with less I/O. Quantum information is just fundamentally different in this way. Anything less than an I/O per qubit will yield a bottleneck that throttles total system performance.

One thing will remain the same for quantum networks and computers: people will want them to be as fast as possible. This means that scalable, reliable quantum systems should be designed around their core I/O capability, rather than have the I/O requirements addressed as an afterthought.

Lesson 3

I/O *quality* matters, so, ideally, we want photons to be collected and transmitted as efficiently as possible.

Fortunately, some qubits offer a “built-in” quantum I/O port. They have a photonic interface that can directly interface fibre-optic photon qubits with stationary qubits (the unmoving qubits in a the quantum system that are made out of, ex spin, ion-trap, superconducting, and so on.) Some of these qubits can interface directly to telecom photons, which means they can be used directly on the existing, vast telecom infrastructure.

Other qubit types without a built-in I/O port need a transducer—a kind of quantum “translator”—to change to the telecom wavelength. But this comes at a price. Transduction isn’t always successful (10% success rates are considered good), and transduction typically adds noise and complexity. This translates into a 10x slowdown of the computation!

Clearly, it would be great to do away with transduction entirely if possible, to maximize the I/O throughput for each I/O port.

Even for many of the qubits that do have a built-in quantum I/O port, it’s actually quite hard to collect individual photons. Just like a lightbulb, light typically naturally emits in all directions. For lightbulb-like quantum systems that operate in a vacuum, the act of collecting a photon isn’t always successful; 5% success rates are considered good. This is an even bigger, 20x slowdown!

However, for specific solid-state qubits, almost all the light can be naturally collected and directed because it’s as if the emitter is already embedded directly inside light-guiding fibre optics. Thus, it would be great to avoid low collection efficiencies entirely by working with a solid-state material to maximize the I/O throughput for each I/O port.

Lesson 4

Connectivity is key.

We can do amazing things if we make the simple assumption that any qubit could be made to interact directly with any other qubit throughout the networked quantum computer through their photons – not just the physical qubits that happen to be in direct proximity. This reduces all kinds of wasteful computational steps that are otherwise taken just to move quantum information around within a box. Interestingly, it also makes the techniques for achieving 'fault -tolerance' much more efficient as it requires far, far fewer resources overall for the same application – think 100x fewer qubits or better!

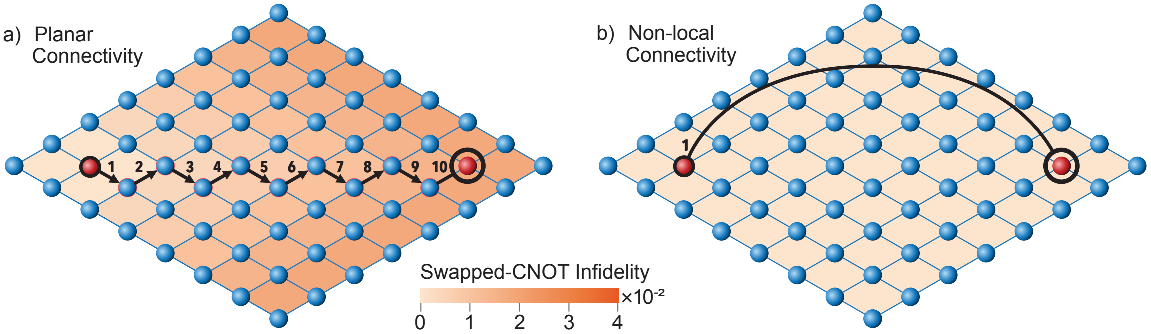

Comparison between two-qubit operations on devices with (a) planar, nearest neighbour connectivity graphs and (b) non-local connectivity. A CNOT gate between two distant qubits on the planar graph requires sequential swap operations. Each consecutive operation accumulates errors and the overall fidelity 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.

Photonic Inc. is addressing these lessons to develop horizontally scalable, fault-tolerant quantum technologies.

The requirements needed to make quantum technologies reliable are still quite high and well known – long quantum lifetimes, high fidelity control, etc. But, importantly, these scalability lessons mean that these targets are not any more difficult to meet.

Photonic's quantum computing and networking platform is ideally positioned to address these challenges and offer a scalable high-performing quantum system.

This is driven by the way in which we leverage silicon colour centres with integrated photonics, specifically the T centre, to achieve the following core benefits:

Networkable and horizontally scalable

The T centre emits at the telecom wavelength, making it a perfect candidate to network large numbers of chips together. It can, therefore, make the system horizontally scalable, much like today's digital supercomputers.

I/O quantity

Photonic's technology is able to achieve non-local connectivity, which means that each qubit has its own I/O and can interact with any qubit within the module/chip, or in other modules.

I/O quality

Again, because T centres emit in the telecom wavelength we avoid the high-loss process that would be required to convert the wavelength. Additionally, since we are using a solid-state material, we maximize the I/O throughput of each I/O.

Connectivity and fault-tolerance

Access to non-local connectivity means we can utilize highly efficient quantum error-correcting techniques and, therefore, significantly minimize the number of error-prone qubits required to represent a single fault-tolerant qubit.

Ultimately, the accessibility of scalable and reliable quantum technologies will enable humanity to solve problems that are currently beyond our capabilities, held back by the inescapable constraints of today's digital computing.

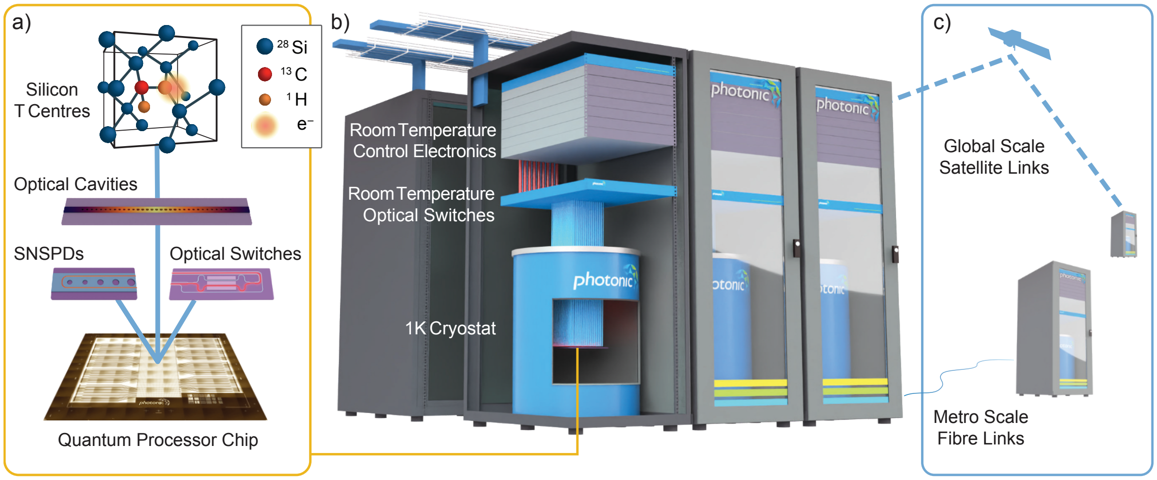

Photonic's scalable quantum technology architecture. A quantum chip is cooled in a 1K cryostat. This chip hosts integrated silicon T centres within optical cavities, photonic switches, and single photon detectors. Optical input-output (IO) ports via telecom fibre connects to a room temperature photonic 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.

For more details on Photonic's approach to this world-changing technology, please visit www.photonic.com.

Dr. Stephanie Simmons
Chief Quantum Officer

Dr. Stephanie Simmons is the Founder and Chief Quantum Officer at Photonic, driving the technical vision for next-generation quantum technologies based on photonically-linked silicon spin qubits. She is a world-leading expert in quantum technologies, silicon spin-photon interfaces, condensed matter spin dynamics and control, silicon-integrated photonics, and quantum optics.