

Quantum Networking 101

Aliro



Quantum Networking 101


Overview and background information.....	1
The Fundamentals of Quantum Networking.....	2
Qubits.....	2
Entanglement.....	3
Superposition.....	3
Teleportation.....	3
No-Cloning Theorem.....	4
How Entanglement-based Quantum Networks Distribute Entanglement.....	5
Quantum Networking Hardware.....	6
Quantum Networking Software.....	6
The Quantum Network Stack.....	7
Protocols for Entanglement Distribution.....	8
Example of Quantum Protocols Being Used for Entanglement Distribution.....	12
How Entanglement-based Quantum Networks Perform Quantum Secure Communication.....	14
Entanglement-based key generation.....	14
QKD networks vs Entanglement-based networks.....	22
Quantum Repeaters.....	23
Common misconceptions along the quantum networking journey.....	23
Closing.....	24
References.....	25



Overview and background information

Throughout this white paper, “quantum network” will refer to entanglement-based quantum networks. Protocols, hardware and software components employed by quantum networks, and the use cases and applications that are enabled by these networks are explained through the example of a 5-node quantum network. While this example will show what quantum networks could actually look like at each layer of the quantum networking stack, it will not and cannot capture what every single quantum network will look like. This is due to the fact that there are a wide range of possibilities for quantum networks that vary in architecture, design, and use cases, and with that variety comes incredible versatility. Entanglement-based quantum networks can interconnect quantum computers, quantum sensors, and other quantum devices. By distributing entanglement between these devices, the distributed entanglement can be used to run powerful end user applications such as Quantum Secure Communications, or QSC.

Entanglement-based quantum networks do not exist purely in science fiction or in the distant future: there are already entanglement-based quantum networks up and running today.



[Click here](#) to watch a video about AliroNet™ and the EPB Quantum Network

Pictured in the graphic above is a schematic for the EPB Quantum Network, which launched in 2023. While this is the first commercial quantum network to be deployed, there are many other quantum networks out there, such as quantum network test-beds and lab-managed networks. To find out more about the EPB Quantum Network and other quantum networks in North America and globally, check out the on-demand webinar [Real World Quantum Network Deployments](#).

The Fundamentals of Quantum Networking

Quantum networks operate in fundamentally different ways from classical networks, leveraging these properties of quantum mechanics. It is these unique properties that enable superior communications security, enhanced computational speed and capability, and increased capacity across a variety of technologies such as atomic clocks, magnetometers, and gravity gradiometers. The foundations of quantum mechanics with particular relevance to quantum networking include qubits, entanglement, superposition, teleportation, and the no-cloning theorem.

Qubits

Quantum bits, or qubits, are the quantum analog of classical bits. Just like classical bits are the basic unit of information by classical devices like your phone, your computer, and the internet are built upon, quantum bits are the basic unit of quantum information that quantum computers, quantum sensors, and quantum networks are built on.

There are many different ways of creating qubits. A few of the most common types of qubits are:

1. Superconducting qubits. These qubits are based on superconducting circuits that exhibit quantum properties. They operate at extremely low temperatures, close to absolute zero. Superconducting qubits are used in quantum processors for computation. Companies such as IBM, Google, and Rigetti use this type of qubit in their quantum computers. These are stationary qubits: they aren't used for quantum communication.
2. Trapped ion qubits. In this approach, individual charged atoms are trapped with electromagnetic fields in a vacuum chamber. The quantum state of the qubit is defined by the internal energy states of the ion or the spin states of the electrons. Trapped ion qubits are highly accurate for quantum operations. They have relatively long coherence times (the time a qubit can maintain its quantum state), making them ideal candidates for complex quantum computing algorithms, such as Shor's algorithm and Grover's algorithm. These are stationary qubits: they generally aren't used for quantum communication.
3. Photonic qubits. At a high level, photonic qubits are individual photons: units of light. These qubits are particularly useful in quantum communication and quantum cryptography, such as in Quantum Secure Communications (QSC). Photonic qubits are ideal for quantum networks because classical networking infrastructure supports their use, they travel at the speed of light, and are less susceptible to decoherence. A photon becomes a qubit when it is encoded with quantum information. Encoding quantum information in photons typically involves defining certain properties of photons such as their polarization, phase, frequency, or path. Special devices are used to accomplish this.

While classical bits and quantum bits each carry information in computing and communications applications, they are quite different in behavior. Qubits display many quantum properties, including entanglement, superposition, teleportation, and the no-cloning theorem. It is these properties of qubits that enable quantum technologies to outperform their classical counterparts. More about how qubits are used in quantum networks can be found in the on-demand webinar [Qubits: Understanding Quantum Information](#).

Entanglement

Qubits can be coupled together in such a way that measuring one qubit will affect the state of another qubit or qubits. This type of coupling is known as entanglement. Qubits are considered to be entangled when measuring one qubit will instantaneously affect the state of the qubit (or qubits) it is entangled with, even if these qubits are many light years apart. This is exactly the spooky action at a distance that upset Einstein and many other physicists in the early 20th century. Measuring one qubit will affect the state of a qubit it is entangled with faster than light can travel between them, however this does not mean entanglement can be used for faster-than-speed-of-light communication.

Entanglement is said to be distributed between users on a quantum network. For example, if two users receive a qubit from an entangled pair, entanglement has been distributed between them. This distributed entanglement can then be consumed by the users to achieve faster and more powerful computation, more precise sensing, bolstered cybersecurity, and other quantum-powered use cases.

Superposition

Quantum superposition refers to the ability of a quantum system, like a qubit, to be in multiple states simultaneously.

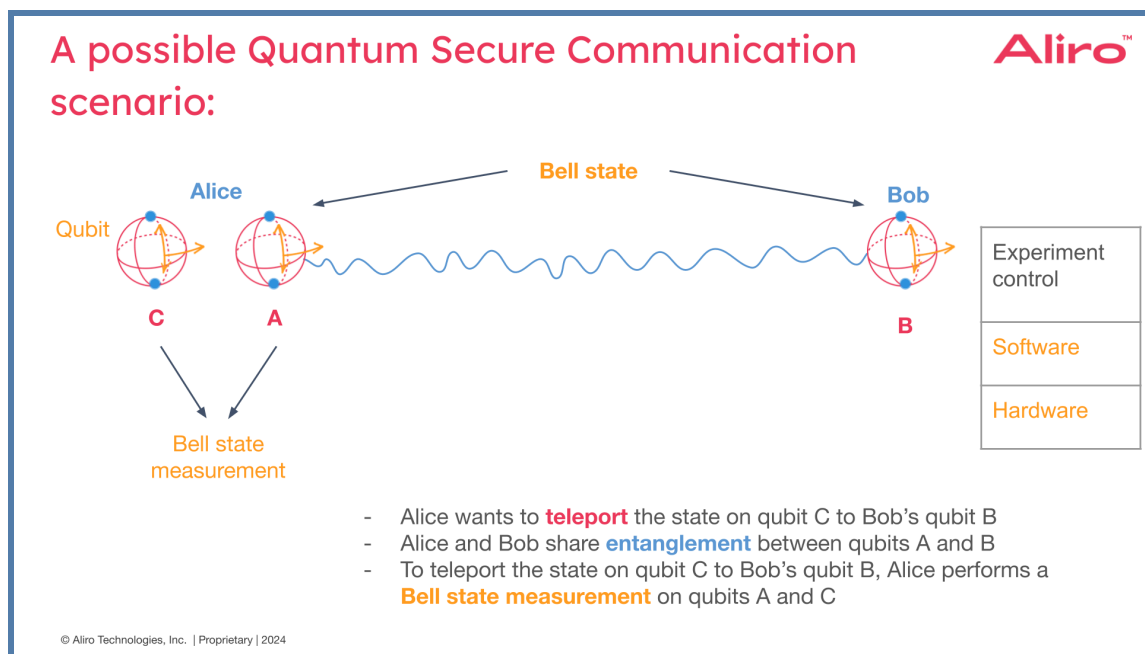
A classical bit can only have one of two values: 0 or 1. Qubits can have the value zero, one, or some combination of these values - and there are infinite such combinations of these values. When qubits are in a combination of values, we say that it is in a superposition of these values. This might sound abstract, so let's use an analogy that mirrors this phenomenon, as depicted in the illustrative example of a skateboarder provided by the [National Institute of Standards and Technology \(NIST\)](#).

Imagine a skateboarder on a halfpipe ramp. In classical terms, the skateboarder could be either at the bottom (location zero) or at the top (location one) of the ramp. In quantum mechanics, however, the skateboarder (akin to a qubit) can exist simultaneously at both locations and everywhere in between. This superposition state is a combination of potential positions.^[NIST]

Prior to measurement a qubit can exist in a superposition state of one of its quantum properties, such as polarization. However, the act of measuring this qubit's quantum property of polarization forces it to 'choose' one of two possible states (either horizontal or vertical), collapsing the superposition. In this way, measurement affects the state of quantum systems. Prior to observation, a quantum system can only be described in terms of probabilities: there's a certain probability of finding a qubit has horizontal polarization, and another probability of finding a qubit has vertical polarization. Once it is measured or observed, the qubit settles into one of its possible states. This phenomenon is famously exemplified in the thought experiment of Schrödinger's cat, where a cat in a box is considered to be simultaneously alive and dead until someone opens the box and observes the cat as alive or dead.

Teleportation

Quantum teleportation can be thought of as a way of moving quantum information from one place to another, but instead of moving a physical qubit, its state is transferred. The process of quantum teleportation uses distributed entanglement in order to make this transfer without physically sending that information through the network's infrastructure.



Above is a simple example of quantum teleportation. There are two nodes, Alice and Bob. Alice and Bob are already sharing entanglement between two of their qubits: qubits A and B. Alice has an additional qubit, qubit C. This is the qubit that has a secret state on it. Alice wants to teleport this state to Bob's location. In order to do that, Alice is going to perform a Bell state measurement between qubits A and C. This act of measuring A and C teleports qubit C to Bob's qubit B.

No-Cloning Theorem

The no-cloning theorem states that it is impossible to create an exact copy of an arbitrary unknown quantum state. This is a rule that emerges directly from the principles of quantum mechanics (the way quantum information behaves) and has profound implications for how information is handled in quantum systems. Quantum information is delicate and complex.

In quantum mechanics, particles such as photons can exist in various states represented by their wave functions. These states might include the position, momentum, spin, or polarization of these particles. Quantum states are probabilistic rather than deterministic, meaning they provide probabilities of finding a particle in a particular state rather than certainties.

The reason quantum states can't be perfectly copied boils down to two key features of quantum mechanics: the linearity of quantum operations and the measurement problem.

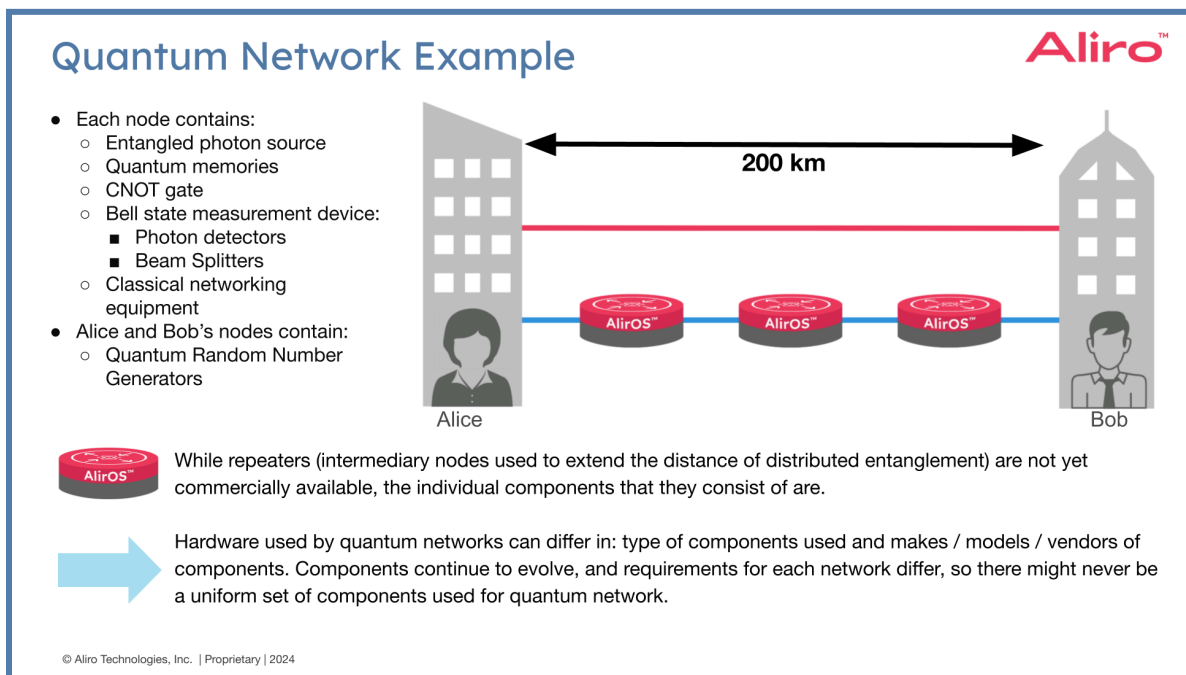
- **Linearity of Quantum Operations:** All operations in quantum mechanics are linear, meaning they must preserve the linear nature of quantum states. When an attempt is made to clone a quantum state, any operation designed to duplicate the state would have to act linearly on the quantum state. However, no linear operation can duplicate all properties of an arbitrary state without violating the structure of quantum mechanics.
- **Quantum Measurement:** The act of measuring a quantum state typically changes the state itself. To clone a quantum state, one would need to measure it to determine its configuration accurately. But quantum measurement affects the state, altering the very information you wish to copy.

The fact that quantum information cannot be copied the way classical bits can be copied may seem like a limitation, but it's also quite useful: because quantum information can't be cloned without changing the state, any eavesdropping attempt on quantum communications can be easily detected.

How Entanglement-based Quantum Networks Distribute Entanglement

Entanglement-based quantum networks are multipurpose networks that can support a wide variety of applications simultaneously. This is possible through distributing entanglement across the network. At a high level, this is accomplished by generating pairs of entangled photons, which are encoded with quantum information through their quantum state, and distributing them to different network nodes. These entangled qubits can then be used to transmit quantum information with a high level of security due to the quantum properties previously discussed. Quantum repeaters can be used to extend the range of quantum networks across far distances.

In the following example, Alice and Bob will use a quantum network for Quantum Secure Communication (QSC). They are located 200 kilometers apart, which means they cannot use a single point-to-point connection for their quantum channel; point-to-point connections typically support links of up to 50 to 100 kilometers. However, by using three quantum repeaters evenly spaced between them, Alice and Bob will be able to utilize the quantum network to secure their communications.



Throughout this white paper, a blue line is used to represent Alice and Bob's quantum channel and a red line represents Alice and Bob's classical channel.

Quantum Networking Hardware

Each of the five nodes in this network contains:

- Entangled photon sources to produce entanglement on each individual node.
- Quantum memories to store quantum information on the node.
- CNOT gates for an entanglement purification protocol, which we'll discuss later.
- Bell state measurement stations.
 - In this case, probabilistic Bell state measurement stations are used. These are made up of single photon detectors and beam splitters.

The Bell state measurement stations are used for the elementary entanglement generation protocol and the entanglement swapping protocol. In addition to the above hardware, Alice's and Bob's nodes will also use quantum random number generators (QRNGs) to provide true randomness for carrying out BBM92, one of the canonical entanglement-based quantum secure communications protocols. IN BBM92, QRNGs are used to generate secure cryptographic keys by providing true randomness at a certain step of the protocol: measurement basis selection. True randomness is essential for ensuring the security of quantum communications, not only for the security against eavesdropping but also for ensuring the cryptographic keys generated are not susceptible to predictive algorithms or any known pattern analysis.

Classical communication is used for heralding results, timing and synchronization, and other operations that don't require entanglement as a resource, but support entanglement distribution. Quantum networks will not be used to completely replace classical networks, but instead will be used together with classical networks to enable higher performance and new applications that neither network can accomplish on its own.

There are two important things to note about quantum network hardware. First, while repeaters (the intermediary nodes used to extend the distance of distributed entanglement) are not commercially available off-the-shelf, the components that they consist of are. Second, hardware used by quantum networks can differ in types of components used as well as the makes, models and vendors of those components.

Components will continue to evolve, and the requirements for each network will differ, so that there may never be a completely uniform set of components used for all quantum networks. While there may never be complete uniformity, standardization, management software, and the implementation of best practices will ensure a smooth evolution of components into the future.

Quantum Networking Software

There are many factors to consider when designing a quantum network. Use of a quantum network simulator with enough accuracy and the capabilities to model, verify, and validate network designs is vital for an efficient and effective quantum network design process. A quantum network simulator is also useful for upgrading and scaling a quantum network.

This example is relatively simple. In networks with many more nodes and in more complex topologies, there will be myriad possible paths for distributing entanglement between end users. With so many hardware components and protocols required to distribute entanglement, control and orchestration software will be vital to keeping the network running smoothly. Software is needed to configure, manage, run and utilize quantum networks efficiently and effectively. Even for very simple quantum networks like this example, with relatively few components, it's nearly impossible to manually configure

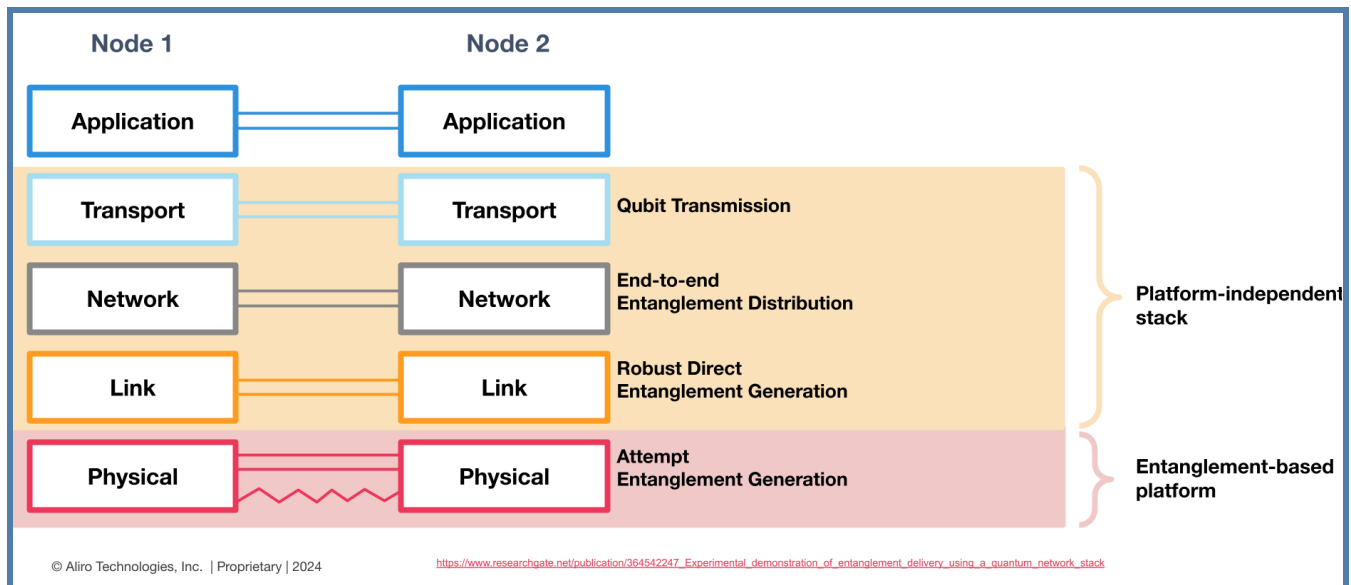
and manage hardware for applications such as Quantum Secure Communication. This software handles other important entanglement distribution-related tasks such as transmission and routing of information in real-world scenarios that must be responsive to noise, photon loss, and probabilistic protocols. Retransmitting information in response to these real-world scenarios, and at multiple layers of the stack, is critical.

Additionally, while this network example is relatively simple - with only one path between Alice and Bob - there could be many paths to choose from. Software can be used to select pathing that optimizes for factors like security and throughput. Quantum networks will vary greatly in terms of hardware deployed, and thus it's also vital that the software used to manage the network be hardware agnostic - giving the network the most flexibility for expanding the network and adding or upgrading hardware components.

The Quantum Network Stack

The quantum network stack is the implementation of protocols needed to accomplish the primary objectives of entanglement-based quantum networking: distributing entanglement, and then utilizing that distributed entanglement for a specific use case. The stack is typically broken into five layers, starting from the bottom:

- The physical layer is used to attempt entanglement generation.
- The link layer is used for robust direct entanglement generation.
- The network layer is used for end-to-end entanglement distribution.
- The transport layer is used for qubit transmission.
- End user applications make up the top layer of the stack.^[NETWORK-STACK]



Each layer of the stack will be examined through the example of using an entanglement-based network and the BBM92 protocol for entanglement-based key generation. While specifics of all the protocols used in the stack could differ, this example will still give a good overview of how a quantum network operates.

This example begins at the bottom of the network stack with the first three layers contributing to entanglement distribution, and works up to the application layer: entanglement-based key generation. For a quantum network to be able to best meet the needs of its users and run their desired applications, there are some specifications that entanglement distribution must meet:

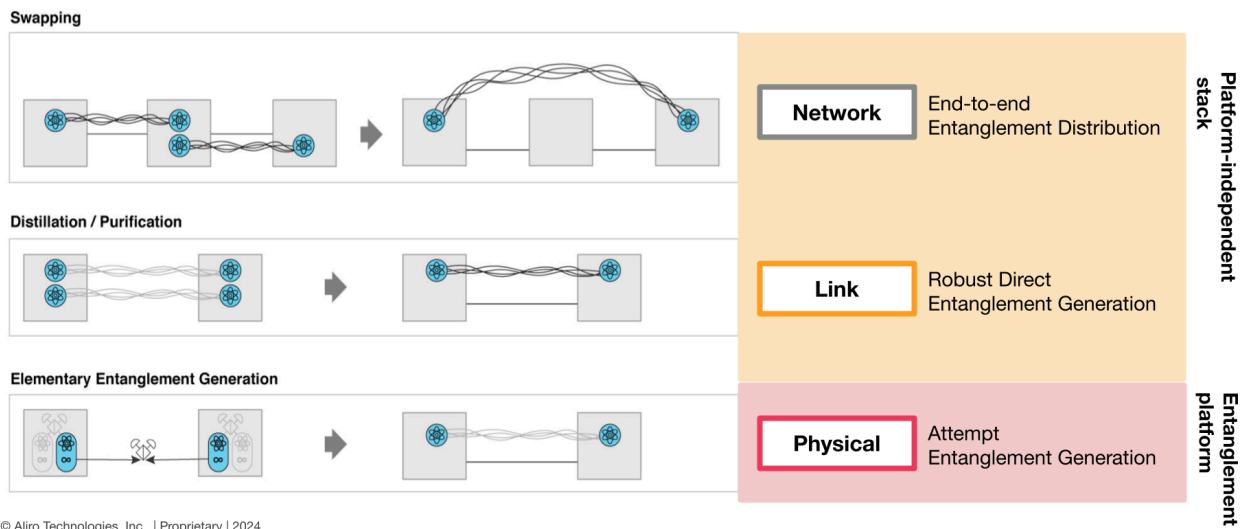
- The distributed entanglement must have a high quality of entanglement, also known as fidelity.
 - Fidelity quantifies the accuracy of a quantum state. It refers to the degree of similarity between the intended quantum state and the actual quantum state after a qubit has been transmitted, stored, or manipulated within the network.
 - Entanglement must meet a certain level of fidelity in order to run the desired end user applications.
- The quantum network must have high enough throughput / high enough message delivery rate for entanglement distribution to be successful.
 - Applications often require many entangled pairs, so it's vital to any application using entanglement that there are enough of these entangled pairs to successfully run them.
- Sometimes it will be necessary to connect distant quantum devices. For example, creating a provably quantum-safe communication channel between two users located vast distances from each other.
- It's important to note that small scale entanglement-based quantum networks are also incredibly useful. For example, clustered quantum computing (achieved through networked quantum processors) will typically take place in a single warehouse, and there are also situations where provably quantum-safe communication between two nearby users is desirable.

Meeting the requirements for distributed entanglement involves a lot of coordination between different parts of the network: the appropriate hardware and software needs to be put into place, protocols and topology must be carefully implemented, logistical considerations need to be taken into account, as well as the integration of existing classical systems. While quantum networks share the same general goals of distributing entanglement and then utilizing that distributed entanglement for a set purpose, different users will have unique requirements that shift the specifications of the network. For example, users may wish to run a wide variety of applications and could face constraints such as budget, logistics, or environmental factors that impact how a network is implemented.

Protocols for Entanglement Distribution

There are three main processes / protocols involved in distributing entanglement for end user applications:

- Elementary entanglement generation, or EEG.
 - Used to distribute entanglement between nearby devices. Note that for local area quantum networks, repeaters or other scaling technologies are not necessary. In this case, EEG alone can be used for entanglement distribution.
- Entanglement purification, sometimes known as entanglement distillation
 - Used to ensure that the quality, or fidelity, of the entanglement is at a high enough level throughout the entire entanglement distribution process so it can be used for end user applications.
- Entanglement swapping.
 - Used to extend the distance of the entanglement distributed by EEG. Using swapping and EEG together aids in distributed entanglement across far distances.



These processes, or protocols, align with the bottom three layers of the quantum network stack. EEG is a physical layer protocol, where entanglement is distributed between neighboring nodes. Entanglement purification is a link layer protocol, creating robust entanglement generation. Entanglement swapping is a network layer protocol, crucial to creating end-to-end entanglement across distance.

Elementary Entanglement Generation

In Elementary Entanglement Generation (EEG), entanglement is distributed between two nearby quantum nodes. There are many protocols, and even families of protocols, for performing EEG. The example shown in the graphic below depicts a meet-in-the-middle scheme. Another common EEG scheme is that of a midpoint source - where an entangled pair is created at a single midpoint and then each of these entangled qubits is transmitted to the two desired nearby nodes.




The first step of this meet-in-the-middle technique is to generate an entangled pair consisting of at least one photonic qubit at each of the two neighboring nodes. This is achieved with an entanglement source. Then a photonic qubit from each of the two entangled pairs is transmitted via optical fiber to a probabilistic Bell state measurement station, consisting of a beam splitter and two single photon detectors.

The two qubits must arrive at the Bell state measurement station at exactly the same time. Quantum memories can be used to help ensure this exact timing. Depending on which side of the beam splitter

the photons exit - which can be confirmed by checking the clicks of the photon detectors - it's possible to determine if entanglement generation was successful. If it was not, then the process is repeated.

Entanglement Purification

In entanglement purification, multiple entangled states with lower fidelity are used to create a single entangled state with higher fidelity. Purification protocols require a certain level of fidelity to be effective, so it's important to ensure fidelity has not degraded too much before the purification is performed.

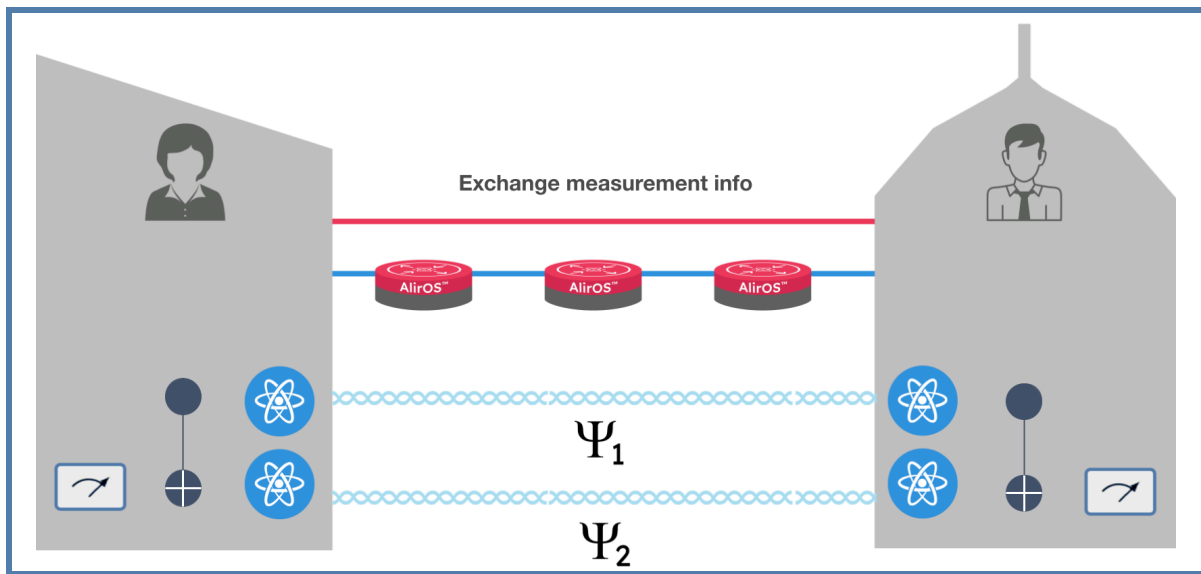


Qubits and the quality of their quantum information are fragile. In this example, the fidelity of the entanglement decays due to noise. Noise refers to many different factors that can cause loss of quantum information, such as:

- environmental factors like temperature and vibrations
- manufacturing defects in components
- natural physical processes

To combat noise and keep fidelity at a high enough level to run end-user applications, purification must be performed many times throughout the entanglement distribution process.

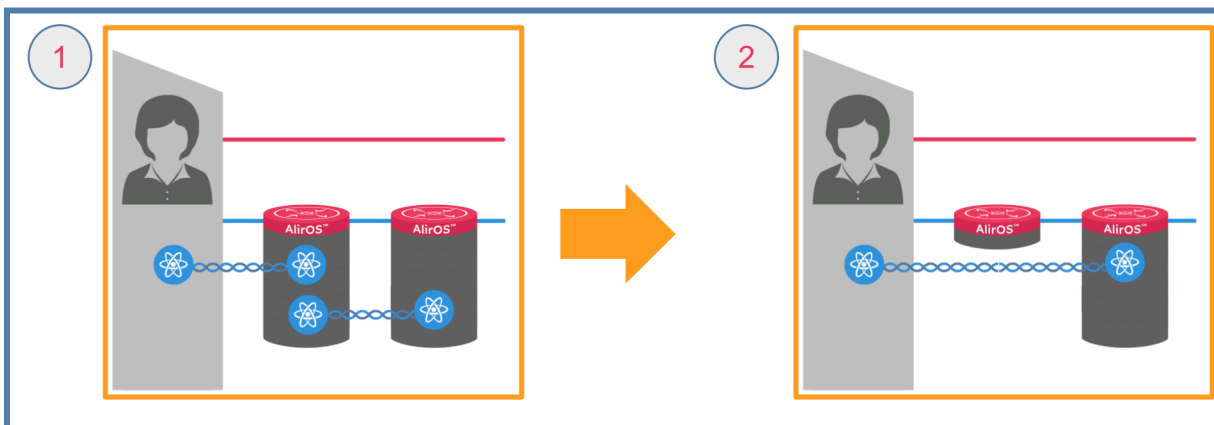
There are many different purification protocols that can be implemented. This example uses a Controlled-NOT (or CNOT) gate in the purification protocol. CNOT gates are essential components of gate-based quantum computers. CNOT gates are often used to entangle and disentangle quantum states in quantum computing algorithms. However, they can also be used in quantum networks.



This example uses a relatively simple version of a purification protocol with CNOT gates: Alice and Bob share two entangled pairs. Ψ_1 is an entangled pair whose fidelity they are trying to increase. Ψ_2 is the entangled pair they are using to help accomplish this improvement. CNOT gates operate on two qubits. It's important to note that when it comes to optical components and using photonic qubits in photonic systems, the CNOT gate will be probabilistic. Alice and Bob first apply CNOT gates to their two qubits respectively, with the control qubit being from Ψ_1 and the target qubit being from Ψ_2 . Alice and Bob next measure their half of the pair of qubits from Ψ_2 and compare the results. If they get the same result, then the distillation is considered a success and they will keep Ψ_1 . The fidelity of Ψ_1 will be at least as high, but often higher, than before. If Alice and Bob get different results, then the distillation or purification is considered a failure. They will discard Ψ_1 and try again.

Entanglement Swapping

Entanglement swapping enables the extension of entanglement between two pairs of qubits over longer distances. In the example, there are two pairs of entangled photons, with one photon from each pair converging at the same midpoint node. This setup facilitates the "swapping" of entanglement, effectively linking the two initially separate pairs into a single, extended entangled pair.



To accomplish this, the two qubits at the middle node are measured in a clever way by performing a Bell state measurement, at which point the two qubits that were previously entangled with qubits at the center node will become entangled to one another - creating a longer distance entanglement. This process is similar to the meet-in-the-middle elementary entanglement generation scheme discussed above. Using EEG and then performing entanglement swapping distributes entanglement to non-neighboring nodes. Using purification ensures that the fidelity of this entanglement remains high.

Example of Quantum Protocols Being Used for Entanglement Distribution

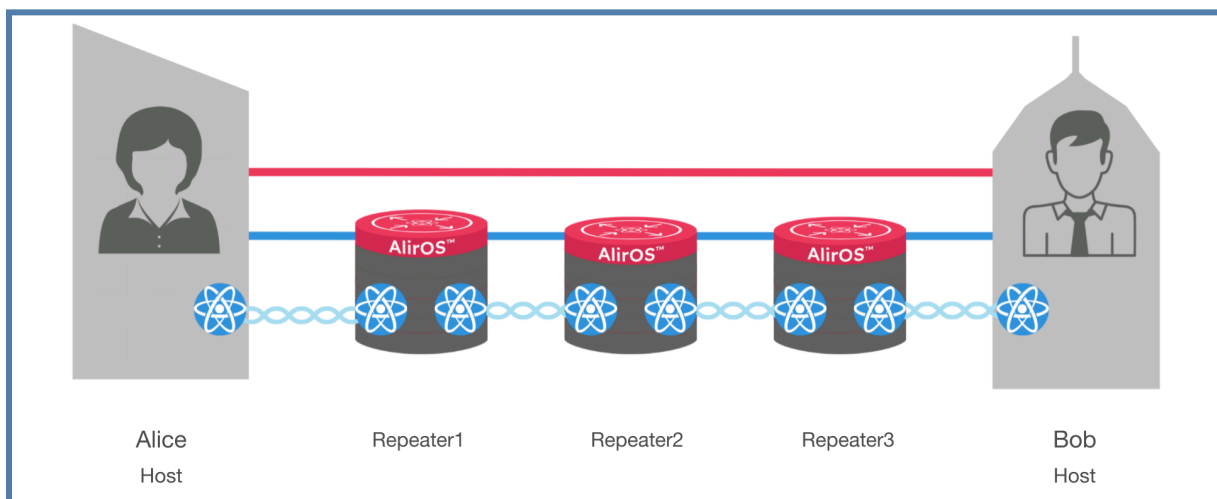
The following example demonstrates entanglement distribution on Alice and Bob's network. For brevity, this example assumes that all of the probabilistic processes work on the first attempt.

The goal of this example is to distribute entanglement between Alice and Bob, located 200 km away.

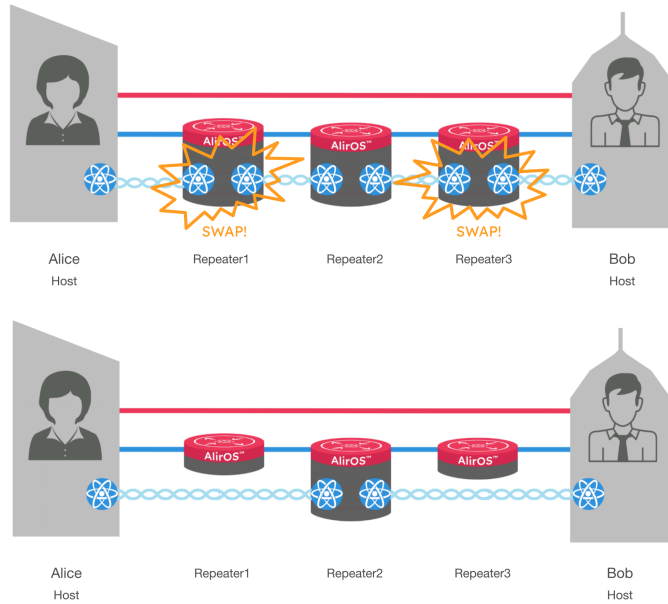
The first step is to perform EEG to create entanglement between each pair of neighboring nodes. Once the EEG protocol for each pair of neighboring nodes has been successful, there will be entanglement between Alice and Repeater1, and between Repeater1 and Repeater2 and so on. Entanglement must be present before entanglement swapping can be performed.

Throughout this example, remember that it may be helpful to perform EEG multiple times between each neighboring node in order to purify the entanglement and ensure fidelity stays high enough for Alice and Bob to make use of the distributed entanglement. The entire entanglement distribution process described in the example, from EEG to purification to swapping, will be performed many times because Alice and Bob need to use more than just one shared entangled pair.

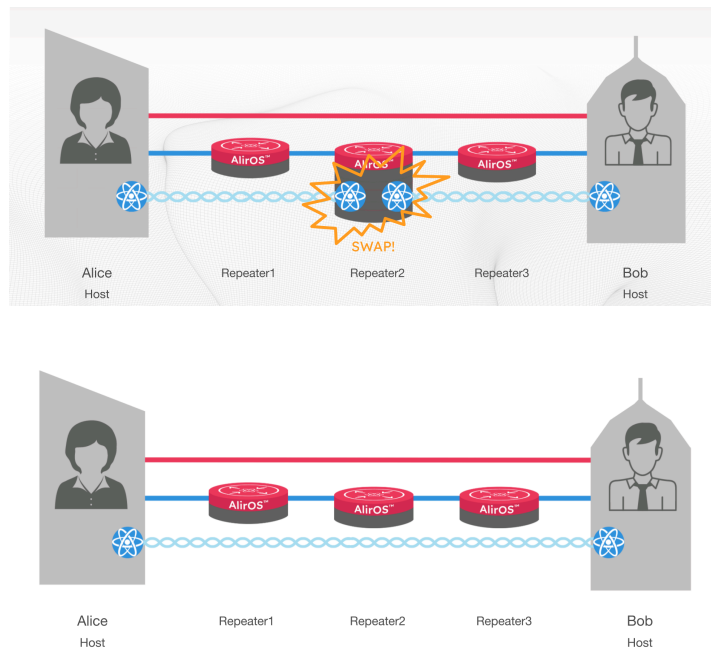
Now that entanglement exists between each pair of neighboring nodes, the nodes are ready for entanglement swapping.



Entanglement swapping at Repeater1 extends entanglement from Alice to Repeater2, which is closer to Bob. A swap at Repeater3 will extend entanglement to Bob and Repeater2 which is closer to Alice. These entanglement swaps could happen one at a time, at the same time, or in any order.



Now another swap is performed at Repeater2, which extends entanglement between Alice and Bob directly.



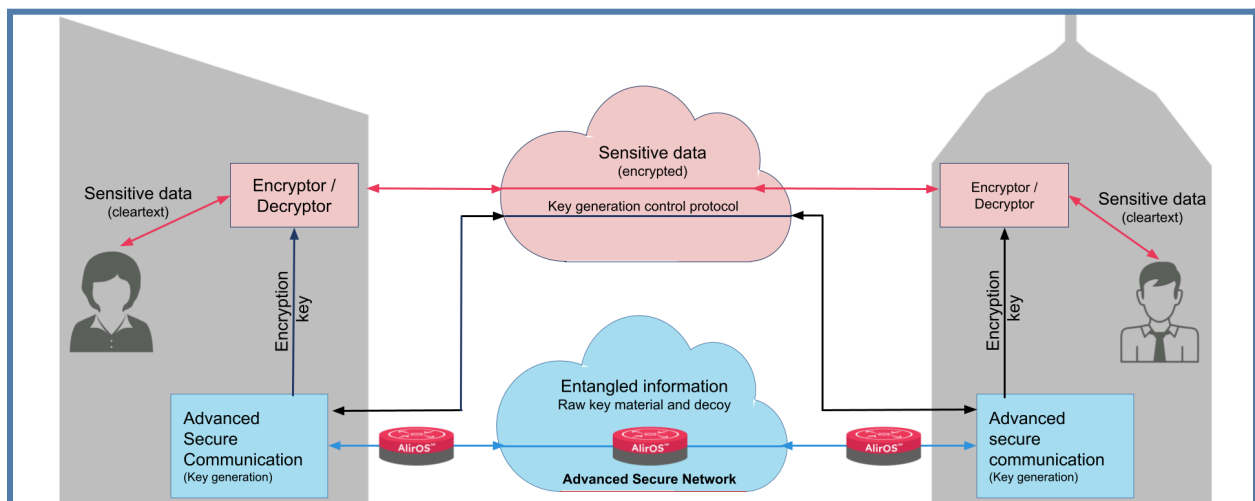
Purification would likely be performed throughout this process many times to keep entanglement quality high, and this process of EEC and swapping would be performed many times so that Alice and Bob share multiple entangled pairs instead of just one entangled pair.

How Entanglement-based Quantum Networks Perform Quantum Secure Communication

While entanglement-based networks have vast capabilities that can be simultaneously enabled or executed, the focus in this example is on Quantum Secure Communication, or QSC. There are many types of information that must remain private, such as information related to defense and military, intellectual property, financial information, medical information, etc. Keeping information secure can be a matter of personal, organizational, and even national security. In this example of QSC, Alice and Bob are using an entanglement-based quantum network to communicate their top secret information. They will use entanglement-based key generation to secure their communications with the BBM92 protocol.

Entanglement-based key generation

Using BBM92 begins with the distribution of entanglement between Alice and Bob, in a process like the previous example. This process is performed many times, in order to distribute a sufficient amount of high-fidelity entanglement between Alice and Bob. Once there is sufficient entanglement between Alice and Bob, the key generation process can begin. Remember that in the example, Alice and Bob are using distributed entanglement for the purposes of Quantum Secure Communication, and specifically for key generation, to generate a key to secure their top secret information.



Entanglement-based key generation utilizes distributed entanglement to generate a shared secret key between end users. This method has information-theoretic security, which means regardless of how strong an adversary's computational resources, whether classical or quantum, the key cannot be broken. This is especially important considering the impending nature of Q-Day. Learn more about Q-Day and how organizations can prepare for it in this on-demand webinar ["What is Q-Day?"](#)

Beyond the strength of the key itself, entanglement-based key generation is also secure in implementation as the key will never be revealed anywhere on the network. This shared quantum secret key can be used in many ways by classical networks. In this example, Alice and Bob use their shared secret key with a symmetric encryption algorithm, such as AES256, to encrypt and decrypt messages and thus secure their communication. In the graphic below, Alice and Bob are connected via classical and quantum channels.

A quantum key is generated using the quantum channel, with help from the classical channel. The quantum key is then used by the classical channel to help secure Alice and Bob's classical communications. BBM92 is the canonical entanglement-based key generation protocol, along with E91.

There are six main steps of BBM92:


- Basis selection
- Measurement
- Conversion to bits
- Sifting
- Calculating qubit error rate
- Generating the final key

The graphic below is an example of generating a key via BBM92.

	Entangled State $ 00\rangle + 11\rangle =$ $ ++\rangle + --\rangle$	Alice Bob	Alice Bob	Alice Bob	Alice Bob	Alice Bob	Alice Bob	Alice Bob	Alice Bob	Alice Bob
1	Basis Selection	X Z	X X	Z X	Z Z	X X	Z X	Z Z	X Z	X X
2	Measurement	+ 1	- -	0 +	1 1	++	1 -	0 0	- 0	++
3	Conversion to Bits	0 1	1 1	0 0	1 1	1 1	1 1	0 0	1 0	1 1
4	Sifting (Same Basis?)									
5	Calculate QBER (Count Errors in a Small Sample)		0%			0%				
6	Generate Final Key				1			0		1

Measurement bases are ways to view and measure qubits. Alice and Bob choose a measurement basis for each of their qubits from the distributed entangled pairs. Choice of each measurement basis is chosen independently and randomly. In this example, quantum random number generators are used for basis selection to ensure true randomness. This example also uses the canonical choices for bases: the Z-basis which has possible measurement results in terms of zero and one, and the X-basis which has possible measurement results in terms of plus or minus.

These bases coincide with the Z-axis and the X-axis of the Bloch sphere representation of qubits respectively.



As an analogy to help visualize these bases, imagine a collection of blocks. These blocks can be at any angle. The Z-basis views the blocks in terms of being at a completely vertical or completely horizontal angle. When the angle of a block is measured in the Z-basis, the result can only be either horizontal or vertical. The X-basis views the blocks in terms of being tilted 45 degrees left of vertical or tilted 45 degrees right of vertical. When the angle of a block is measured in the X-basis, the result can only be one of these angles.

The probability of which result will be returned after measurement depends on how close the angle of the block is to the angle of the basis state. For example, a block that's mostly vertical will be more likely to be measured as vertical than horizontal when measuring in the Z-basis. With different measurement bases the same qubits are still being observed and measured, but through different lenses. Alice and Bob measure each received qubit in the selected basis. The measurement process collapses the qubit's quantum state - which is a superposition of multiple states - into exactly one of the two possible basis states.


Note: Remember that in every protocol discussed in this paper, measurements have been made with the same measurement stations: performing teleportation, EEG, entanglement swapping, and measuring the target qubit in the purification protocol. These measurements have been carried out in part by single photon detectors.

If Alice and Bob measure their qubit on the same basis, then they will get the same result.

For example, if they both measure in the Z-basis and Alice gets results zero, then Bob must also get result zero. Similarly, if they both measure in the X-basis and Alice gets the result minus, then Bob must also get the result minus. However, if Alice and Bob measure in different bases, then there can be any combination of measurement results. For example, if Alice measures in the Z-basis and Bob measures in the X-basis, possible results are:

Alice measuring 0 and Bob measuring +, Alice measuring 0 and Bob measuring -, Alice measuring 1 and Bob measuring +, and Alice measuring 1 and Bob measuring -.

Once measurement is complete, Alice and Bob convert all of the measurement results from quantum states into classical bits, or zeros and ones. When measuring in the Z-basis, the measurement results are already in bits. When measuring in the X-basis, + is converted to 0, and - is converted to 1. Assuming that there is no noise, when Alice and Bob measure a qubit in the same basis, then they will always receive the same value.



If Alice and Bob measure in different bases, then there is a 50% chance that they will receive a different value. This is because there's a 50% chance of getting measurement results 0 and -, 1 and + which yield different values. There is also a 50% chance that they will receive the same bit values. This is because there is a 50% chance of getting measurement results 0 and + or 1 and -, which yield the same values.

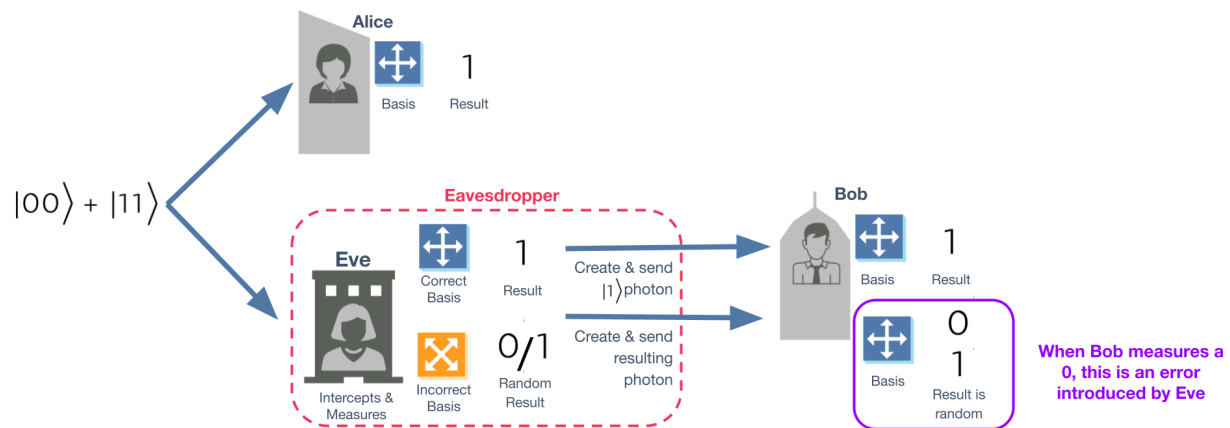
The relationship between bases selected and matching bit values is very important for the security of the BBM92 protocol. To perform sifting, Alice and Bob share their choices of measurement basis over a public authenticated classical channel and discard all results in which they use different measurement bases. The resulting bits are considered to be sifted. In the absence of noise or an eavesdropper, this means that Alice and Bob should agree on all remaining values, as they will agree on all measurement bases for these bits.

Alice and Bob choose a sample of the remaining sifted results and compare them to estimate quantum bit error rate, or QBER. There is a tradeoff here between taking a larger sample to get a better estimate of QBER, and taking a smaller sample that leaves more bits to be used for key material.

Alice and Bob are estimating the percentage of mismatched bit pairs out of the sifted material. In the perfect case, where there is no noise and no eavesdropper, the qubit error rate will be 0% because all of Alice and Bob's sifted measurement results will agree under these perfect conditions. If there is no noise but the qubit error rate is greater than zero, then Alice and Bob know with 100% certainty that an eavesdropper is present.

An eavesdropper, Eve, must choose a measurement basis when measuring the intercepted qubits, just as Alice and Bob do, to discover the quantum information they contain. However, Alice and Bob's qubit measurement bases are not sent over the network until after their measurement occurs. This means that Eve will not have the ability to intercept that information prior to their measurement, and thus could choose a different basis than Alice and Bob when Alice and Bob do agree on a measurement basis.

However, if Eve measures an intercepted qubit in a different basis, then there is a 50% chance that Alice and Bob will receive different values leading to non-zero qubit error rate. This is depicted in the graphic below of how an eavesdropper can contribute to qubit error rate.



Obtaining any information about the key material requires an eavesdropper to measure the system in some way. Because of quantum mechanics, the act of measuring disturbs the system and produces observable detectable anomalies. All such anomalies contribute to the qubit error rate, but neither Alice nor Bob can distinguish whether these anomalies came from environmental noise or errors, or an eavesdropper.

In real-world quantum networks, noise will often be present and the qubit error rate will be greater than zero. Because the source of each error cannot be determined, we assume the worst case scenario (at least from a security perspective) that all errors are caused by an eavesdropper. The BBM92 protocol tolerates up to 11% qubit error rate, which corresponds to a fidelity of 89%. However, there's a tradeoff between security confidence and throughput, so the exact value will differ based on individual user requirements. If qubit error rate exceeds the user-established tolerance limit, the protocol determines that too much information was leaked to potential eavesdropper. The generated key is discarded and a new one is generated, typically on a different quantum channel.

If QBER is within acceptable bounds, no eavesdropper is detected. Alice and Bob then use the remaining bits that were not consumed in estimating QBER to generate a final key. While the remaining bits could also be used as the final generated key, typically classical post-processing techniques will be performed on the remaining bits, such as information reconciliation and privacy amplification protocols, in order to generate an even more secure final key.

The example in the table below shows the steps of BBM92.

In this example, there's no noise and there's no eavesdropper. In reality, a network will likely experience noise and potentially an eavesdropper. There are also only nine entangled pairs in the example. In practice, far more than nine entangled pairs will be used. For example, generating a key of 256 bits requires more than 256 entangled pairs, as bits are lost to sifting and when sampling to estimate qubit error rate.

	Entangled State $ 00\rangle + 11\rangle =$ $ ++\rangle + --\rangle$	Alice Bob	Alice Bob	Alice Bob	Alice Bob	Alice Bob	Alice Bob	Alice Bob	Alice Bob	Alice Bob
1	Basis Selection	X Z	X X	Z X	Z Z	X X	Z X	Z Z	X Z	X X
2	Measurement	+ 1	- -	0 +	1 1	+ +	1 -	0 0	- 0	+ +
3	Conversion to Bits	0 1	1 1	0 0	1 1	1 1	1 1	0 0	1 0	1 1
4	Sifting (Same Basis?)									
5	Calculate QBER (Count Errors in a Small Sample)		0%			0%				
6	Generate Final Key				1			0		1

Step 1. In row 1 of the table, first step of BBM92, Alice and Bob use their quantum random number generators to select a measurement basis for each qubit that they received. For the first pair, Alice's quantum random number generator selected the X-basis, while Bob's quantum random number generator selected the Z-basis. For the second pair, both Alice and Bob will measure in the X-basis. This process continues until enough bases have been selected to create a sufficiently sized key, as well as to account for qubit loss during sifting and for estimating the QBER. For brevity, this example only performs this process nine times. In a real-world scenario, this would be performed many more times.

Step 2. Next, Alice and Bob measure their qubits in the selected bases. Row 2 of the table shows the results Alice and Bob receive. Remember that when Alice and Bob use the same measurement basis and there is no noise or eavesdropper, they will agree on measurement results. When Alice and Bob use different measurement bases, there's only a 50% chance that they will agree on measurement results. For the first pair, Alice and Bob measured + and 1, respectively. For the second pair Alice and Bob both measured -, and so on for each qubit pair.

Step 3. Alice and Bob convert their measurement results to bits with + becoming 0 and - becoming 1. When Alice and Bob have the same measurement basis (and there is no noise or eavesdropper present) they will also agree on bit values. Row 3 in the table above shows the result of converting to bits in this example.

Step 4. Next Alice and Bob perform sifting: sharing their choices of measurement basis over a public authenticated classical channel, keeping the bits in which they use the same measurement basis, and throwing away the bits in which they use different measurement bases. As seen in row 4 of the table, for the first pair of bits Alice used the X-basis and Bob used the Z-basis, and so they throw away those bits. For the second pair of bits, Alice and Bob both use the X-basis, so they'll keep these bits. This process continues for all pairs of bits - in this case, all nine pairs.

Step 5. Now Alice and Bob use a sample of the remaining bits to estimate qubit error rate.

They choose the second and fifth pair as their sample. They find that the values agree for both pairs and so qubit error rate is estimated to be 0% and fidelity to be 100%.


Step 6. Finally, Alice and Bob use the remaining bits as their generated final key. They have generated the key 1 0 1. In this example, no classical post-processing was performed. In practice, post-processing techniques such as information reconciliation and privacy amplification would be used. Notice that in this example, the nine entangled pairs generate a key of only three bits. While this is not an exact ratio of entangled pairs to resulting bits for key material, this example shows that many entangled pairs will be needed to facilitate Quantum Secure Communication with the BBM92 protocol.

A shared secret key can be used in many ways by a classical network. In the example, Alice and Bob generated their key to be used by a symmetric encryption algorithm in order to encrypt and decrypt classical messages.

Shared secret keys can be used for other things beyond encryption and decryption, such as for passwords, generating new keys via use of a key derivation function, identity authentication via challenge-response protocols, message authentication via message authentication codes, and many other useful security functions.

This example covered the three bottom layers of the quantum networking stack, but skipped to the application layer with Quantum Secure Communication via BBM92. The transport layer was not utilized in this example. This doesn't mean the transport layer is not important. It's incredibly useful for many applications. Quantum teleportation is a transport layer protocol; it's an entanglement-based protocol that utilizes distributed entanglement to communicate quantum information across the network, and is even considered an application all on its own.

Many distributed quantum applications require communication of quantum information between devices. Teleportation is the protocol that can be used to accomplish this. The way in which information is sent via teleportation also has useful properties for cybersecurity purposes. Quantum teleportation can be used to communicate quantum information to endpoints on the network without that information ever being exposed on the network itself. This means that even if intermediary nodes between communicating users become compromised, the quantum information that is being sent between them will not be exposed. The security of quantum information being sent is said to be device independent.



Teleportation is illustrated in the graphic above. Entanglement has been distributed across a network, and Alice and Bob share an entangled pair. Alice has a qubit that she wants to send to Bob. By consuming the entanglement with a teleportation protocol, Alice can send the qubit to Bob. The entangled pair is measured during this process, essentially collapsing the qubits into classical bits and

consuming the entanglement. For more detail on how quantum teleportation works, see the on-demand webinar [“The Role of Quantum Teleportation in Entanglement-based Secure Networks.”](#)

The quantum network in the examples used throughout this whitepaper has a linear topology (i.e., there is only one path between the two end users) and it is 200 kilometers long. It employs hardware, software, and protocols to carry out a specific task: Quantum Secure Communications, specifically using BBM92 for quantum key generation. This example does not encompass all quantum networks. Quantum networks can vary in many important ways, such as:

- Topology
- Size (local area, metro area, and wide area quantum networks)
- Hardware components
- Protocols (for example, meet-in-the-middle schemes for certain endpoints or varying schemes for EEG)
- Use cases and applications

Part of the reason quantum networks vary so much is that they can be used for widely varying scenarios. These same entanglement-based networks are multipurpose networks that can be utilized for different use cases simultaneously, such as distributed quantum computing and quantum secure communications. However, a quantum network that connects together quantum computers in a warehouse versus a quantum network that connects quantum devices together across the globe for various use cases will have distinct differences.

Three use cases for entanglement-based quantum networks are quantum secure communications, distributed quantum computing and distributed quantum sensing. Many valuable applications of these use cases have been identified, and many more will be developed as quantum networks evolve into a global-scale quantum internet.

Quantum Secure Communications

Quantum Secure Communications is a family of security solutions that are enabled by entanglement-based quantum networks. Entanglement-based key generation and quantum teleportation enable additional security of communicating information. Shared entanglement between users can also be used to directly communicate classical and or quantum information in a provably secure manner, and to detect eavesdroppers of classical and or quantum information.

Distributed Quantum computing

Distributed quantum computing is the connection of multiple quantum processing units or multiple quantum computers using entanglement-based quantum networks to connect them via distributed entanglement. The quantum computers can be located in close proximity or far distances apart. By connecting quantum computers together, it's possible to achieve much greater computing power and additional security. Details about this use of entanglement-based quantum networks to network QPUs for increased computational power can be found in the on-demand webinar [“Entanglement-based Networks for Increasing Quantum Computing Performance.”](#)

Distributed quantum sensing

Distributed quantum sensing is the connection of multiple quantum sensors via entanglement-based networks. These sensors are connected by distributed entanglement. The sensors can be close together or at distant locations. By connecting the sensors together, it's possible to achieve much

greater accuracy, sensitivity, and precision than individual quantum sensors or networked classical sensors can achieve.

QKD networks vs Entanglement-based networks

QKD networks, or quantum key distribution networks have been around for decades. The more recently available and deployed entanglement-based networks differ from QKD networks in several important ways:

- As their name suggests, QKD networks are used for quantum key distribution. They are single purpose and can be used for quantum key distribution only.
- Entanglement-based networks can be used for multiple use cases simultaneously such as distributed quantum computing, distributed quantum sensing, Quantum Secure Communications - including key generation and a host of other applications.
- QKD networks are prepare-and-measure networks. For example, if Alice and Bob are communicating via QKD, Alice will prepare a quantum state and then send it to Bob, who will measure it. The security lies in the fact that any attempt to eavesdrop on quantum states as they are transmitted will inevitably alter those states in a detectable way. This ensures that any interception or eavesdropping can be detected by the legitimate communicating parties, allowing them to use only uncompromised keys for secure communication.
- In contrast, Entanglement-based networks distribute entanglement across distances, and their security relies on the physics of entanglement itself. The use of entanglement can provide stronger security assurances against a wider range of attacks, including those involving sophisticated quantum hacking techniques that prepare-and-measure networks are vulnerable to.
- QKD networks employ trusted relay nodes when extending the distance of quantum information transmission. These are not relay nodes that are provably trustworthy, but must be assumed trustworthy whether this is the case or not. If they become compromised so too will any information that comes through them.
- Entanglement-based quantum networks use repeaters instead of relay nodes, which don't have that same security vulnerability. The information being communicated is never exposed on the network, so even if a repeater becomes compromised, the information transmitted will not become compromised.
- BBM92 is often considered the entanglement based version/analog of BB84, which is the canonical quantum key distribution algorithm. BBM92 addresses several of the vulnerabilities that come with BB84, such as trusted relay nodes.

Quantum Repeaters

A question that frequently comes up in discussing entanglement-based quantum networks is why classical repeaters and routers can't be used for quantum communication. After all, classical and quantum networks use photons over optical fiber to facilitate communication. There are a few reasons why classical repeaters can't be used for quantum networks.

The way classical repeaters typically work is by copying and amplifying classical signals. This is simply impossible to do with quantum information, due to the no-cloning theorem. The no-cloning theorem shows that it's not possible to create an identical copy of an arbitrary quantum state. The process classical repeaters use for copying and amplifying classical signals cannot create identical copies of arbitrary quantum states.

Classical repeaters are designed for classical information processing. They're not designed for quantum information processing. For example, they can't store quantum-encoded photons or carry out quantum protocols. Quantum information experiences decoherence, the degradation of quantum information from noise and other factors. Classical repeaters can't maintain the fragile quantum information intended for transmission.

Common misconceptions along the quantum networking journey

There are several common misconceptions that organizations often have about deploying entanglement-based quantum networks.

1. Quantum networks will completely replace classical networks.

This isn't the case. Quantum networks are used in tandem with classical networks. Many quantum networking protocols require both a quantum channel and a classical channel. In addition to both classical and quantum networking channels to accomplish use cases like Quantum Secure Communication, quantum networks do not require a massive overhaul or forklift upgrade of classical systems.

2. Underestimating the technology readiness level of quantum networks.

Quantum networks are not a futuristic technology, many years away from practicality. Quantum networks are ready for deployment today. For examples of quantum networking projects in North America, see the on-demand webinar [“Real World Quantum Network Deployments”](#).

3. Overestimating the level of effort and the timeline of deployment for a quantum network.

Projects can go from not even having requirements for the network gathered, to having a local area quantum network up and running and equipped with Quantum Secure Communications in under a year.

4. Quantum networks are single-purpose networks and can only perform one application.

This misconception is half true: QKD networks are single-purpose quantum networks: they can only be used for generating and distributing quantum keys for secure communication. However, entanglement-based quantum networks are multipurpose networks, capable of simultaneously enabling many quantum use cases and applications

Closing

Entanglement-based quantum networks are being built today by a variety of organizations for a variety of use cases – benefiting organizations internally, as well as providing great value to an organization’s customers. Telecommunications companies, national research labs, intelligence organizations, and systems integrators are just a few examples of the organizations Aliro is helping to leverage the capabilities of quantum networking.

Building entanglement-based quantum networks is no easy task. It requires:

- Emerging hardware components necessary to build the quantum network.
- The software necessary to design, simulate, run, and manage the quantum network.
- A team with expertise in quantum physics and classical networking.
- Years of hard work and development.

This may seem overwhelming, but Aliro is uniquely positioned to help you build your quantum network. The steps you can take to ensure your organization is meeting the challenges and leveraging the benefits of the quantum revolution are part of a clear, unified solution already at work in quantum networks like the EPB Quantum NetworkSM powered by Qubitekk in Chattanooga, Tennessee.

AliroNetTM, the world’s first full-stack entanglement-based quantum network solution, consists of the software and services necessary to ensure customers will fully meet their quantum networking goals. Each component within AliroNetTM is built from the ground up to be compatible and optimal with quantum networks of any scale and architecture. AliroNetTM is used to simulate, design, run, and manage quantum networks as well as test, verify, and optimize quantum hardware for network performance. AliroNetTM leverages the expertise of Aliro personnel in order to ensure that customers get the most value out of the software and their investment.

Depending on where customers are in their quantum networking journeys, AliroNetTM is available in three modes that create a clear path toward building full-scale entanglement-based secure networks: (1) Emulation Mode, for emulating, designing, and validating quantum networks, (2) Pilot Mode for implementing a small-scale quantum network testbed, and (3) Deployment Mode for scaling quantum networks and integrating end-to-end applications. AliroNetTM has been developed by a team of world-class experts in quantum physics and classical networking.

To get started on your Quantum Networking journey, reach out to the Aliro team. info@alirotech.com

References

[NIST] Fortier, Tara. “Demystifying Quantum: It's Here, There and Everywhere.” *National Institute of Standards and Technology*

<https://www.nist.gov/blogs/taking-measure/demystifying-quantum-its-here-there-and-everywhere>.

[NETWORK-STACK] Pompili, M. & Donne, C. & Raa, I. & Vecht, B. & Skrzypczyk, M. & Ferreira, G. & Kluijver, L. & Stolk, A. & Hermans, S. & Pawelczak, P. & Kozłowski, W. & Hanson, R. & Wehner, S.. (2022). Experimental demonstration of entanglement delivery using a quantum network stack. *npj Quantum Information*. 8. 10.1038/s41534-022-00631-2.

https://www.researchgate.net/publication/364542247_Experimental_demonstration_of_entanglement_delivery_using_a_quantum_network_stack)