Quantum Internet: a vision for the road ahead

Stephanie Wehner,[1*], David Elkouss,[1], Ronald Hanson,[1,2]

[1]QuTech, Delft University of Technology, PO Box 5046, 2600 GA Delft, The Netherlands
[2]Kavli Institute of Nanoscience, Delft University of Technology,
PO Box 5046, 2600 GA Delft, The Netherlands
[*]To whom correspondence should be addressed; E-mail: s.d.c.wehner@tudelft.nl.

The internet - a vast network that enables simultaneous long-range classical communication - has had a revolutionary impact on our world. The vision of a quantum internet is to fundamentally enhance internet technology by enabling *quantum* communication between any two points on earth. Such a quantum internet may operate in parallel to the internet that we have today, and connect quantum processors in order to achieve capabilities that are provably impossible using only classical means. Here, we propose stages of development towards a full-blown quantum internet, and highlight experimental and theoretical progress to attain them.

Introduction

A quantum internet enables us to solve problems that are fundamentally out of reach for the classical internet [1]. The most well-known application of a quantum internet is quantum key distribution (QKD), which enables two remote network nodes to establish an encryption key

[1]Some researchers still go one step further and believe all communication will eventually be done over quantum channels (*1*).

whose security only relies on the laws of quantum mechanics. A quantum internet, however, has many other applications ranging from secure access to remote quantum computers (*2*), clock synchronization (*3*), and even extending the baselines of telescopes (*4*) (see Figure 1). Moreover, as quantum internet research expands, other useful applications will likely be discovered in the next decade.

Central to all these applications is that a quantum internet enables us to send quantum bits (qubits), which are fundamentally different than classical bits. While classical bits can take only two values, '0' and '1', qubits can be in a superposition of being '0' and '1' at the same time. Importantly, qubits cannot be copied, and in fact any attempt to do so can be detected. It is this feature that makes qubits naturally well suited for security applications, but at the same time makes transmitting qubits over long distances a truly formidable endeavour. Since qubits canot be copied or amplified, repetition or signal amplification are ruled out as a means to overcome imperfections, and a radically new technological development is needed in order to build a quantum internet (*5*) (see Figure 2).

We are now at an exciting moment in time, akin to the eve of the classical internet. In late 1969, the first message was sent over the nascent four node network then still referred to as ARPANET. Recent technological progress (*6–9*) now predicts that we may see the first small-scale implementations of a quantum internet within the next five years.

At first glance, realizing a quantum internet (see Figure 3) may seem even more difficult than building a large scale quantum computer. After all, we might imagine that in full analogy to the classical internet, the ultimate version of a *quantum* internet consists of fully-fledged quantum computers that can exchange an essentially arbitrary number of qubits. Thankfully, it turns out that many quantum network protocols do not require large quantum computers to be realized: a quantum device with a single qubit at the end point is already sufficient for many applications. What's more, errors in quantum internet protocols can often be dealt with using classical rather

than quantum error correction, imposing fewer demands on the control and quality of the qubits. The reason why quantum internet protocols can outperform classical communication with such relatively modest resources is due to the fact that their advantages rely solely on inherently quantum properties such as quantum entanglement, which can be exploited already with very few qubits. In contrast, a quantum computer must feature more qubits than can be simulated on a classical computer in order to offer a computational advantage. Given the challenges posed by the development of a quantum internet, it is useful to reflect on what capabilities are needed to achieve specific quantum applications, and what technology is required to realize them.

Here, we identify stages of development towards a full-blown quantum internet. These stages are functionality driven: central to their definition is not experimental difficulty itself, but the essential question of what is needed to actually enable useful applications. Each stage is interesting in its own right, and distinguished by a specific quantum functionality that is sufficient to support a certain class of protocols. To illustrate, we give examples of known application protocols in each stage where a quantum internet is already known to bring advantages.

Realizing a quantum internet demands significant development to realize quantum repeaters as well as end nodes. It is clear that in the short term, one may optimize both repeaters and end nodes relatively independently. That is, one can imagine a quantum internet using relatively simple end nodes, while using repeaters powerful enough to cover larger distances. Similarly, a near term quantum internet may be optimized for shorter - for example, pan European - distances, while employing much more powerful end nodes capable of realizing a larger set of protocols. Ideally, all these designs ensure forward compatibility to achieve the ultimate goal of a full-blown world wide quantum internet. We note that while the intermediate repeater nodes need to be able to support the functionality of each stage, an application centric view makes no other statements regarding the capabilities of the repeater nodes of the network.

Finally, we discuss progress towards implementing a quantum internet, which poses signif-

icant challenges to physics, engineering and computer science.

Stages of functionality and applications

Let us formulate the functionality driven stages of quantum internet development. Each stage is distinguished by an increasing amount of functionality, at the expense of increasing experimental difficulty. We say that an experimental implementation has reached a certain stage only if the functionality of that stage and all previous stages (Figure 4) is available to all the end nodes using the network.

Crucial to the distinction between the stages is that the subsequent stage offers a fundamentally new functionality not available in the previous one, rather than simply improving parameters or offering "more of the same" by increasing the number of qubits. For the sake of clarity, the stages and tests below target systems that prepare and transmit qubits, but it is also possible to phrase both in terms of qudits or continuous variables. For each stage, we describe some of the application protocols that are already known and that can be realized with the functionality provided in that stage. It is conceivable that a simpler protocol, or better theoretical analysis, may be found in the future that solves the same task, but is less demanding in terms of functionality. In parallel to the daunting experimental challenges in making quantum internet a reality, there is thus a challenge for quantum software developers to design protocols that can realize a task in a stage that can be implemented more easily. We identify relevant parameters for each stage to establish a common language between hardware and software developers. These parameters can be estimated using a series of simple tests, allowing us to certify the performance of an experimental implementation in attaining a specific stage, as well as the performance of protocols depending on these parameters. So far, most application protocols have only been analyzed for perfect parameters. As such, the exact requirements of many application protocols on these parameters have not yet been determined, and deserve much needed future investigation.

While functionality driven stages make demands on the communication links and quantum repeaters, it will not be important in this section how these links are realized: they may be realized by direct transmission in fiber, by being relayed by any kind of quantum repeater, or even by means of teleportation using pre-shared entanglement. What matters is that these links can be used to generate the necessary quantum states for a specific stage.

Trusted Repeater Networks. The first stage differs significantly from the others in the sense that it does *not* allow the end to end transmission of qubits. Nevertheless, from a technological perspective, trusted repeater networks can form an interesting stepping stone towards a quantum internet, spurring infrastructure deployment and engineering developments, and - depending on the underlying technology - trusted repeaters (*10*) can be upgraded to true quantum repeaters later on.

Specifically, a trusted repeater network (sometimes also called a trusted node network) has at least two end nodes, and a sequence of short distance links connecting nearby intermediary repeater nodes. Each pair of adjacent nodes uses quantum key distribution (QKD) (*11–13*) to exchange encryption keys. These pairwise keys allow the end nodes to generate their own key provided that all intermediary nodes are trusted (*14*). A first step towards upgrading such networks could be measurement-device independent QKD (*15–17*), which is a QKD protocol secure even with untrusted measurement devices that can be implemented with standard optical components and sources (*17*) and which already incorporate some useful ingredients for later stages such as two-photon Bell measurements.

Prepare and Measure Networks. This stage is the first to offer end-to-end quantum functionality. It is sufficient to perform end-to-end QKD without the need to trust intermediary repeater nodes, and already allows a host of protocols for other interesting tasks. Informally, this stage allows any node to prepare a one qubit state, and transmit the resulting state to any other node who then measures it (see Figure 5 for a definition). Transmission and measurement

are allowed to be post-selected, that is, a signal that the qubit is lost may be generated instead. For instance, the receiving node is allowed to ignore non-detection events and conclude such qubits are lost. If the sender can prepare an entangled state of two qubits, then this stage also includes the special case where the sender transmits the first and second qubit to two different nodes in the network (or to another node and itself). Such entanglement distribution is then also post-selected.

It is important to note that such a post-selected prepare and measure functionality is *not* equivalent to transmitting arbitrary qubits across the network (*18*). The task of transmitting arbitrary qubits demands the ability to transfer an *unknown* state $|\Psi\rangle$ (which the sender does not know how to prepare), *deterministically* to the receiver, that is, no post-selection on detection events is allowed.

The classical reader may wonder what is the use of transmitting qubits at all, if there is a procedure for the sender to prepare the state $|\Psi\rangle$. After all, we might imagine that the sender simply sends classical instructions for this procedure to the receiver, who then prepares the qubit himself. The difference when sending different quantum states $|\Psi\rangle$ themselves is that an eavesdropper, or indeed the receiver, can not make a copy of $|\Psi\rangle$ without disturbing the quantum state. This means that attempts to gain information from $|\Psi\rangle$ by an eavesdropper may be detected, enabling QKD.

Application Protocols: This stage is already sufficient to realize protocols for many interesting cryptographic tasks, as long as the probability of loss p, and the inaccuracies in transmission ε_T and measurement ε_M (Figure 5) are sufficiently low. The most famous of such tasks is quantum key distribution (QKD), which provides a solution to the task of generating a secure encryption key between two distant end nodes (Alice and Bob) (*11–13*). Importantly, QKD is secure even if the eavesdropper who trying to learn the key has access to an arbitrarily large quantum computer to attack the protocol, and remains secure at any point in the future even if

such a quantum computer becomes available later on. This is provably impossible using classical communication. The BB84 QKD (*11*) protocol can be realized using only single qubit preparations and measurements tolerating some amount of post-selection p (*19*). For known protocols in this stage, $\varepsilon_T + \varepsilon_M \leq 0.11$ is sufficient, and can be estimated by testing for only a small number of states (*20*). In practise, single qubit preparation can be replaced by attenuated laser pulses, using also decoy state BB84 to guarantee security (*21*). QKD is commercially available at short distance using standard telecom fibers (*22*), and a variety of protocols are known (see (*23*) for a survey).

Another class of protocols in that stage lies the domain of two-party cryptography. Here, there is no eavesdropper, but rather Alice and Bob themselves do not trust each other. An example of such a task is secure identification, where Alice (a potentially impersonating user) may wish to identify herself to Bob (a potentially malicious server or ATM) without revealing her authentication credentials (*24, 25*). It is known that even using quantum communication such tasks cannot be implemented securely without imposing assumptions on the power of the adversary (*26–28*). Classical protocols rely on computational assumptions, whose security against an attacker who holds a quantum computer is unclear. Nevertheless, it is possible to achieve provable security for all such relevant tasks by sending more qubits than the adversary can store easily within a short time frame, known as the bounded (*29*), or more generally noisy-storage model (*30, 31*). Significantly, this assumption only needs to hold during the execution of the protocol, and security is preserved into the future even if the adversary later obtains a better quantum memory. There exist protocols for which it is again sufficient to prepare and measure single qubits, where the sufficient values of $p, \varepsilon_M, \varepsilon_T$ (see Figure 5) depend on the storage assumption (*32*).

Other known protocols in this stage include position verification (*33*), weakened forms of two-party cryptographic tasks that can form building blocks like imperfect bit commit-

ments (*34*) and coin-flipping (*35*). Here, requirements in terms of p, ε_M and ε_T have not been analyzed. For no tasks, is a full set of necessary and sufficient conditions on these parameters known to date.

Entanglement Distribution Networks. The third stage allows the end-to-end creation of quantum entanglement in a deterministic or heralded fashion, as well as local measurements. Importantly, the end nodes require no quantum memory for this stage (see Figure 5).

The term deterministic entanglement generation refers to the fact that the process succeeds with (near) unit probability. Heralding is a slightly weaker form of deterministic entanglement generation in which we signal the successful generation of entanglement by an event that is *independent* of the (immediate) measurement of the entangled qubits themselves. Here, the generation of entanglement is deterministic conditioned on such a successful heralding signal. Specifically, this prohibits post-selecting on detection events when measuring the entangled qubits. We remark that this stage also includes networks that allow the generation of multi-partite entangled states, followed by immediate measurements, but no memory. The generation of multi-partite entanglement is however not required to attain this stage.

Application protocols: The main advance over the previous stage is that this stage allows the realization of device-independent protocols, where the quantum devices are largely untrusted. Specifically, the concept of device independence (*36, 37*) models the end nodes as black boxes, to which we can give classical instructions to perform specific measurements, and receive the resulting measurement outcomes. No guarantees are given about the actual quantum state or measurements performed by the device, which may even be constructed by the adversary. It is important to note that the classical software used to control such quantum devices *is* trusted, and it is assumed that the quantum device merely exhibits input/output behaviour. In particular, devices can record their inputs and outputs (*38*), but cannot transmit the key back to the adversary. The coordination allowed by entanglement now also in principle allows players to "cheat"

an online bridge game (*39*).

Low errors in preparation (ε_P) and measurement (ε_M) as $\varepsilon_P + \varepsilon_M \leq 0.057$ (Figure 5) are sufficient to ensure the implementability of device independent QKD (*36*), where necessary and sufficient conditions for the parameters to implement general tasks in this stage are unknown.

Quantum Memory Networks. The fourth stage is distinguished by the capability of the end nodes to have local memory, while simultaneously allowing universal local control (see Figure 5). This allows the implementation of much more complex protocols requiring temporary storage of a quantum state during further quantum or classical communication. Examples include protocols for solving distributed systems tasks. This stage also implies the ability to perform entanglement distillation, and generate multi-partite entangled states from bipartite entanglement by exploiting the ability for local memory and control. A crucial difference between this stage and the previous one is thus also that we are now able to transfer *unknown* qubits from one network node to another, for example, by performing deterministic teleportation. We note that this capability is not guaranteed in the previous stage: technology that can be used to deterministically relay qubits over long distances by means of large-scale quantum error correction, implies the technological capability of realizing a good local quantum memory. We emphasize that a quantum memory network does not require operations to be performed with an accuracy that would be above threshold for fault tolerant computation.

An important parameter in application protocols is the number of rounds k (Figure 5). In order to realize useful application protocols, the storage time t thus needs to stand in relation to the communication time in the network instead of an absolute time. This means that networks of nodes which are far apart do in fact need to exhibit longer memory times in order to attain this stage, and the quality of the memory is time dependent. The fact that this time t is related the maximum time that it takes any two nodes to communicate is because a stage is attained only if the functionality is available to *any* two nodes in the network, even the two which are

farthest apart.

Application protocols: The availability of quantum memories, and the deterministic transmission of qubits opens up many new protocols in this stage. We start with cryptographic tasks: To allow clients to make use of these computers securely, that is, without revealing the nature or outcome of their computation, it is possible to perform secure assisted quantum computation (*40*), or blind quantum computation (*2, 41*). Here a simple quantum device capable of preparing and measuring single qubits is sufficient to perform a computation on a large-scale quantum computer such that the quantum computer cannot gain information about the program and result. Note that the fact that we need one large-scale quantum computer does not imply that a quantum computing network (the highest stage) is required to run such protocols: we only need a quantum internet that allows a client to communicate with the computing server. Recall a network attains a specific stage only if the functionality is available to all nodes.

Other cryptographic tasks in this domain are tools such as protocols for the sharing of classical (*42*) or quantum (*43*) secrets, including verifiable secret sharing schemes (*44*), anonymous transmissions (*45*). Evidently, the number of qubits determines the size of the secrets or qubits transmitted, but no fault tolerance is in principle required.

This stage also opens the door to interesting applications outside the domain of cryptography. For example, proposals exist for to exploiting long-distance entanglement to extend the baseline of telescopes (*4*), for basic forms of leader election (*46*), and for improving the synchronization of clocks (*3*). Depending on the demands made on such synchronization, the proposed protocols could be realized with quantum memory or few-qubit fault tolerant networks (see below).

Necessary and sufficient parameter requirements for solving the above mentioned tasks are not yet known in general. It is also conceivable that an improved analysis considering whether deterministic qubit delivery is really necessary, or whether maybe post-selected transmission of

qubits is enough, can push some of the protocols above to a lower stage. Initial results for blind quantum computation indicates this might indeed be the case.

Few Qubit fault-tolerant Networks. The next stage differs by demanding that the local operations can be performed fault-tolerantly, which is considerably more challenging. Fault tolerance is not necessary for many known quantum internet protocols, but fault-tolerant operations available would allow the execution of local quantum computation of high circuit depth, as well as an (in theory) arbitrary extension of storage times to execute protocols with an arbitrary number of rounds of communication.

The term few qubits here refers to the fact that the number of qubits available is still small enough such that the end nodes themselves can be simulated effectively on a classical computer. This does not imply that the entire network can be simulated efficiently, or that there would exist equivalent classical protocols: the effects of entanglement cannot generally be replicated classically.

We remark that we are here only interested in the performce of the fault tolerant scheme, not how it is realized. Fault tolerance implies that all error parameters (Figure 5) of a quantum memory network can be made negligible by adding more resources. As a guideline to relevant experimental parameters we refer to works in *distributed* quantum computing (see e.g. (*47*)).

Application protocols: Having access to fault-tolerant gates allows higher accuracy clock synchronization (*3*), and protocols that require many rounds of communication and high circuit depth to be useful. This includes distributed quantum computing, as well as applications for full-scale quantum computing networks (see below), restricted to few qubits. This could be of great practical interest, especially for applications in the domain of distributed systems, but as with the implementation of quantum algorithms on quantum computers, the power of having only a limited number of qubits at our disposal is an important subject of investigation.

Quantum Computing Networks. The final stage consists of quantum computers that can

arbitrarily exchange quantum communication. In some sense, it breaks with our paradigm that the next stage is not "more of the same". However, in this case we really do gain a new ability, namely finding solutions to computational problems that can no longer be found efficiently on classical computers.

Application protocols: It is clear that this ultimate stage of a quantum internet allows in principle all protocols to be realized. We remark that small-scale versions of the protocols below can also be realized in the few-qubit fault-tolerant stage, and further development may yield more sophisticated protocols and analysis that place them in lower stages.

First, we again focus on cryptography. In this stage, it is possible to perform coin flipping with an arbitrarily small bias (*48, 49*). We can also solve genuinely quantum tasks, like secure multi-party quantum computation, which forms an extension of classical secure function evaluation to the quantum regime. Classically, this means that node j holds an input string x_j, and all n nodes jointly want to compute $y = f(x_1, \ldots, x_n)$. The goal is that malicious nodes cannot infer anything more about the inputs x_j of the honest nodes than they can by observing the output y. An example of such a problem is secure voting, in which $x_j \in \{0, 1\}$ corresponds to the choice one of two possible candidates, and f is the majority function. The quantum version of this primitive (*50*) allows each party to hold a quantum state $|\Psi_j\rangle$ as input, and the parties jointly wish to compute a quantum operation U.

Next, we focus on distributed systems, formed when several computing devices are connected, sometimes colloquially referred to as a cloud. Many challenges arise in the coordination and control of such systems that may be less familiar to a physicist. As a very simple example, consider a bank transaction being recorded redundantly on several backup servers. If one or more of the backup servers fail during the update, then they may later show inconsistent data (e.g., \$1M vs. \$0). Tool protocols for achieving consensus between processors are widely deployed in practice, for example in Google's Chubby system (*51*). Outside the domain of the

internet itself, examples include the reliability in smart grids, flight control systems, and sensor arrays.

While this area is presently much less developed in the quantum domain (*52*), several protocols are known that show that a quantum internet has great potential for solving the problems in distributed systems much more efficiently than what is possible classically. Very intuitively, the reason why quantum communication can help solve these problems is that entanglement allows coordination among distant processors that greatly surpasses what is possible classically. It is this fact that yields advantages for distributed systems tasks such as consensus and agreement.

One of the most striking examples of a quantum advantage in distributed systems can be found for the task of byzantine agreement. Here, the goal is to allow n processors to agree on a common bit, while some fraction of them may be faulty. The term byzantine refers to the very demanding model of arbitrarily correlated faults, in which the faulty processors essentially collude to thwart the protocol. In (*53*) it is shown that in some regimes, there exists a quantum protocol to solve this task using only a constant number of rounds of quantum communication, while the amount of classical communication scales as $O(\sqrt{n/\log n})$ where n is the number of processors. The protocols given in (*53*) require many qubits, thus demanding the final stage of a quantum internet. The objective of leader election is to elect a unique leader from a number of distributed processors, which is an important tool for example for deciding which processor gets to employ a particular resource. This tasks is particularly challenging in an anonymous network, in which no node has an identifier. In this setting, there is no exact classical algorithm for leader election for general network topologies, while quantumly leader election is possible (*54*). The protocol proposed in (*54*) requires each end node to process a number of qubits that scales with the number of processors (end nodes). To be used in networks of reasonable size, we thus require a quantum computing network. A number of other leader election protocols have been proposed in a variety of models (*55, 56*).

Finally, this stage allows distributed computational tasks to be solved by transmitting in some cases even exponentially fewer (*57*) qubits than classical bits. A notable example is fingerprinting (*58*). However, these protocols generally require a large number of qubits at each end node to achieve a significant advantage. Specific variants of such protocols with energy constraints can also be realized at lower stages (*59*). Finally, we remark that the presence of entanglement also brings new security issues for existing classical protocols (*60*), requiring new insights and analysis.

Implementation status and challenges

The current status of long-distance quantum networks is at the lowest stage - trusted-repeater networks - with several commercial systems for quantum key distribution on the market. The first extended trusted repeater networks have already been implemented over metropolitan distances (*61–64*) and a long distance implementation has recently been completed (*65*). The hardware required at the lowest stage (mainly light sources, optical links and detectors) has been described in detail in previous literature (see e.g. (*14, 23*)). Realizing the first stage with end-to-end quantum functionality - prepare-and-measure networks - over long distances demands the use of quantum repeaters to bridge long distances via intermediate qubit storage or error correction, as well as routers to forward the quantum state to the desired node. Several recent experiments have demonstrated elements belonging to this and higher stages at short distances suggesting that higher-functionality networks are within reach. To put these experiments into the right perspective, we briefly summarize the main requirements for three types of quantum internet hardware.

Photonic communication channels. Photonic channels establish quantum links between the distant repeater stations and between the end nodes. Two types of photonic channels can be distinguished: free-space channels (potentially via satellites (*66, 67*)) and fiber-based chan-

nels. Each of these has its own advantages and disadvantages, and a future quantum internet - similar to the current classical internet - may employ a combination of them. We require these channels to exhibit minimal photon loss and decoherence. The effect of photon loss on fidelity can in general be dealt with by photon heralding protocols, but photon loss unavoidably affects the communication rate across the network. For photons in the telecom frequency bands loss in fibers can be as low as 0.2 dB/km. Decoherence can in general be overcome through entanglement distillation (*68–70*), which requires additional levels of qubit processing. Finally, the bandwidth of the channels is of practical importance: multiplexing in frequency, time, spatial and/or polarization degrees of freedom allows for increases of the communication rates.

End nodes. For the quantum internet to reach its full potential, the end nodes need to meet the following requirements:

1. Robust storage of quantum states during the time needed to establish entanglement between end nodes. Importantly, this robustness must persist under quantum operations performed on the end node.

2. High-fidelity processing of quantum information within the node. For the more advanced tasks multiple qubits will be required, making the end nodes similar to small-scale quantum computers.

3. Compatibility with photonic communication hardware: efficient interface to light at the relevant wavelength (telecom bands for fiber-based networks).

Several experimental platforms are currently being pursued for the end nodes. Each of these combines well-controlled matter-based qubits with a quantum optical interface via internal electronic transitions. The generation of photon-mediated entanglement between distant matter qubits has been achieved with trapped ions (*71*), atoms (*72, 73*), NV centers in diamond (*74*) and semiconductor quantum dots (*75, 76*) over distances up to 1.3km (*77*). By using

measurement-based schemes with heralding, high-fidelity entangled states could be created in these experiments even though substantial photon loss was present. The major challenge in extending these point-to-point entangled links into true networks is the robust storage of quantum states. The intrinsic coherence times of most above-mentioned platforms are very long (for instance, more than a second for ions and NV centers). However, crosstalk due to unwanted couplings or imperfect individual addressability can severely affect the coherence of a memory qubit under operations on another qubit in the same node (*78, 79*).

A promising approach is to employ different types of qubits within a node. For instance, trapping different species of ions allows for individual addressing of the ions via their different electronic transition frequencies (*80–82*). In a similar fashion, carbon-13 nuclear spins near a diamond NV center provide a robust register of memory qubits that do not interact with the laser control fields on the NV electron spin (*83*). In a very recent experiment, such hybrid network nodes enabled the generation of two remote entangled states on which entanglement distillation could then be performed (*84*). If several of such robust memories can be successfully integrated into a multi-qubit network node, the highest stages of the quantum internet may come into reach.

Another challenge for most of the above systems is that these do not intrinsically couple to light in the telecom band. To fulfil requirement 3, wavelength conversion at the single-photon level can be employed. Pioneering experiments using non-linear optics (see e.g. (*85, 86*)) have already demonstrated the feasibility of such conversion; the current challenge is to realize a robust and high-efficiency (say $> 50\%$) converter that exhibits a high signal-to-noise ratio (say > 100).

As an alternative to the above systems with intrinsic optical interface, the end nodes could be formed from a quantum processor with qubit frequencies in the microwave domain, such as a superconducting qubit circuit, in combination with a microwave-to-optical conversion process.

The physics of such a conversion, for instance by employing mechanical resonators (*87, 88*) or atomic transitions (*89*), is currently being investigated in many labs.

Quantum repeater requirements. Quantum repeater stations need to improve the rate of photonic qubit transfer. The requirements for quantum repeaters are similar to but less strict than for the end nodes. In particular, depending on the exact architecture (see (*90*) for a review), the storage of quantum states may only be required for the time needed to establish entanglement between the nearest active nodes; this storage time can deviate significantly from that required for the end nodes. Also, the qubit processing capabilities required are limited, and therefore systems different from the ones above can be considered. As a prime example, an ensemble of atoms/ions either in gas phase or in a solid can be used as an on-demand quantum memory (*91*). If the memory can herald the arrival of a photon and store the photon's quantum state, photon loss can be efficiently mitigated. Storage and on-demand retrieval have already been achieved (*92–95*), although efficiencies are still to be improved. Such memories also allow for multiplexing within a single device. Furthermore, they are compatible with current-day down-conversion sources for entangled photon pairs. Current challenges are to combine heralding and on-demand high-efficiency retrieval with long coherence times.

It is noteworthy that a radically different approach to quantum repeaters has emerged in recent years in which the quantum state of interest is encoded in multiple photons such that error correction performed at the repeater stations can erase errors due to photon loss and decoherence during transmission (*96–99*). The main advantage of such a scheme is that the classical two-way communication of standard repeater schemes (necessary to convey the heralding signal of whether or not the photons arrived at the stations) becomes obsolete. The communication rates of these schemes are therefore potentially much higher. However, the experimental demands seem daunting at present: for encoding the qubit the near-deterministic generation of a many-photon cluster state is required, which is far beyond the state of the art (*100*). Furthermore,

since these schemes are running quantum error correction schemes, they will only work if the error thresholds associated with the desired transmission qualities are met, thus placing more stringent requirements on the control and readout fidelities within the repeater nodes. That being said, theory research (*101*) in this direction is likely to yield more insights and experimental progress may bring such novel schemes closer to reality in the future.

Finally, we remark that the end nodes that are currently being developed may also function themselves as repeaters.

Network stack requirements. In order to enable wide-spread use and application development, it is essential to develop methods that allow quantum protocols to connect to the underlying hardware implementation transparently, and to make fast and reactive decisions for generating entanglement in the network in order to mitigate limited qubit lifetimes (Figure 8). Classically, this is achieved by a series of layered protocols such as TCP/IP (*102*) that provide an abstraction which ultimately allows application protocols to exchange data between two end nodes without having to know any details on how this connection is actually realized. No such network stack presently exists for a quantum internet, and only some basic elements have been noted (*103*). As a trivial example on why a new stack is required for a quantum network, we note that the first novel feature is a mapping between classical control information (header) and the underlying qubits. In contrast, classically a header and data may be nicely combined in one piece of data to be transmitted. Another example is the use of error detection at the link layer of the classical network stack that does not easily translate to a realistic quantum network. Clearly, error detection can theoretically be realized by using quantum error correcting codes, but this method may be prohibitively expensive in practise and other methods (see e.g. (*104*)) may be more suitable. These are just two simple examples of the challenges involved in designing such a network stack, calling for significant development.

Although it is hard to predict what the exact physical components of a future quantum

internet will be, it is likely that we will see the birth of the first multi-node quantum networks in the next few years. This development brings the exciting opportunity to test all the ideas and functionalities that so far only exist on paper, and may indeed be the dawn of a future large-scale quantum internet.

1. D. Castelvecchi (2018). Https://www.nature.com/articles/d41586-018-01835-3.
2. A. Broadbent, J. Fitzsimons, E. Kashefi, *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on* (IEEE, 2009), pp. 517–526.
3. P. Komar, *et al.*, *Nature Physics* **10**, 582 (2014).
4. D. Gottesman, T. Jennewein, S. Croke, *Physical Review Letters* **109**, 070503 (2012).
5. H. J. Kimble, *Nature* **453**, 1023 (2008).
6. A. I. Lvovsky, B. C. Sanders, W. Tittel, *Nature Photonics* **3**, 706 (2009).
7. T. Northup, R. Blatt, *Nature Photonics* **8**, 356 (2014).
8. W. Gao, A. Imamoglu, H. Bernien, R. Hanson, *Nature Photonics* **9**, 363 (2015).
9. A. Reiserer, G. Rempe, *Reviews of Modern Physics* **87**, 1379 (2015).
10. L. Salvail, *et al.*, *Journal of Computer Security* **18**, 61 (2010).
11. C. H. Bennett, G. Brassard, *International Conference on Computer System and Signal Processing, IEEE, 1984* (1984), pp. 175–179.
12. A. K. Ekert, *Physical Review Letters* **67**, 661 (1991).
13. S. Wiesner, *ACM Sigact News* **15**, 78 (1983).

14. V. Scarani, *et al.*, *Reviews of Modern Physics* **81**, 1301 (2009).

15. E. Biham, B. Huttner, T. Mor, *Physical Review A* **54**, 2651 (1996).

16. S. L. Braunstein, S. Pirandola, *Physical Review Letters* **108**, 130502 (2012).

17. H.-K. Lo, M. Curty, B. Qi, *Physical Review Letters* **108**, 130503 (2012).

18. M. M. Wilde, *Quantum information theory* (Cambridge University Press, 2013).

19. C. Branciard, E. G. Cavalcanti, S. P. Walborn, V. Scarani, H. M. Wiseman, *Physical Review A* **85**, 10301 (2012).

20. A. Gilchrist, N. K. Langford, M. A. Nielsen, *Physical Review A* **71**, 062310 (2005).

21. H.-K. Lo, X. Ma, K. Chen, *Physical review letters* **94**, 230504 (2005).

22. A. Extance, *Fibre Systems* (2017). Https://www.fibre-systems.com/feature/quantum-security.

23. E. Diamanti, H.-K. Lo, B. Qi, Z. Yuan, *npj Quantum Information* **2** (2016).

24. I. Damgaard, S. Fehr, L. Salvail, C. Schaffner, *Theoretical Computer Science* **560**, 12 (2014).

25. F. Dupuis, O. Fawzi, S. Wehner, *IEEE Transactions on Information Theory* **61**, 1093 (2014).

26. D. Mayers, *Physical Review Letters* **78**, 3414 (1997).

27. H.-K. Lo, H. F. Chau, *Physical Review Letters* **78**, 3410 (1997).

28. H.-K. Lo, *Physical Review A* **56**, 1154 (1997).

29. I. Damgaard, S. Fehr, L. Salvail, C. Schaffner, *SIAM Journal on Computing* **37**, 1865 (2000).

30. S. Wehner, C. Schaffner, B. M. Terhal, *Physical Review Letters* **100**, 220502 (2008).

31. R. Koenig, S. Wehner, J. Wullschleger, *IEEE Transactions on Information Theory* **58**, 1962 (2012).

32. N. H. Y. Ng, S. K. Joshi, C. C. Ming, C. Kurtsiefer, S. Wehner, *Nature Communications* **3**, 1326 (2012).

33. J. Ribeiro, F. Grosshans, *arXiv preprint arXiv:1504.07171* (2015).

34. A. Chailloux, I. Kerenidis, *Proceedings of the 52th Annual Symposium on Foundations of Computer Science* (2011).

35. D. Aharonov, A. Ta-Shma, U. V. Vazirani, A. C. Yao, *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing* (2000).

36. A. Acín, *et al.*, *Physical Review Letters* **98**, 230501 (2007).

37. D. Mayers, A. Yao, *Foundations of Computer Science, 1998. Proceedings. 39th Annual Symposium on* (IEEE, 1998), pp. 503–509.

38. J. Barrett, R. Colbeck, A. Kent, *Physical Review Letters* **110**, 010503 (2013).

39. S. Muhammad, *et al.*, *Physical Review X* **4**, 021047 (2014).

40. A. M. Childs, *Quantum Information & Computation* **5**, 456 (2005).

41. D. Aharanov, M. Ben-Or, E. Eban, *Proceedings of Innovations in Computer Science* (2008), pp. 453–469.

42. M. Hillery, V. Bužek, A. Berthiaume, *Physical Review A* **59**, 1829 (1999).

43. R. Cleve, D. Gottesman, H.-K. Lo, *Physical Review Letters* **83**, 648 (1999).

44. C. Crépeau, D. Gottesman, A. Smith, *Proceedings of EUROCRYPT* (2005).

45. M. Christandl, S. Wehner, *International Conference on the Theory and Application of Cryptology and Information Security* (Springer, 2005), pp. 217–235.

46. A. Ambainis, H. Buhrman, Y. Dodis, H. Röhrig, *Proceedings of IEEE Complexity* (2004).

47. N. H. Nickerson, J. F. Fitzsimons, S. C. Benjamin, *Physical Review X* **4**, 041041 (2014).

48. C. Mochon, *arXiv preprint arXiv:0711.4114* (2007).

49. A. Chailloux, I. Kerenidis, *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on* (IEEE, 2009), pp. 527–533.

50. C. Crépeau, D. Gottesman, A. Smith, *Proceedings of the Thiry-fourth Annual ACM Symposium on Theory of Computing*, STOC '02 (ACM, 2002), pp. 643–652.

51. M. Burrows, *Proceedings of the 7th symposium on Operating systems design and implementation* (USENIX Association, 2006), pp. 335–350.

52. V. S. Denchev, G. Pandurangan, *SIGACT News* **39**, 77 (2008).

53. M. Ben-Or, A. Hassidim, *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing* (2005), pp. 481–485.

54. S. Tani, H. Kobayashi, K. Matsumoto, *Proceedings of STACS: 22nd Annual Symposium on Theoretical Aspects of Computer Science* (2005), pp. 581–592.

55. M. Ganz, *Quantum Information Processing* **16**, 1 (2017).

56. N. Aharon, J. Silman, *New Journal of Physics* **12**, 033027 (2010).

57. H. Buhrman, R. Cleve, S. Massar, R. de Wolf, *Reviews of Modern Physics* **82**, 665 (2010).

58. H. Buhrman, R. Cleve, J. Watrous, R. De Wolf, *Physical Review Letters* **87**, 167902 (2001).

59. J. M. Arrazola, N. Lütkenhaus, *Physical Review A* **89**, 062305 (2014).

60. C. Crépeau, L. Salvail, J.-R. Simard, A. Tapp, *Proceedings of ASIACRYPT* (2011), pp. 407–430.

61. M. Peev, *et al.*, *New Journal of Physics* **11**, 075001 (2009).

62. M. Sasaki, *et al.*, *Optics Express* **19**, 10387 (2011).

63. D. Stucki, *et al.*, *New Journal of Physics* **13**, 123001 (2011).

64. S. Wang, *et al.*, *Optics Express* **22**, 21739 (2014).

65. R. Courtland, *IEEE Spectrum* **53**, 11 (2016).

66. G. Vallone, *et al.*, *Physical Review Letters* **115**, 040502 (2015).

67. J. Yin, *et al.*, *Science* **356**, 1140 (2017).

68. C. H. Bennett, *et al.*, *Physical Review Letters* **76**, 722 (1996).

69. D. Deutsch, *et al.*, *Physical Review Letters* **77**, 2818 (1996).

70. W. Dür, H.-J. Briegel, J. Cirac, P. Zoller, *Physical Review A* **59**, 169 (1999).

71. D. Moehring, *et al.*, *Nature* **449**, 68 (2007).

72. S. Ritter, *et al.*, *Nature* **484**, 195 (2012).

73. J. Hofmann, *et al.*, *Science* **337**, 72 (2012).

74. H. Bernien, *et al.*, *Nature* **497**, 86 (2013).

75. A. Delteil, *et al.*, *Nature Physics* **12**, 218 (2016).

76. R. Stockill, *et al.*, *Physical Review Letters* **119**, 010503 (2017).

77. B. Hensen, *et al.*, *Nature* **526**, 682 (2015).

78. D. Hucul, *et al.*, *Nature Physics* **11**, 37 (2015).

79. W. Pfaff, *et al.*, *Science* **345**, 532 (2014).

80. C. Ballance, *et al.*, *Nature* **528**, 384 (2015).

81. T. Tan, *et al.*, *Nature* **528**, 380 (2015).

82. I. Inlek, C. Crocker, M. Lichtman, K. Sosnova, C. Monroe, *Physical Review Letters* **118**, 250502 (2017).

83. A. Reiserer, *et al.*, *Physical Review X* **6**, 021040 (2016).

84. N. Kalb, *et al.*, *Science* **356**, 928 (2017).

85. S. Tanzilli, *et al.*, *Nature* **437**, 116 (2005).

86. S. Zaske, *et al.*, *Physical Review Letters* **109**, 147404 (2012).

87. R. W. Andrews, *et al.*, *Nature Physics* **10**, 321 (2014).

88. J. Bochmann, A. Vainsencher, D. Awschalom, A. N. Cleland, *Nature Physics* **9**, 712 (2013).

89. S. Probst, *et al.*, *Physical Review Letters* **110**, 157001 (2013).

90. W. J. Munro, K. Azuma, K. Tamaki, K. Nemoto, *IEEE Journal of Selected Topics in Quantum Electronics* **21** (2015).

91. N. Sangouard, C. Simon, H. De Riedmatten, N. Gisin, *Reviews of Modern Physics* **83**, 33 (2011).

92. N. Kalb, A. Reiserer, S. Ritter, G. Rempe, *Physical Review Letters* **114**, 220501 (2015).

93. C. Kurz, *et al.*, *Nature Communications* **5**, 5527 (2014).

94. H. Tanji, S. Ghosh, J. Simon, B. Bloom, V. Vuletić, *Physical Review Letters* **103**, 043601 (2009).

95. A. Delteil, Z. Sun, S. Fält, A. Imamoğlu, *Physical Review Letters* **118**, 177401 (2017).

96. K. Azuma, K. Tamaki, H.-K. Lo, *Nature Communications* **6** (2015).

97. M. Pant, H. Krovi, D. Englund, S. Guha, *Physical Review A* **95**, 012304 (2017).

98. S. Muralidharan, J. Kim, N. Lütkenhaus, M. D. Lukin, L. Jiang, *Physical Review Letters* **112**, 250501 (2014).

99. W. J. Munro, A. M. Stephens, S. J. Devitt, K. A. Harrison, K. Nemoto, *Nature Photonics* **6**, 777 (2012).

100. I. Schwartz, *et al.*, *Science* **354**, 434 (2016).

101. T. Rudolph, *APL Photonics* **2**, 030901 (2017).

102. V. G. Cerf, R. E. Kahn, *IEEE Transactions on Communications* **22**, 637 (1974).

103. R. Van Meter, J. Touch, *IEEE Communications Magazine* **51**, 64 (2013).

104. C. Pfister, M. A. Rol, A. Mantri, M. Tomamichel, S. Wehner, *Nature Communications* **9** (2018).

105. M. Takeoka, S. Guha, M. Wilde, *Nature Communications* **5**, 5235 (2014).

106. S. Pirandola, R. Laurenza, C. Ottaviani, L. Banchi, *Nature Communications* **8** (2017).

107. L.-M. Duan, M. Lukin, J. I. Cirac, P. Zoller, *Nature* **414**, 413 (2001).

108. C. Simon, *et al.*, *Physical Review Letters* **98**, 190503 (2007).

109. L. Jiang, *et al.*, *Physical Review A* **79**, 032325 (2009).

110. J. J. Wallman, S. T. Flammia, *New Journal of Physics* **16**, 103032 (2014).

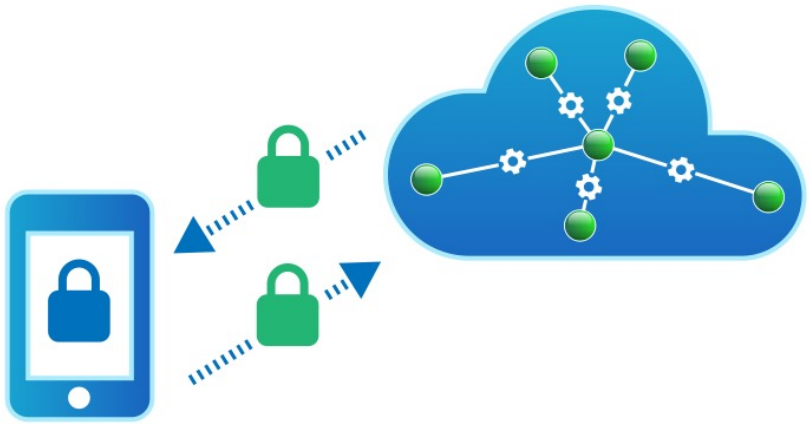


Figure 1: One application of a quantum internet is to allow secure acccess to remote quantum computers in the cloud (*2*). Specifically, a simple quantum terminal capable of preparing and measuring only single qubits can use a quantum internet to access a remote quantum computer in such a way that the quantum computer can learn nothing about which computation it has performed. Almost all other applications of a quantum internet can be understood from two special features of quantum entanglement. First, if two qubits at different network nodes are entangled with each other, then such entanglement enables stronger than classical correlation and coordination. For example, for any measurement on qubit one, if we made the same measurement on qubit two, then we instantaneously obtain the same answer, even though that answer is random and was not determined ahead of time. Very roughly, it is this feature that makes entanglement so well suited for tasks that require coordination. Examples include clock synchronization (*3*), leader election and achieving consensus about data (*52*), or even using entanglement to help two online bridge players coordinate their actions (*39*). The second feature of quantum entanglement is that it cannot be shared. If two qubits are maximally entangled with each other, then it is impossible by the laws of quantum mechanics for a third qubit to be just as entangled with any of them. This makes entanglement inherently private, bringing great advantages to tasks that require security such as, for example, generating encryption keys (*12*) or secure identification (*24, 25*). .

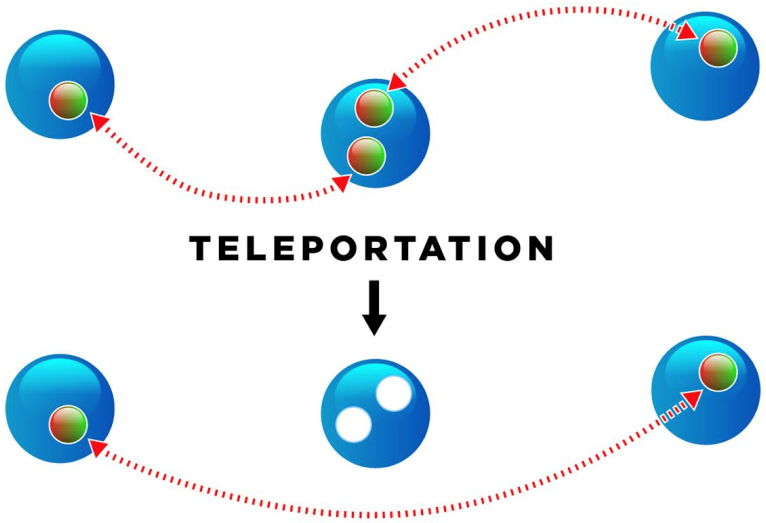


Figure 2: Quantum repeaters work fundamentally different from classical repeaters. In its simplest form, a quantum repeater works by first generating entanglement (dashed line) between the repeater (middle) and each of the end nodes (left and right) individually. Intuitively, this can be done because the distance of each end point to the repeater is still sufficiently small to allow direct entanglement generation by transmitting photons over telecom fiber. Subsequently, the repeater teleports one of the qubits entangled with node one onto node two. This procedure is known as entanglement swapping, and allows the creation of entanglement over distances where direct transmission is infeasible. After establishing long distance entanglement, a data qubit may now be sent using quantum teleportation.

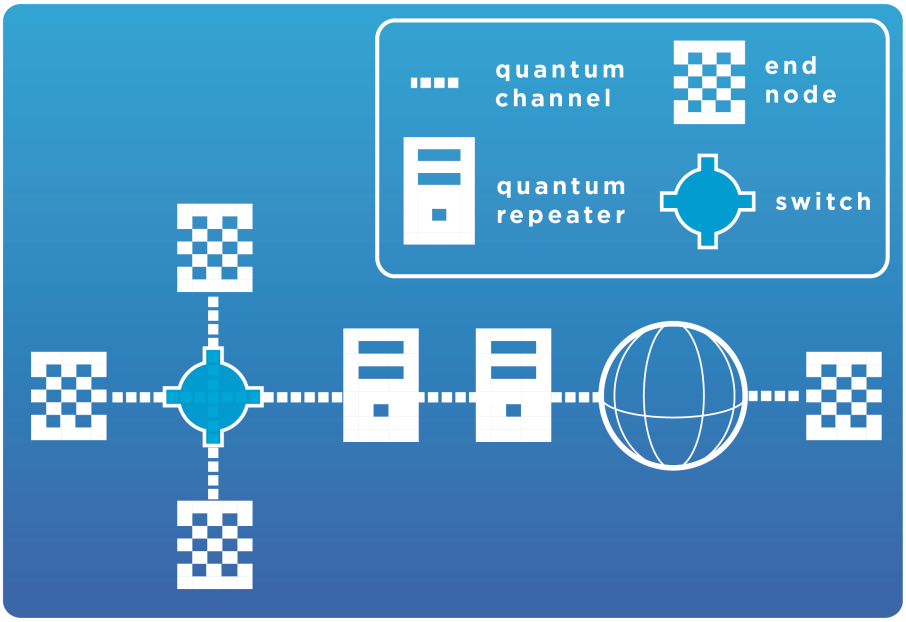


Figure 3: A quantum internet consists of three essential quantum hardware elements. First, we need a physical connection (quantum channel) that supports the transmission of qubits. Examples are standard telecom fibers as they are presently used to communicate classical light. Second, we need a means to extend these short distances. Quantum channels are inherently lossy. For instance, the transmissivity of fiber optical channels scale exponentially with the distance. This scaling has strong implications for applications. For instance, both for entanglement and key distribution, the achievable rates can at most be proportional to the transmissivity (*105, 106*). Hence, in order to reach longer distances, intermediate nodes called *quantum repeaters* are necessary (see e.g. (*96, 107–109*), and for reviews (*90, 91*)). Such a repeater is placed at certain intervals along the optical fiber connection, in theory allowing qubits to be transmitted over arbitrarily long distances. In the future, powerful repeaters may also double as long-distance routers in a quantum network. The final element are the *end nodes*, that is, the quantum processors connected to the quantum internet. These may range from extremely simple nodes that can only prepare and measure single qubits, to large scale quantum computers. End nodes may also act as quantum repeaters themselves, although this is not a requirement. A quantum internet is not meant to replace classical communication, but rather to supplement it with quantum communication. We hence assume all nodes can communicate classically, for example over the classical internet, in order to exchange control information.

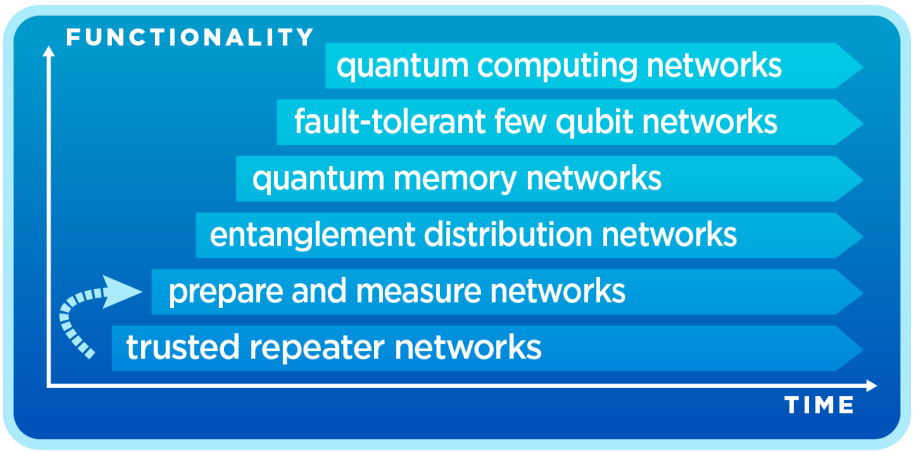


Figure 4: Stages of quantum internet development. A specific implementation of a quantum internet may, as a classical network, be optimized for distance, functionality, or both. The term network commonly refers to a situation that goes beyond point-to-point communication: the objective of a network is to provide any end nodes (connected to the network) with the means to exchange data, making three end nodes the smallest instance of a true network. Outside the lab, only trusted repeater networks have been realized in metropolitan areas (*61–64*). Two single far away end nodes (*65*) have also been connected via satellite.

Stage	Additional Functionality	Protocol Parameters
Prepare and Measure	For any two end nodes i, j, any one qubit state $\|\Psi\rangle$ and any one qubit projective measurement M, there exists a way for i to prepare $\|\Psi\rangle$, transfer it to j, such that either [1] j performs measurement M on $\|\Psi\rangle$ or [2] j concludes the qubit was lost.	Distances ε_T and ε_M from the ideal transmission and measurement operations (Figure 7). Probability p that the state is not lost.
Entanglement Distribution	For any two end nodes i, j: • the network allows the heralded creation of a maximally entangled state $\|\Phi_{ij}\rangle$ • nodes i and j can deterministically perform any single qubit measurements M_i and M_j	Distances ε_P from the ideal preparation, and ε_M from the idealized measurement (Figure 7).
Quantum Memory	For any two end nodes i, j, the network allows the execution entanglement generation and the following additional tasks in any order: • Preparation of a one qubit ancilla state $\|\Psi\rangle$ by end node i or j • Measurements of any subset of the qubits at node • Application of an arbitrary unitary U at node • Storage of the qubits for a minimum time $k \cdot \ell_m \cdot t$, where t is defined as the time that is required to generate one EPR pair and send a classical message from node i to j maximized over all pairs of nodes, and ℓ_m is the time that it takes for the execution of a depth m quantum circuit at the end node.	Number of rounds k. Circuit depth m. Number of physical qubits q. For each of the operations, an estimate ε_j from the ideal operation (Figure 7).
Few qubit fault-tolerant	Fault-tolerant execution of a universal gate set on q logical qubits, where $q \geq 1$ is small enough such that the local processor can efficiently be simulated on a classical computer.	Number of logical qubits q.
Quantum Computing	q is larger than the number of qubits that can effectively be simulated on a classical computer.	Number of logical qubits q.

Figure 5: Formal definitions of the stages, and relevant parameters for protocol design. Higher stages include all functionality available at the previous ones. It is an open question to determine necessary and sufficient conditions for these parameters to realize general protocols.

Stage	Examples of known protocols
Quantum computing	Leader election, Fast byzantine agreement, Weak coin flipping with arbitrarily small bias
Fault-tolerant few qubit	Clock synchronization, Distributed quantum computation
Quantum Memory	Blind quantum computing (using remote quantum servers), Improved coin flipping, Anonymous quantum transmissions, Extending baseline of telescopes, Secret sharing, Simple leader election and agreement protocols, Time limited clock synchronization
Entanglement Distribution	Device independence for QKD and other protocols in the prepare and measure stage
Prepare and Measure	Quantum key distribution (QKD), Two-party cryptography, Position verification, Imperfect coin flipping

Figure 6: Examples of known protocols and their requirements. It is a challenge for quantum network programmers to find protocols for the same tasks, that can be realized with lower stages of a quantum internet. It is an interesting open question what minimum stage is required in order to realize a specific task.

Performance of quantum internet protocols. A general quantum internet protocol is comprised of a series of operations consisting of state preparation, transmission, unitary operations and measurements. In reality, each of these operations is noisy, so instead of executing a sequence of ℓ ideal operations $\mathcal{I} = \mathcal{I}_\ell \circ \ldots \circ \mathcal{I}_1$, we are executing the real (noisy) protocol $\mathcal{R} = \mathcal{R}_\ell \circ \ldots \circ \mathcal{R}_1$. To assess the performance of the real protocol execution, it is sufficient to estimate the diamond norm distance (*20*)

$$D_\diamond(\mathcal{R}, \mathcal{I}) = \max_{\rho_{SE}} D(\mathcal{R} \otimes \mathsf{id}_E(\rho_{SE}), \mathcal{I} \otimes \mathsf{id}_E(\rho_{SE})) \ ,$$

where $D(\tau, \sigma)$ is the well known trace distance (*18*) which determines how well two states τ and σ can be distinguished by *any* physical process, and S denotes the system that the protocol acts on which may be part of a larger system SE. Since $D_\diamond$ is (unlike the fidelity) a metric, it is straightforward to show that having estimated individual errors $\|\mathcal{R}_j - \mathcal{I}_j\|_\diamond \leq \varepsilon$ allows an estimate of the overall error as

$$D_\diamond(\mathcal{R}, \mathcal{I}) \leq \ell \cdot \varepsilon \ .$$

For unitary operations and projective measurements, the diamond norm distance is directly related to the average gate fidelity (*110*). We remark that if the ideal operation $\mathcal{I}(\rho) = \Phi$ simply aims to prepare a state Φ, and the real operation prepares $\mathcal{R}(\rho) = \tilde{\Phi}$ then the diamond norm distance satisfies $D_\diamond(\mathcal{R}, \mathcal{I}) \leq \sqrt{1 - F(\Phi, \tilde{\Phi})}$, where F is the fidelity. Evidently, the end-user —who desires to run application protocols— should be able to perform tests that give confidence for any possible operation instead of having to test the exact unitaries and measurements in any conceivable protocol.

Figure 7: Performance measures of a quantum internet. It is an interesting open question to devise integrated tests which are more refined then estimating the errors ε of the individual operations.

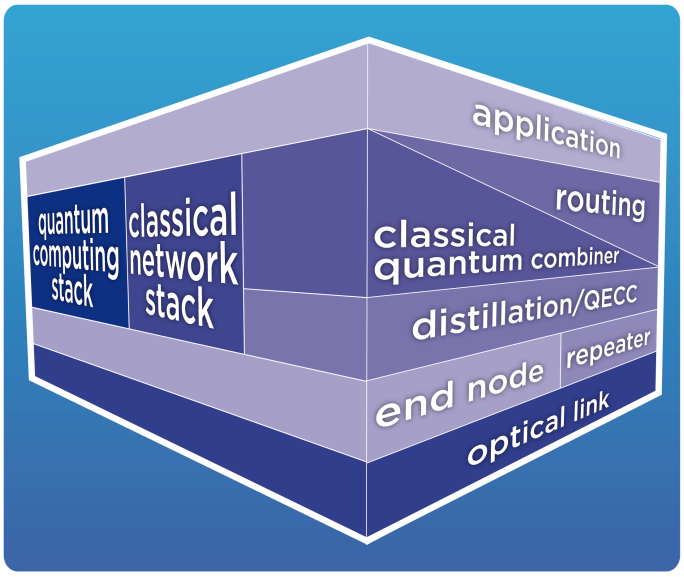


Figure 8: Essential elements of a future quantum network stack.