Quantum Internet Research Group
Internet-Draft
Intended status: Informational
Expires: September 10, 2019

W. Kozlowski
QuTech

March 9, 2019

Architectural Principles for a Quantum Internet
draft-irtf-qirg-principles-00

Abstract

The vision of a quantum internet is to fundamentally enhance Internet technology by enabling quantum communication between any two points on Earth. To achieve this goal, a quantum network stack must be built from the ground up as the physical nature of the communication is fundamentally different. The first realisations of quantum networks are imminent, but there is no practical proposal for how to organise, utilise, and manage such networks. In this memo, we attempt lay down the framework and introduce some basic architectural principles for a quantum internet. This is intended for general guidance and general interest, but also to provide a foundation for discussion between physicists and network specialists.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 10, 2019.

Table of Contents

1. Introduction

Quantum networks are distributed systems of quantum computers that utilise fundamental quantum mechanical phenomena such as superposition, entanglement, and quantum measurement to achieve capabilities beyond what is possible with classical networks. This new networking paradigm offers promise for a range of new applications such as tamper-proof communications [1], distributed

quantum computation [2], and quantum sensor networks [3]. The field of quantum communication has been a subject of active research for many years and the most well-known application of quantum computers that has already been deployed, quantum key distribution (QKD), is a protocol used for secure communications.

Fully quantum networks capable of transmitting and managing entangled states in order to send, receive, and manipulate distributed quantum states are now imminent [4] [5]. Whilst a lot of effort has gone into physically connecting the devices and bringing down the error rates there are no concrete proposals for how to run these networks. To draw an analogy with a classical network, we are at a stage where we can physically connect our devices and send data, but all sending, receiving, buffer management, connection synchronisation, and so on, must be managed by the application itself at what is essentially assembly level. Furthermore, whilst physical mechanisms for forwarding quantum states exist, there are no protocols for managing it.

2. Model of computation

In order to understand the framework for quantum networking a basic understanding of quantum information is necessary. The following sections aim to introduce the bare minimum necessary to be understand the principles of operation of a quantum network. This exposition was written with a classical networking audience in mind. It is assumed that the reader has never before been exposed to any quantum physics.

2.1. Qubit

The differences between quantum computation and classical computation begin at the bit-level. A classical computer operates on the binary alphabet { 0, 1 }. A quantum bit, a qubit, exists over the same binary space, but unlike the classical bit, it can exist in a so-called superposition of the two possibilities:

$$a\,|0\rangle + b\,|1\rangle,$$

where |X> denotes a quantum state, here the binary 0 and 1, and the coefficients a and b are complex numbers called probability amplitudes. Physically, such a state can be realised using a variety of different technologies such as electron spin, photon polarisation, atomic energy levels, and so on.

Upon measurement, the qubit loses its superposition and irreversibly collapses into one of the two basis states, either |0> or |1>. Which of the two states it ends up in is not deterministic. The

probability of measuring the state in the |0> state is |a|^2 and similarly the probability of measuring the state in the |1> state is |b|^2. This randomness is not due to our ignorance of the underlying mechanisms, but rather it is a fundamental feature of a quantum mechanical system [6].

The superposition property plays an important role in fundamental gate operations on qubits. Since a qubit can exist in a superposition of its basis states, the elementary quantum gates are able to act on all states of the superposition at the same time. For example, consider the NOT gate:

NOT (a |0> + b |1>) -> a |1> + b |0>.

2.2. Multiple qubits

When multiple qubits are combined in a single quantum state the space of possible states grows exponentially and all these states can coexist in a superposition. For example, the general form of a two qubit register is

a |00> + b |01> + c |10> + d |11>

where the coefficients have the same probability amplitude interpretation as for the single qubit state. Each state represents a possible outcome of a measurement of the two qubit register. For example, |01>, denotes a state in which the first qubit is in the state |0> and the second is in the state |1>.

Performing single qubit gates affects the relevant qubit in each of the superposition states. Similarly, two qubit gates also act on all the relevant superposition states, but their outcome is far more interesting.

Consider a two qubit register where the first qubit is in the superposed state (|0> + |1>)/sqrt(2) and the other is in the state |0>. This combined state can be written as:

(|0> + |1>)/sqrt(2) x |0> = (|00> + |10>)/sqrt(2),

where x denotes a tensor product (the mathematical mechanism for combining quantum states together). Let us now consider the two-qubit CNOT gate. The CNOT gate takes as input two qubits, a control and target, and applies the NOT gate to the target if the control qubit is set. The truth table looks like

IN	OUT
00	00
01	01
10	11
11	10

Now, consider performing a CNOT gate on the ensemble with the first qubit being the control. We apply a two qubit gate on all the superposition states:

CNOT (|00> + |10>)/sqrt(2) -> (|00> + |11>)/sqrt(2).

What is so interesting about this two-qubit gate operation? The final state is *entangled*. There is no possible way of representing that quantum state as a product of two individual qubits, they are no longer independent and their behaviour cannot be fully described without accounting for the other qubit. The states of the two individual qubits are now correlated beyond what is possible to achieve classically. Neither qubit is in a definite |0> or |1> state, but if we perform a measurement on either one, the outcome of the partner qubit will *always* yield the exact same outcome. The final state, whether it's |00> or |11>, is fundamentally random as before, but the states of the two qubits following a measurement will always be identical.

Once a measurement is performed, the two qubits are once again independent. The final state is either |00> or |11> and both of these states can be trivially decomposed into a product of two individual qubits. The entanglement has been consumed and if the same measurement is to be repeated, the entangled state must be prepared again.

3. Entanglement as the fundamental service

Entanglement is the fundamental building block of quantum networks. To see this, consider the final state from the previous section:

(|00> + |11>)/sqrt(2).

Neither of the two qubits is in a definite |0> or |1> state and we need to know the state of the entire register to be able to fully describe the behaviour of the two qubits.

Now consider sending one of the qubits to another device. This device can be anywhere: on the other side of the room, in a different

country, or even on a different planet. Provided negligible noise has been introduced, the two qubits will forever remain in the entangled state until a measurement is performed. The physical distance does not matter at all for entanglement.

This lies at the heart of quantum networking, because it is possible to leverage these non-classical correlations in order to design completely new types of algorithms that are not possible to achieve with just classical communication. Examples of such applications are quantum cryptography, blind quantum computation, or distributed quantum computation.

As a trivial example consider the problem of reaching consensus between two nodes. The two nodes want to agree on the value of a single bit. In a quantum network they can simply request the network to generate the state (|00> + |11>)/sqrt(2) for them and that is essentially all that needs to be done. Once any of the two nodes performs a measurement the state of the two qubits collapses to either |00> or |11> so whilst the outcome is random, the two nodes will always measure the same value. We can also build the more general multi-qubit state (|00...> + |11...>)/sqrt(2) and perform the same algorithm between an arbitrary number of nodes.

However, it is impossible to entangle two qubits without ever having them directly interact with each other (e.g. by performing a local two-qubit gate, such as the CNOT). A local interaction is necessary to create entanglement and thus such states cannot be created between two quantum computers that cannot transmit quantum states to each other. Therefore, it is the entanglement property of multi-qubit states that draws the line between a genuine quantum network and a collection of quantum computers connected over a classical network.

A quantum network is defined as a collection of nodes that is able to distribute entangled states amongst themselves. A quantum computer that is able to communicate classically with another quantum computer is not a member of a quantum network.

This is a crucial difference between classical and quantum networks. Classical applications transmit data over the network to synchronise distributed state. Quantum network applications obtain distributed states, synchronised at the physical level via entanglement, from the network to perform quantum algorithms.

More complex services and applications can be built on top of entangled states distributed by the network.

4. Achieving quantum connectivity

4.1. No-cloning theorem

To build a network we must first physically connect all the nodes with quantum channels that enable them to distribute the entanglement. Unfortunately, our ability to transfer quantum states is complicated by the no-cloning theorem.

The no-cloning theorem states that it is impossible to create an identical copy of an arbitrary unknown quantum state. Since performing a measurement on a quantum state destroys its superposition, there is no practical way of learning the exact state of a qubit in an unknown state. Therefore, it is impossible to use the same mechanisms that worked for classical networks for error-correction, amplification, retransmission, and so on as they all rely on the ability to copy the underlying data. Since any physical channel will always be lossy, connecting a quantum network is a challenging endeavour and its architecture must at its core address this very issue.

4.2. Direct transmission

The most straightforward way to distribute an entangled state is to simply transmit one of the qubits directly to the other end across a series of nodes while performing sufficient error correction to bring losses down to an acceptable level. Despite the no-cloning theorem and the inability to directly measure a quantum state error-correcting mechanisms for quantum communication exist [7]. However, even in the most optimistic scenarios the hardware requirements to fault-tolerantly transmit a single qubit are beyond near-term capabilities. Nevertheless, due to the promise of fault-tolerance and its favourable poly-logarithmic scaling with distance, this may eventually become a desirable method for entanglement distribution.

4.3. Bell pairs and entanglement swapping

4.3.1. Bell Pairs

An alternative relies on the observation that we do not need to be able to distribute any arbitrary entangled quantum state. We only need to be able to distribute any one of what are known as the Bell Pair states. Bell Pair states are the entangled two-qubit states:

|00> + |11>,
|00> - |11>,
|01> + |10>,
|01> - |10>,

where the constant 1/sqrt(2) normalisation factor has been ignored for clarity. Any of the four Bell Pair state above will do as it is possible to transform any Bell Pair into another Bell Pair with local operations performed on only one of the qubits. That is, either of the nodes that hold the two qubits of the Bell Pair can apply a series of single qubit gates to just their qubit in order to transform the ensemble between the different variants.

Distributing a Bell Pair between two nodes is much easier than transmitting an arbitrary quantum state over a network. Since the state is known error-correction is easier and error-detection combined with reattempts becomes a valid strategy.

The reason for using Bell Pairs specifically as opposed to any other two-qubit state, is that they are the maximally entangled two-qubit set of basis states. Maximal entanglement means that these states have the strongest non-classical correlations of all possible two-qubit states. Furthermore, since single-qubit local operations can never increase entanglement, less entangled states would impose some constraints on distributed quantum algorithms. This makes Bell Pairs particularly useful as a generic building block for distributed quantum applications.

4.3.2. Teleportation

The observation that we only need to be able to distribute Bell Pairs relies on the fact that this enables the distribution of any other arbitrary entangled state. This can be achieved via quantum state teleportation. Quantum state teleportation consumes an unknown quantum state that we want to transmit and recreates it at the desired destination.

To achieve this, a Bell Pair needs to be distributed between the source and destination. The source then entangles the transmission qubit with its end of the Bell Pair and performs a measurement. This consumes the Bell Pair's entanglement turning the source and destination qubits into independent states. However, this process transforms the Bell Pair's qubit at the destination into the transmission qubit's original state. Note he process requires the source to also communicate its two-bit measurement result so that the destination can correct for the randomness of the outcome.

The unknown quantum state that was transmitted never entered the network itself. Therefore, the network needs to only be able to reliably produce Bell Pairs between any two nodes in the network.

4.3.3. Bell Pair links and entanglement swapping

Reducing the problem to one of generating a Bell Pair state has facilitated the problem, but it has not solved it.

The technology to generate a Bell Pair between two directly connected quantum nodes already exists and has been demonstrated in laboratory conditions [8]. Interestingly, neither of the two qubits of the pair need to be transmitted any further.

A Bell Pair between any two nodes in the network can be constructed from Bell Pairs generated along each individual link on the path between the two end-points. Each node along the path can consume the two Bell Pairs on the two links that it is connected to in order to produce a new Bell Pair between the two far ends. This process is known as entanglement swapping. Pictorially it can be represented as follows:

```
 x~~~~~~~~~~~~~x x~~~~~~~~~~~~~x
[ ]-----------[   ]-----------[ ]
```

where x~~x denotes a Bell Pair with individual qubits represented by x, -- denotes a quantum link, and [ ] denotes a node. The diagram above represents the situation after the middle node has generated a Bell Pair with two of its directly connected neighbours. Now, the middle node performs an entanglement swap operation (the exact details of the mechanism are beyond the scope of this memo). This operation consumes the two Bell Pairs and produces a new Bell Pair between the two far ends of this three-node network as follows:

```
 x~~~~~~~~~~~~~~~~~~~~~~~~~~~~~x
[ ]-----------[   ]-----------[ ]
```

The outcome is guaranteed to be a Bell Pair between the two end nodes, but which of the four possible Bell Pairs is produced is not deterministic. However, the middle node will know which one was produced as the entanglement swap is a measurement operation that yields two classical bits. The final state can be inferred from this two-bit readout. Therefore, the middle node needs only to communicate the outcome over a classical channel to one or both ends who can apply a correction to transform the pair into any of its other forms (if so desired).

4.3.4. Distillation

Neither the Bell Pair or the swapping operations are lossless operations. Therefore, with each link and each swap the quality of the state degrades. However, it is possible to create higher quality

Bell Pair states from two or more lower quality Bell Pair states. Therefore, once the quality loss over a given distance become prohibitive, additional redundancy may be used to restore the state quality.

4.4. Direct transmission vs. swapping

Direct state transmission whilst simpler conceptually is much more demanding to implement reliably in practice which means that any near-term practical realisation is more likely to succeed if it is based on the Bell Pair and entanglement swapping architecture. This is the architecture that we will focus on in the rest of this memo for practical reasons.

Nevertheless, we are not entirely discarding the direct transmission proposal. Whilst it does enable the fault-tolerant transmission of unknown quantum states, it might still be more beneficial to use it to distribute Bell Pairs instead. Distributing Bell Pairs via direct transmission means that one can leverage the advantages of entanglement swapping which allows for paralellisation as the Bell Pairs can be built up from both ends simultaneously. Furthermore, the generic nature of the Bell Pair means that a network may provision resources better before it receives any request.

5. Architecture of a quantum internet

5.1. Model of a quantum network

A generic quantum network of three nodes could be represented as

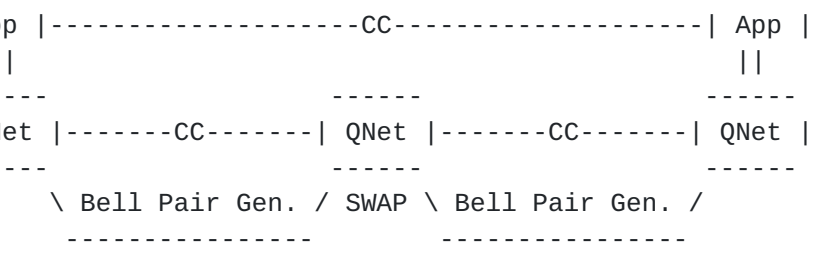

Where "App" is some application running over a quantum network, --CC-- denote classical communication links (e.g. over the public Internet or a private LAN), and "QNet" is a generic network stack. Architectures for the network stack have been proposed already [9], but their discussion is beyond the scope of this memo. However, they all map onto this generic diagram. Nodes within a quantum network that are capable of performing the entanglement swap operation are often referred to as quantum repeaters and we shall adopt this terminology from this point on. End-hosts connecting at the edge of the network are not necessarily repeaters themselves.

The key message here is that a network stack relies on the hardware being able to provide two services: Bell Pair generation across a link, and swap operation. In any network model it is assumed that the physical device is capable of providing both of these services and offers a suitable interface for their usage.

Strictly speaking quantum memories are not needed for a functional quantum network as long as the network is able to generate the Bell Pairs, swap the entanglement, and deliver the final Bell Pair to the application in a usable form. However, in general, to be able to provide the two services above, the hardware will also need to be able to store the qubits in memory which is highly non-trivial.

Furthermore, it is also assumed that the applications are able to communicate classically, and that the nodes themselves are also connected over some classical channel. The classical links between the nodes need not always have an associated quantum link, but it is assumed that any quantum link has a classical link running in parallel.

5.2. Physical constraints

The model above has effectively abstracted away the particulars of the hardware implementation. However, certain physical constraints need to be considered in order to build a practical network. Some of these are fundamental constraints and no matter how much the technology improves, they will always need to be addressed. Others are artefacts of the early stages of a new technology.

5.2.1. Fidelity

The quality of a quantum state is described by a physical quantity called fidelity. Fidelity is the measure of how close a quantum state is to the quantum state we desire it to be in. It expresses the probability that one state will pass a test to identify as the other.

Fidelity is an important property of a quantum system that stems from the fact that no physical operation is perfect. Furthermore, applications will in general require the fidelity of a quantum state to be above some minimum threshold in order to guarantee the correctness of their algorithm and it is the responsibility of the network to provide such a state.

Additionally, entanglement swap operations, even if perfect, lead to a further reduction in the fidelity of the final state. Two imperfect Bell Pairs when combined will produce a slightly worse Bell Pair. Whilst distillation is one of the available mechanisms to

correct for these errors it requires additional Bell Pairs to be produced. There will be a trade-off between how much distillation is to be done versus what fidelity is acceptable.

This is a fundamental constraint as perfect noiseless operations and lossless communication channels are unachievable. Therefore, no Bell Pair will be generated with perfect fidelity and the network must account for this.

5.2.2. Memory lifetimes

In addition to discrete operations being imperfect, storing a qubit in memory is also highly non-trivial. The main difficulty in achieving persistent storage is that it's extremely challenging to isolate a quantum system from the environment. The environment introduces an uncontrollable source of noise into the system which affects the fidelity of the state. This process is known as decoherence. Eventually, the state has to be discarded once its fidelity degrades too much.

The memory lifetime depends on the particular physical setup, but the highest achievable values currently are on the order of hundreds of milliseconds. These values have increased tremendously over the lifetime of the different technologies and are bound to keep increasing. However, if quantum networks are to be realised in the near future, they need to be able to handle short memory lifetimes. An architecture that handles short lifetimes may also be more cost-efficient in the future.

5.2.3. Rates

Entanglement generation on a link between two connected nodes is not a very efficient process and it requires many attempts to succeed. A fast repetition rate for Bell Pair generation is achievable, but only one in a few thousands will succeed. Currently, the highest achievable rates of success are of the order of 10 Hz. Combined with short memory lifetimes this leads to very tight timing windows to build up network-wide connectivity. Achievable rates are likely to increase with time, but just like with quantum memories, it may be more cost-efficient in the future to provide low-rate links in some parts of the network.

5.2.4. Communication qubit

Some physical architectures are not able to generate entanglement using any memory qubit that they have access to. In these systems, entanglement is generated using a communication qubit and once a Bell Pair has been generated, the qubit state is transferred into memory.

This may impose additional limitations on the network. In particular if a given node has only one communication qubit it cannot simultaneously generate Bell Pairs over two links. It must generate entanglement over the links one at a time.

5.2.5. Homogeneity

Currently all hardware implementations are homogeneous and they do not interface with each other. In general, it is very challenging to combine different quantum information processing technologies due to their sensitivity to losses. Coupling different technologies with each other is of great interest as it may help overcome the weaknesses of the different implementations, but this is not a near-term goal.

5.3. Architectural principles

Given that the most practical way of realising quantum network connectivity is using Bell Pair and entanglement swapping repeater technology what sort of principles should guide us in assembling such networks such that they are functional, robust, efficient, and most importantly: they work. Furthermore, how do we design networks so that they work under the constraints imposed by the hardware available today, but do not impose unnecessary burden on future technology. Redeploying network technology is a non-trivial process.

As this is a completely new technology that is likely to see many iterations over its lifetime, this memo must not serve as a definitive set of rules, but merely as a general guide based on principles and observations made by the community. The benefit of having a community built document at this early stage is that expertise in both quantum information and network architecture is needed in order to successfully build a quantum internet.

5.3.1. Goals of a quantum internet

When outlining any set of principles we must ask ourselves what goals do we want to achieve as inevitably trade-offs must be made. So what sort of goals should drive a quantum network architecture? The following list has been inspired by the history of the classical Internet, but it will inevitably evolve with time and the needs of its users. The goals are listed in order of priority which in itself may also evolve as the community learns more about the technology.

1. Support distributed quantum applications

   The primary purpose of a quantum internet is to run distributed quantum algorithms and it is of utmost importance that they can

run well and efficiently. Therefore, the needs of quantum applications should always be considered first.

If a network is able to distribute entanglement it is officially quantum. However, if it is unable to distribute these states with a sufficiently high fidelity at a reasonable rate for a majority of potential applications it is not practical.

2. Support tomorrow's distributed quantum applications

There are many applications already proposed to run over a quantum internet. However, more algorithms will be invented as the community grows as well as the robustness and the reliability of the technology. Any proposed architecture should not constrain the capabilities of the network for short-term benefit.

3. Hardware heterogeneity

There are multiple proposals for realising practical quantum repeaters and they all have their advantages and disadvantages. It is also very likely that the most optimal technologies in the future will be hybrid combinations of the many different solutions currently under development. It should be an explicit goal of the architecture to allow for a large variety of hardware implementations.

4. Be flexible with regards to hardware capabilities and limitations

This goal encompasses two important points. First, the architecture should be able to function under the physical constraints imposed by the current generation hardware. Second, it should not make it difficult to run the network over any hardware that may come along in the future. The physical capabilities of repeaters will improve and redeploying a technology is extremely challenging.

5. Security, availability, and resilience

Whilst the priority for the first quantum networks should be to simply work, we cannot forget that ultimately they have to also be secure. There are three key security considerations at the network level, confidentiality, integrity, and authenticity.

Confidentiality and integrity - it is vital that the network can provide a reasonable guarantee of the minimum fidelity of a delivered Bell Pair as the application's own security mechanisms rely on this. Uncertainty about the fidelity of a Bell Pair may potentially expose its data to an eavesdropper.

Authenticity - it is important that any application can have confidence that the other end of the Bell Pair has been delivered to the desired partner.

Additionally a practical and usable network is able to continue to operate despite losses and failures, and will be robust to malicious actors trying to disable connectivity. These may be simply considered different aspects of security, but it is worthwhile to address them explicitly at the architectural level already.

6. Easy to manage and monitor

Quantum networks rely on complex physical phenomena and require hardware that is challenging to build. Furthermore, the quantum resources will at first be very scarce and potentially very expensive. This entails a need for a robust management solution. It is important that a good management solution needs to come with adequate monitoring capabilities.

Good management solutions may also be key to optimising the networks which in turn may be crucial in making them economically feasible. Unlike user data that is transmitted over classical networks, quantum networks only need to generate generic Bell Pairs. This leaves a lot of room for pre-allocating resources in an efficient manner.

5.3.2. The principles of a quantum internet

The principles support the goals, but are not goals themselves. The goals define what we want to build and the principles provide a guideline in how we might achieve this. The goals will also be the foundation for defining any metric of success for a network architecture, whereas the principles in themselves do not distinguish between success and failure.

1. Bell Pairs are the fundamental building block

The key service that a quantum network provides is the distribution of entanglement between the nodes in a network. This point additionally specifies that the entanglement is primarily distributed in the form of the entangled Bell Pair states which should be used as a building block in providing other services, including more complex entangled states.

2. Fidelity is part of the service

In addition to being able to deliver Bell Pairs to the communication end-points, the Bell Pairs must be of sufficient fidelity. However, different applications will have different requirements for what fidelity they can work with. It is the network's responsibility to balance the resource usage with respect to the application's requirements. It may be that it is cheaper for the network to provide lower fidelity pairs that are just above the threshold required by the application than it is to guarantee high fidelity pairs to all applications regardless of their requirements.

3. Bell Pairs are indistinguishable

Any two Bell Pairs between the same two nodes are indistinguishable for the purposes of an application provided they both satisfy its required fidelity threshold. This point is crucial in enabling the reuse of resources of a network and for the purposes of provisioning resources to meet application demand.

4. Time as an expensive resource

With the current technology, time is the most expensive resource. It is not the only resource that is in short supply (memory, and communication qubits are as well), but ultimately it is the lifetime of quantum memories that imposes the most difficult conditions for operating an extended network of quantum nodes. Current hardware has low rates of Bell Pair generation, short memory lifetimes, and access to a limited number of communication qubits. All these factors combined mean that even a short waiting queue at some node could be enough for the Bell Pairs to decohere.

However, time is only expensive once quantum operations are underway. If no quantum operations are currently being processed then the network can use this time to prepare and provision resources.

As hardware improves, the need for carefully timing quantum operations may become smaller. It is currently unknown what the cost of these improvements will be, but it is conceivable that there is value in having relatively cheap and undemanding links connected at the edges of a network which will have very short memory lifetimes and low rates of Bell Pair generation.

5. Limit classical communication

This point offers a practical guideline to the issue of timing. A bottleneck in many quantum networked algorithms is the classical communication needed between quantum operations to synchronise state.

For example, some quantum protocols may need to perform a correct for the random outcome of a quantum measurement. For this, they will block the state from further operations until a classical message is received with the information necessary to perform the correction. The time during which the quantum state is blocked is effectively wasted. It reduces the time available for subsequent operations possibly rendering the state useless for an application.

Trade-offs that allow a protocol to limit the number of blocking classical communication rounds once quantum operations have commenced will in general be worth considering.

6. Parallelise quantum operations

A further point to address the issue of timing constraints in the network. The Bell Pairs on the individual links need not be generated one after another along the path between the communication end-points. The order does not matter at all. Furthermore, the order of the swap operations is flexible as long as they don't reduce the fidelity too much. Parallelising these operations is key to optimising quantum protocols.

7. Avoid time-based coordination when possible

A solution to timing constraints is to synchronise clocks and agree on the timing of events. However, such solutions have several downsides. Whilst network clock synchronisation may be accurate enough for certain purposes it introduces an additional element of complexity, especially when multiple nodes in different networks must be synchronised. Furthermore, clock synchronisation will never be perfect and it is conceivable that hardware capabilities advance so much that time-based mechanisms under-utilise resources in the more efficient parts of the network.

Nevertheless, it may not be possible to avoid clocks, but such solutions should be adequately justified.

8. Pre-allocate resources

Regardless of what application is running over the network it will have the same needs as any other application: a number of

Bell Pairs of sufficient fidelity. Whilst the fidelity is a variable number, the indistinguishability of Bell Pairs means that there is lots of flexibility in how a network may provision resources to meet demand. The additional timing constraints mean that pre-allocation of resources will be central to a usable quantum network.

6. Security Considerations

Even though no user data enters a quantum network security is explicitly listed as a goal in this memo. However, as this is an informational memo it does not propose any concrete mechanisms to achieve these goals.

In summary:

- Confidentiality and integrity in the quantum context is the network's guarantee on the minimum fidelity of the delivered Bell Pair states. Uncertainty about the fidelity of a Bell Pair may potentially expose an application to an eavesdropper.

- Authenticity in a quantum network is the guarantee that the other end of the Bell Pair is with the requested partner and not any other third party.

7. IANA Considerations

This memo includes no request to IANA.

8. Acknowledgements

The author would like to acknowledge funding received the Quantum Internet Alliance.

The author would further like to acknowledge Stephanie Wehner, Carlo Delle Donne, Matthew Skrzypczyk, and Axel Dahlberg for useful discussions on this topic prior to the submission of this memo.

9. Informative References

[1] Bennett, C. and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing", Theoretical Computer Science 560, 7-11, 2014, <http://www.sciepub.com/reference/53249>.

[2] Crepeau, C., Gottesman, D., and A. Smith, "Secure multi-party quantum computation. Proceedings of Symposium on Theory of Computing", Proceedings of Symposium on Theory of Computing , 2002, <https://arxiv.org/abs/quant-ph/0206138>.

[3] Giovanetti, V., Lloyd, S., and L. Maccone, "Quantum-enhanced measurements: beating the standard quantum limit", Science 306(5700), 1330-1336, 2004, <https://arxiv.org/abs/quant-ph/0412078>.

[4] Castelvecchi, D., "The Quantum Internet has arrived (and it hasn't)", Nature 554, 289-292, 2018, <https://www.nature.com/articles/d41586-018-01835-3>.

[5] Wehner, S., Elkouss, D., and R. Hanson, "Quantum internet: A vision for the road ahead", Science 362, 6412, 2018, <http://science.sciencemag.org/content/362/6412/eaam9288.full>.

[6] Aspect, A., Grangier, P., and G. Roger, "Experimental Tests of Realistic Local Theories via Bell's Theorem", Phys. Rev. Lett. 47 (7): 460-463, 1981, <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.47.460>.

[7] Muralidharan, S., Kim, J., Lutkenhaus, N., Lukin, M., and L. Jiang, "Ultrafast and Fault-Tolerant Quantum Communication across Long Distances", Phys. Rev. Lett. 112 (25-27), 250501, 2014, <https://arxiv.org/abs/1310.5291>.

[8] Humphreys, P., Kalb, N., Morits, J., Schouten, R., Vermeulen, R., Twitchen, D., Markham, M., and R. Hanson, "Deterministic delivery of remote entanglement on a quantum network", Nature 558, 268-273, 2018, <https://arxiv.org/abs/1712.07567>.

[9] Meter, R. and J. Touch, "Designing quantum repeater networks", IEEE Communications Magazine 51, 64-71, 2013, <https://ieeexplore.ieee.org/document/6576340>.

Author's Address

Wojciech Kozlowski
QuTech
Building 22
Lorentzweg 1
Delft 2628 CJ
Netherlands

Phone: +31 (0)15 2787077
Email: w.kozlowski@tudelft.nl