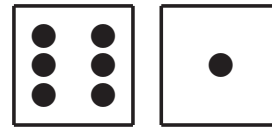

Quantum Communications

"Quantum communications" refers to two related applications: first, the use of quantum states to ensure true randomness in number selection and to communicate encryption keys to other parties, known respectively as quantum random number generation and quantum key distribution; second, the use of quantum effects themselves, such as the spin of photons, to encode a message, which is known as quantum internet or quantum networking.

There are four reasons to be excited by quantum communications and all are strategically relevant:

1. Properly implemented, quantum communications applications enjoy *information-theoretic security*, which means that no adversary, regardless of their computing resources or background knowledge, can decipher communications that have been covertly intercepted. Not even a quantum computer can decrypt such communications! This is because the security is a property of the underlying mathematics and quantum physics, rather than the putative "hardness" of a particular math problem.

   Quantum security guarantees to protect institutions against the future. Those continuing to use *computationally secure* post-quantum classical alternatives for distributing their keys rely on assumptions that may be proven incorrect. For instance, a mathematician may discover a new algorithm that unscrambles post-quantum encryption.

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

2. Quantum communications systems, unlike classical ones, reveal when a communication has been intercepted. That interception could be a surveilor, or it might be ordinary environmental interference, such as electronic noise or malfunctioning hardware. (Users of such systems typically cannot determine if the message failure was an accident of the environment or the actual presence of an eavesdropper.) The detection of interception capability results from the nature of quantum states. The act of interception interferes with quantum states, and this interference can be detected, unlike in classical communications, where interception is both easy and stealthy.

   For this reason, properly implemented quantum communications systems are not susceptible to proxying attacks. (You may also see these attacks referred to as "machine-in-the-middle" or "man-in-the-middle" attacks.) That's because if an attacker does intercept a photon carrying a particular quantum state, it is impossible for the attacker to both measure the photon's quantum state and retransmit a photon with the same quantum state.

3. In a fully quantum network that uses quantum states themselves to communicate, communication security becomes end-to-end. Users no longer have to rely on network trust, and can shut out eavesdroppers from both the content of their communications *and the metadata about those conversations*. Because governments extensively use metadata to study adversaries, this metadata-denying affordance of quantum internet schemes may be what is driving quantum network investments in Europe and China.

4. Just as Grover's algorithm speeds up some kinds of computations when performed on a quantum computer, some kinds of multi-party mathematical protocols enjoy a similar speedup when the parties communicate over a quantum network.

These benefits of quantum communications – information-theoretic security, awareness of message interception, the possibility of metadata secrecy, and certain kinds of optimizations – are driving both interest in quantum communications and its early commercializa-

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

tion. Indeed, the first quantum key distribution systems reached the market in 2005.[1]

Although quantum communication was discovered before quantum computing, another way to think about quantum communications systems is as a quantum computer with a "flying qubit" that travels from one party to the second, or with two flying qubits that travel from a common sender to two different receiving parties.

Quantum communications builds upon the technologies of quantum sensing discussed in Chapter 2, including single-photon detectors, the ability to perform low-noise measurements of quantum states, and even superconducting quantum devices.[2]

This chapter sets the stage for interest in quantum communications by briefly explaining the rise of signals intelligence (SIGINT) (Section 7.2 (p. 264)) capabilities of governments and the proliferation of these powers to nongovernmental actors. SIGINT is information derived from communications systems, radars, and weapons systems.[3] The chapter continues by explaining three quantum communications technologies, all of which can contribute to the confidentiality and integrity of communications.

First, quantum random number generation techniques use quantum uncertainty to create truly random numbers. Computer systems use high-quality random numbers in security, in simulations, and statistical models.

Second, quantum key distribution techniques use randomness to make secure encryption keys and ensure their confidentiality and integrity when they are transmitted to multiple parties. Although these protocols are called *quantum* key distribution, they are ultimately used to secure *classical* communications, for instance over the regular Internet or even the telephone.

Finally, a quantum internet would preserve quantum states and allow quantum computation between parties in different physical locations – possibly over great distances. This would provide both security against interception and secrecy of metadata. If the quantum

[1] Garfinkel, "Quantum Physics to The Rescue: Cryptographic Systems Can Be Cracked. And People Make Mistakes. Take Those Two Factors out of The Equation, and You Have Quantum Cryptography and a New Way to Protect Your Data" (2005).

[2] Takemoto et al., "Quantum Key Distribution Over 120 km Using Ultrahigh Purity Single-Photon Source and Superconducting Single-Photon Detectors" (2015).

[3] Director of National Intelligence, "What Is Intelligence?" (2019).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

networking necessary to achieve the ideal of a quantum internet were achieved, one could likely use the technology to connect disparate, small quantum devices into a larger cluster computer, or connect multiple quantum computers together to create a larger quantum computer.

7.1 Information-Theoretic Security

To understand the power of information-theoretic security is to understand the sublime attraction of quantum methods for protecting communications. Because many readers will not be familiar with the concept of information-theoretic security, we present below three math problems: one that is easy, one that was hard in 1977 when it was posed but was solved in 1994, and one that is information-theoretic secure, which means that it cannot be solved with the information that we present, even by an attacker who has unlimited computer power.

7.1.1 An Easy Math Problem

Here is an easy math problem. The variables p and q are positive integers and p is less than q ($p < q$).

$$p \times q = 15 \tag{1}$$

That is, what two numbers multiplied by each other equal 15? The answer is 3 and 5. This is an easy problem.

Recall that 15 is the number factored by IBM's quantum computer in 2001 (Section 5.2 (p. 188)). A simple way to think about this problem is to imagine that you have 15 cubes in a single line and you want to arrange them into a rectangle. If you did that, what would the dimensions of that rectangle be?

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

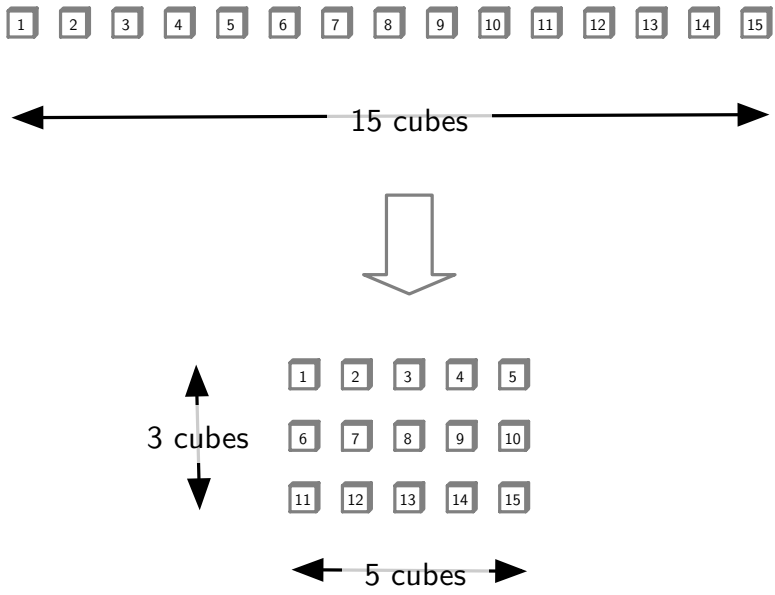

It turns out that there is only one way to make that rectangle, and that's with three rows of five cubes each.[4]

7.1.2 A Hard Math Problem

Here is a math problem that was posed in 1977 but was not solved until 1991, when it was cracked by an international team of 600 volunteers using more than a thousand computers. Instead of trying to factor the 2-digit number 15, try to break this number down to its prime factors p and q:

$$\begin{aligned} p \times q = {} & 114381625757888867669235779976146612010218 2 \\ & 96721242362562561842935706935245733897830 59 \\ & 71235639587050589890751475992900268795435 41 \end{aligned} \tag{2}$$

This 129-digit number is called RSA-129. It was chosen by Ron Rivest in 1977 as a puzzle to accompany the publication of a Martin Gardner column in *Scientific American*.[5] Like the number 15 in

[4] Turning the rectangle 90° so that it's five rows of three cubes each doesn't count as another "way" in this situation, because we required that the first factor be less than the second.

[5] Gardner, "Mathematical Games: A New Kind of Cipher That Would Take Millions of Years to Break" (1977).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

equation 1, RSA-129 has two factors, here called p and q.[6] But what are p and q in this case? That was the problem posed by Rivest.

RSA-129 has a curious property: if you factor the number into its two primes, you can use the result to decrypt a secret message that Rivest wrote and encrypted back in 1977.

Factoring RSA-129 was computationally infeasible in 1977. Rivest didn't know how long it would be until computers were fast enough that it would be feasible. Gardner's column claims that Rivest estimated it would take "40 quadrillion years" to factor such a number. But that estimate was based on a single 1977 computer running with the best factoring algorithm of the day: in the following years computers got faster, factoring algorithms got better; it also became possible to connect many computers together to work on the same number at the same time. This is what we mean when we say that factoring RSA-129 was *computational infeasible* in 1977, or alternatively, that RSA-129 was *computationally secure* then. Finding the factors of RSA-129 is left as an exercise for the reader.

7.1.3 An Impossible Math Problem

Now here is a math problem that you can't solve no matter how much computational power you have:

> There is a line that passes through the points (x_1, y_1) and (x_2, y_2). Find the value of y where the line passes through the y-axis (that is, when $x = 0$), given that one of the points is (3,5).

That is, solve for y in this equation given $x = 0$, knowing that $x_1 = 3$ and $y_1 = 5$:

$$y = mx + b \tag{3}$$

This equation can't be solved to give a unique solution for y: you aren't provided with enough information. The equation $y = mx + b$ describes a line on a graph, where m is the slope of the line and b is y-intercept. It's the y-intercept that you are trying to find. You can't find the y-intercept because you only have one point on the graph.

[6] Mathematicians frequently reuse variable names like p and q in different equations, just as lawyers reuse labels like "plaintiff," "defendant," and "the Court" in different lawsuits.

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

This is an example of a problem that is information-theoretic secure (see the sidebar "Secret Sharing" on page 266).

Today nearly every use of encryption on the planet is protected using ciphers that are computationally secure. As we saw in Chapter 5, these algorithms can be cracked simply by trying every possible decryption key and recognizing the message when it is properly decrypted. Quantum computers promise to make this process faster. Even post-quantum encryption algorithms are still merely computationally secure: we know that with enough computer power, these algorithms can be cracked. There might also be short-cuts to cracking these algorithms that haven't yet been discovered, just as better approaches for factoring were discovered after 1977 that made it easier to factor RSA-129.

Adopters of a properly implemented quantum encryption system do not have to rely on *computationally secure* algorithms for distributing their keys. Instead, they use qubits, safe with the knowledge that if the qubits are intercepted by an adversary, then the legitimate sender and recipient will be able to determine this fact.

There are actually two ways to use quantum cryptography, one that is secure given what we know about quantum computers today, and a second that is secure given our understanding of quantum physics and the physical laws of the universe:

1. With **Quantum Key Exchange**, flying qubits are used to exchange an encryption key that is then used with a conventional quantum-resistant symmetric encryption algorithm, such as AES-256. Because we believe that AES-256 cannot be cracked on a quantum computer, this approach is believed to be secure for the foreseeable future. That is, the key exchange is information-theoretic secure, but the bulk encryption is only computationally secure.[7]

2. With **Quantum networking** or **"quantum internet,"** flying qubits are used to exchange *all* of the information end-to-end between the parties. This approach is information-theoretic

[7]Note that AES-256 is only computationally secure against our current notions of quantum computing. It might not be secure against a computer based on quantum gravity, or strange matter, multiverse computation, or some kind of physics that we haven't yet imagined. Specifically, it might not be secure against a device that could solve NP-hard problems in polynomial time.

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

secure if the laws of quantum computing are correct. Put another way, it is secure as long as it is impossible to predict the future with absolute accuracy.

7.2 Golden Ages: SIGINT and Encryption Adoption

Signals Intelligence is one of the oldest intelligence gathering disciplines (Table 7.1). Many histories of SIGINT start with the use of wireless during World War I by both German and Allied forces: radio offered the advantage of instantaneous communications to troops in the field, potentially anywhere in the world, but suffered from risk that the enemy could be privy to the communications as well. Radio was too powerful to ignore, but too dangerous to use without some mechanism for protecting communications. Military users resolved this conflict by turning to encryption.[8]

In recent years events surely have altered the balance between those who wish to eavesdrop on communications and those who wish to keep their communications private. However, there is no clear accounting as to which side is now ahead.

7.2.1 The Golden Age of SIGINT

On the SIGINT side, many governments have developed audacious, comprehensive, systematic programs to capture communications and personal data in order to identify people, to attribute actions to parties and adversaries, to perform link analysis (the evaluation of relationships among people, adversaries, and others), and to capture communications content. For instance, it is alleged that in 2011 the Iranian government used compromised encryption certificates to access the email accounts of hundreds of thousands of Iranians who used Google's Gmail.[9]

In recent years, there have been repeated accounts in the US media of both Chinese and Russian successes in exfiltrating data

[8] In fact, the use of both encryption and cryptanalysis by militaries predates the invention of radio by at least 2500 years. For a history of code-making and code-breaking, we recommend David Kahn's updated classic (Kahn, *The Codebreakers: The Comprehensive History of Secret Communication From Ancient Times to The Internet* (1996)) as well as the more manageable (Singh, *The Code Book: The Science of Secrecy From Ancient Egypt to Quantum Cryptography* (2000)). For a contemporaneous account of code-breaking during World War I, we recommend Yardley, *The American Black Chamber* (1931).

[9] Hoogstraaten et al., *Black Tulip Report of The Investigation into The DigiNotar Certificate Authority Breach* (2012).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

Table 7.1. A sampling of the intelligence gathering disciplines (Director of National Intelligence, "What Is Intelligence?" (2019)).

GEOINT Geospatial Intelligence Gathered from satellite, aerial photography, and maps.

HUMINT Human Intelligence Gathered from a person. Includes diplomatic reporting, espionage, interrogation, traveler debriefing, and other activities.

IMINT Imagery Intelligence Analysis of images for their intelligence value. The National Geospatial-Intelligence Agency has primary responsibility for IMINT.

MASINT Measurement and Signature Intelligence Intelligence typically reviewed through the use of scientific measurement instruments. The Defense Intelligence Agency has primary responsibility for MASINT.

OSINT Open-Source Intelligence Analysis of information sources that are generally available, including news media and social media. The Director of National Intelligence's Open Source Center and the National Air and Space Intelligence Center are major contributors to OSINT.

SIGINT Signals Intelligence Intelligence gathered by analyzing "signals," which may include the analysis of intentional communications (COMINT – communications intelligence) and analysis of unintentional electronic emanations (ELINT – electronic intelligence). "The National Security Agency is responsible for collecting, processing and reporting SIGINT."

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

Secret Sharing

Secret sharing is an information-theoretic approach to splitting a secret into multiple parts. Invented independently in 1977 by G. R. Blakley[a] and Adi Shamir,[b] one primary use of secret sharing is splitting cryptographic keys used for data backups. Doing this renders the backup unusable unless multiple parties receiving the secret shares get together and reassemble the secret, allowing the backup to be decrypted.

Secret sharing works by representing the secret as a mathematical function that cannot be solved with the information present alone in each of the shares. In the example below, the secret is the y-intercept, which is where the straight line crosses the Y axis. Each share is a point on the line. Two points uniquely define a line, so without a second share, there is no way to identify the y-intercept.

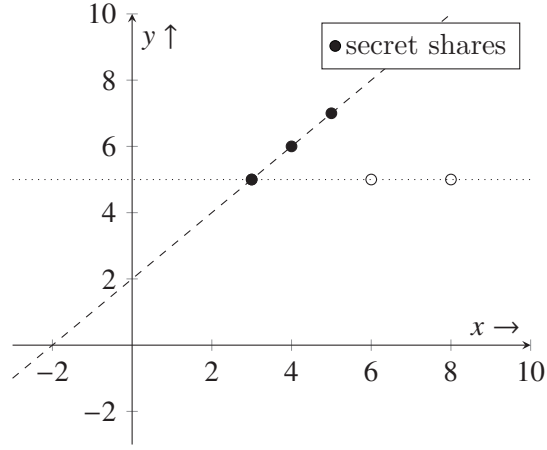

Here we see an example of secret sharing at work. The secret is $y = 2$ (the dashed line). The shares are $x_1, y_1 = (3, 5)$, $x_2, y_2 = (4, 6)$ and $x_3, y_3 = (5, 7)$. Combining any two secrets allows reconstructing the line. Notice that if the shares had been $(3, 5)$, $(6, 5)$ and $(8, 5)$, then the secret would have been $y = 5$. Thus, there is no way for a person receiving the share of $(3, 5)$ to know the value of the secret without combining their share with a share that someone else received.

[a] Blakley, "Safeguarding Cryptographic Keys" (1979).
[b] Adi Shamir, "How to Share a Secret" (1979).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

from both public and private US information systems. With respect to China, the breach of the US Office of Personnel Management database resulted in the theft of records on more than 20 million current and past federal employees, including fingerprint records and lengthy, detailed forms used when applying for a security clearance. Chinese hackers are also reported to have stolen the credit reports on over a hundred million Americans. Between these two attacks, China can presumably identify and target people who are both likely involved in intelligence efforts and who are economically vulnerable. This data surveillance has real consequences for US efforts and is believed to have enabled China to identify multiple CIA assets in Africa.[10] Turning to Russia, the former superpower has many satellites, terrestrial assets, and near-shore submarines, all of which can be used for collection of SIGINT. At the end of 2020, the US intelligence stated that a supply chain attack on the US company Solar Winds, which makes software to help organizations monitor their computer systems, was "likely Russian in origin."[11] More than ten thousand US companies and government agencies were compromised as a result of the attack.

Books and reports that synthesize government programs into single readings, like Barton Gellman's *Dark Mirror*,[12] can seem like paranoid science fiction. In that book, for instance, Edward Snowden refuses to reveal whether he has a blender, for fear that the appliance's electrical signal would reveal his location to intelligence agencies. There is no way to know from public sources if Snowden's fears are justified. But we do know that in 2014 a smart refrigerator was taken over by hackers and used to send spam,[13] and that in 2019 the FBI's Oregon office warned that hackers can take over the microphones and cameras in smart TVs and use them for surveillance.[14] More recently, *New York Times* cybersecurity reporter Nicole Perlroth published the bestseller *This Is How They Tell Me the World*

[10] Zach, "China Used Stolen Data to Expose CIA Operatives in Africa and Europe" (2020).

[11] Cybersecurity and Infrastructure Security Agency, "Joint Statement by The Federal Bureau of Investigation (FBI), The Cybersecurity and Infrastructure Security Agency (CISA), The Office of The Director of National Intelligence (ODNI), and The National Security Agency (NSA)" (2021).

[12] Gellman, *Dark Mirror: Edward Snowden and The American Surveillance State* (2020).

[13] Starr, "Fridge Caught Sending Spam Emails in Botnet Attack" (2014).

[14] Steele, "Oregon FBI Tech Tuesday: Securing Smart TVs" (2019).

Ends which details decades of offensive hacking efforts by China, Iran, Israel, North Korea, Russia, and the US to access information and booby-trap information protection systems.[15]

Peter Swire, who served under two presidential administrations and was responsible for reviewing intelligence community activities after the Snowden documents were dumped, argues that we live in "The Golden Age of Surveillance."[16] Not only do nation states like China, Russia, and the US have well-funded institutions with technically gifted employees searching for new ways to monitor, but important other factors have also begun to enhance surveillance powers.

As information traverses the Internet, operators of servers can log *metadata* about activity. US law currently makes it much easier for law enforcement to obtain metadata than content. Perhaps this is because the content/metadata distinction was in part driven from the days when a telephone's content was recorded with a pair of alligator clips onto a reel-to-reel tape recorder and metadata was captured with a *dialed number recorder* that literally recovered the numbers that a person dialed *and nothing else*.

Metadata is commonly believed to be less sensitive than content. However, there is a good argument to be made that metadata is more revealing than content. Metadata is easier to structure in computer databases and analyze. Consider the act of watching and interacting with a YouTube video. The *content* of the session includes:

- The visual content of the video, including the individual frames, the images of the people in the frames, the images of the buildings, etc.
- The audio content of the video, including the sounds, music, and other information.
- The text of any comments left on the video.

But if you were an analyst, consider the knowledge that could be derived from the same video's metadata:

- The video's unique identifier and its title.
- The time that the video was recorded, uploaded, and edited.

[15]Perlroth, *This Is How They Tell Me The World Ends: The Cyberweapons Arms Race* (2021).

[16]Swire, "The Golden Age of Surveillance" (2015).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

- The unique identifiers of each person that watched the video, their geographic location, their internet protocol (IP) address, and the time that it was watched.
- Whether the viewers clicked "thumbs up" or "thumbs down" on the video.
- Whether the viewers shared the video with friends and, if so, whom.
- The identifiers of any individuals in the video found with face recognition software.

The additional information available from metadata – particularly surrounding the identity of the community of users interested in the video and the people to whom they send it, might be far more important than the video's actual content.

The lines between content and metadata are not sharp. A transcript of the video might be considered content, but keywords extracted from the transcript might be considered metadata. While we classify the comments as content, the timings between individual keystrokes when the comments were left might be considered metadata – even if software can recover the actual typed words using those timings.

Metadata can thus indicate location, the identities of friends, and provide many hints about the content of communications and actual activities online. In many cases, the metadata/content distinction is functionally irrelevant, because operators of servers and services directly examine the content of our email, photographs, and other communications in the dual interests of security (anti-spam) and commercialization (behavioral-based advertising). The private sector plays a critical role by assembling dossiers of both proprietary company data and open source information on people; such products can then be sold to both marketers and (even foreign) government agencies.

The move to the "cloud" means that governments can obtain troves of data about people that previously would have been confined to a home or a business with legal process (or simply by guessing or otherwise obtaining the user's password). Individual users of technology also contribute to surveillance power by documenting their lives on social networks, and by carrying mobile trackers and dutifully

storing contact books in them, which give companies and intelligence agencies alike access to location data and fodder for link analysis.

As much as technological trends have benefited nation states, these capabilities have devolved to many private sector actors as well.[17]

Especially concerning to some is the use of state collection capabilities to support domestic industries and silence critics living abroad. In the 1990s, for example, France was accused of using its intelligence apparatus to spy against Boeing, Textron, and Bell.[18] More recently businesses have raised concerns about intellectual property exfiltration by China, which then shares the information with commercial rivals in China. Businesses are concerned about China and other nations using a range of surveillance capabilities to collect information on dissidents, regime critics, and refugees who live outside of the country. For example, in 2010 Google revealed that its Gmail system had been hacked by China and that information from the email accounts of human rights activists had been pilfered.[19] Businesses are also concerned about the convergence of organized crime and government in Russia, which not only directly engages in financial fraud but also creates platforms and even a market for others to do so.[20]

7.2.2 The Golden Age of Encryption

The Golden Age of Surveillance is accompanied by a corresponding golden age of encryption *adoption by default*. Since 1991, users with significant technical ability have been able to use strong encryption in the form of Phil Zimmerman's Pretty Good Privacy,[21] although even later versions that were heralded as being easy to use were

[17] Weinbaum et al., *SIGINT for Anyone: The Growing Availability of Signals Intelligence in The Public Domain* (2017); Koller, *The Future of Ubiquitous, Realtime Intelligence: A GEOINT Singularity* (2019).

[18] Doyle, "Business Spy War Erupts between US and France: Paris Forced to Come Clean on Hi-Tech Dirty Tricks" (1993); Greve, "Boeing Called A Target Of French Spy Effort" (1993).

[19] Zetter, "Google to Stop Censoring Search Results in China After Hack Attack" (2018).

[20] Organized Crime and Corruption Reporting Project, "The Russian Laundromat Exposed" (2017); US Agency for International Development, Bureau for Africa, "Government Complicity in Organized Crime" (2019).

[21] Garfinkel, *PGP: Pretty Good Privacy* (1994).

still too difficult for most people.[22] Since then, technologists have sought to change the security landscape by implementing encryption by default in seamless ways. Perhaps most notable is the shift of addresses on the World Wide Web from being prefixed by `http://` to `https://`, which provides users greater confidentiality and integrity in their web browsing. Prior to this change, users' web browsing was sent over the Internet without encryption, allowing adversaries and telecommunications providers alike to monitor users' website visits or even change the content of web pages as they were being viewed.[23] Email likewise has moved from communications where most messages sent over the Internet backbone were sent entirely in plain-text to a system where such messages are largely encrypted (although email encryption is not generally end-to-end – see "Is Your Email Encrypted?" on page 272). Likewise, the popular messaging app WhatsApp offers end-to-end encryption. When WhatsApp was acquired by Facebook, the creators left to support Signal, another messaging application offering end-to-end encryption. Likewise, Apple's iPhone and its newest laptops and desktops use encryption for storage and for text messages sent between Apple users. Although such techniques can be defeated through the use of so-called 0-day attacks,[24] companies like Apple are typically quick to fix such vulnerabilities when they become public.

Central to this rise in encryption is that the user need not understand, configure, or even activate it because encryption is on by default. This offers a lesson for the confidentiality and integrity gains possible in quantum communications: for these innovations to be realized, they must not only be easy to use, they must be secure and integrated into the fabric of communications systems and consumer-facing applications.

7.3 Quantum Random Number Generation (QRNG)

All of these encryption systems we discussed in the last section are based on more-or-less the same technology stack: the AES encryption algorithm to encrypt the messages, a secure random number

[22] Whitten and Tygar, "Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0" (1999).

[23] The advent of free encryption certificate services and a policy from Google that sites with TLS would get higher rankings in search results caused a rush to adopt the `https://` prefix.

[24] Perlroth, *This Is How They Tell Me The World Ends: The Cyberweapons Arms Race* (2021).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

Is Your Email Encrypted?

Much email sent today is between two Gmail users. These messages are encrypted by the Transport Layer Security (TLS) as they travel from the sender's web browser to Google's web-mail service. Although the messages are not encrypted in the memory of Google's servers, they are encrypted when they are written to Google's disks where the messages are stored.[a] Likewise, the email messages are encrypted when they are sent from Google's servers to the Gmail recipient.

Mail that gets sent from Gmail to other mail providers, such as Microsoft's Office 365 cloud platform, are frequently encrypted using the SMTP STARTTLS protocol.[b]

This kind of protection is not as strong as the so-called *end-to-end* encryption offered by the S/MIME and PGP encryption systems. However, STARTTLS is significantly easier to use because each user does not need to create or otherwise obtain a public/private keypair.

[a] Google LLC, "Encryption at Rest" (2021).
[b] Rose et al., *Trustworthy Email* (2019).

generator to create the AES key, and public key cryptography to get the per-message key from the message sender to the recipient. Earlier in this book we discussed the role of the AES and public key cryptography algorithms. In this section we will discuss the role of random numbers.

Cryptography depends on strong random numbers. For instance, a RSA-2048 key is generated from prime numbers that are over 300 digits long: these prime numbers are found by guessing random numbers and checking them to see if they are prime. (Unlike factoring, there are mathematical tricks that are used to rapidly determine if a number is prime or not.) Likewise, the AES-256 keys are themselves random numbers.

Random numbers thus form the very basis of the security provided by encryption. If a 256-bit key is random, then that means every key is equally probable. But if an attacker can somehow interfere with the randomness of the number generation process, it can dramatically reduce the possible number of encryption keys. For such an attack, the strength of AES-256 with a key that is not very random might not be strong at all.

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

The NIST Randomness Beacon

In 2013, the US National Institute of Standards and Technology deployed its "Randomness Beacon," a web-based service that posted random numbers in blocks of 512 bits every minute. Like an electronic lottery machine, the bits posted to the NIST website are unpredictable.

The randomness service is an endless source of numbers that can be used in situations where a random choice needs to be made, and the person making the choice wants to demonstrate that they made the choice fairly. In football games, for example, the receiving team is chosen by a coin toss – but how do we know the coin is fair? In this and similar situations where a decision must be made on a random choice, the NIST service can be relied upon by both parties to ensure a selection that is unbiased.

Example applications that NIST proposed included selection for random screening at security checkpoints, selection of test and control groups in scientific trials, selection of people for random tax audits, assignment of judges to cases, and so forth. Because the beacon is public, and because each bitsream is added to a hash chain (or blockchain), the system can be audited by any party. Of course, being public comes with a risk as well: the bits should not be used in cases were both randomness and secrecy are required. To drive in this lesson, the NIST website states:[a]

WARNING:
DO NOT USE BEACON GENERATED VALUES
AS SECRET CRYPTOGRAPHIC KEYS.

[a] See beacon.nist.gov/home

Modern computers generate random numbers by using an initial *random seed* which is then used with a deterministic random bit generator, also called a pseudo-random number generator (PRNG). Typically, the random seed is created by combining many events that, if not completely random, are at least unpredictable. For example, the early PGP program instructed users to type on the keyboard and used the inter-character timing as a source of randomness. Other

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

sources of randomness include the arrival time of packets at a network interface, inputs to digital cameras, and even seismic sensors. In practice, the quality of random numbers is determined by the samples taken from the "random" source, the quality of the mixing, and the quality of the PRNG. If any of these produce output that is somewhat predictable, or for which there is correlation between successive values, then a knowledgeable adversary can gain advantage when attempting to decrypt a message that was encrypted with such "poor quality" randomness.

Concerns about the strength of random number generators has been raised many times in the past. One such case from the US involves the Dual Elliptic Curve Deterministic Random Bit Generator (Dual_EC_DRBG).[25] When Dual_EC_DRBG was proposed, security professional Bruce Schneier and others raised concerns that the algorithm might include a "secret backdoor" that would allow the US government to predict the algorithm's "random" outputs.[26] These concerns were confirmed in 2013.[27] Following the disclosure, NIST issued guidance stating "NIST strongly recommends that, pending the resolution of the security concerns and the re-issuance of SP 800-90A, the Dual_EC_DRBG, as specified in the January 2012 version of SP 800-90A, no longer be used."[28] In 2015, the Director of Research at the National Security Agency said that the agency's "failure to drop support for the Dual_EC_DRBG" after vulnerabilities were identified in 2007 was "regrettable."[29,30]

In 2019 cryptographers stated that two Russian-designed encryption systems, Streebog and Kuznyechik, might also contain a secret backdoor that would give an advantage to a knowledgeable attacker trying to decrypt a message protected with the algorithm. In this

[25] Barker and Kelsey, *Recommendation for Random Number Generation Using Deterministic Random Bit Generators (Revised)* (2007).

[26] Schneier, "Did NSA Put a Secret Backdoor in New Encryption Standard?" (2007).

[27] Perlroth, "Government Announces Steps to Restore Confidence on Encryption Standards" (2013); Buchanan, *The Hacker and The State: Cyber Attacks and The New Normal of Geopolitics* (2020).

[28] Information Technology Laboratory, "Supplemental ITL Bulletin for September 2013" (2013).

[29] Wertheimer, "Encryption and The NSA Role in International Standards" (2015).

[30] This story and others surrounding the quest to produce high-quality random numbers at scale is discussed in Garfinkel and Leclerc, "Randomness Concerns When Deploying Differential Privacy" (2020), from which this story and its references are taken.

case, the weakness was not in the random number generator, but in the algorithms' so-called "substitution boxes."[31]

Quantum states provide the best source for strong, unbiased randomness. Scientists have developed several different methods to derive strong randomness from quantum events, including the path that photons take when light is split, the polarization of individual photons, and the phase of quantum states and processes.[32] A notional device bears similarity to the dual-slit experiment discussed in Section B.1.3, "Light: It Acts Like a Wave" (p. 490). The device works by cycling a particle or photon in and out of superposition. Measurement disturbs the superposition, causing decoherence and the production of a random bit. That bit is then used as a basis to generate random numbers. One way to think of these machines is as a quantum computer with a single qubit that is constantly computing the answer to the question "is the qubit 0 or 1?"

Number generation in such a scheme faces two sets of challenges. The first is the cycle speed of the prepare-superposition process and the speed of the measurement-decoherence process, which together determines how fast these systems can produce random bits. These machines may also be impacted by errors produced by classical noise and the reliability and tolerances of the quantum source and of the measurement mechanism, which can bias the results.

Properly implemented, QRNG produces strong randomness.[33] In fact, it probably produces the strongest possible random numbers, since modern physics holds that quantum processes are the ultimate source of all nondeterminism that we observe in the universe. QRNG has also been commercially available for years. In fact, after scientists created a QRNG system at the Australian National University in 2011,[34] the investigators found they had more random numbers than they would ever need for experiments. So they created a free QRNG service on the web.[35] In 2020, IBM and Cambridge Quantum Computing offered QRNG as a cloud service. And NIST is deploy-

[31] Perrin, "Partitions in The S-Box of Streebog and Kuznyechik" (2019).
[32] X. Ma et al., "Quantum Random Number Generation" (2016).
[33] Acin and Masanes, "Certified Randomness in Quantum Physics" (2016); Bierhorst et al., "Experimentally Generated Randomness Certified by The Impossibility of Superluminal Signals" (2018).
[34] Symul, Assad, and Lam, "Real Time Demonstration of High Bitrate Quantum Random Number Generation with Coherent Laser Light" (2011).
[35] See qrng.anu.edu.au/

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

ing Entropy as a Service (EaaS), a public, quantum-based source of random numbers.

Using these remote, cloud-based services requires some reliance on the provider, but there are measures that can be taken to reduce the risk. Instead of using the source directly, it can be combined with a secret key and then used in a cryptographically strong PRNG – a CSPRNG! This approach works as long as the secret key is kept secret and as long the PRNG is really a CSPRNG. That's the use case that NIST envisions for its EaaS. The EaaS project is explicitly designed to serve Internet of Things (IoT) devices by providing random numbers that these devices can use to create strong encryption keys. The idea is that IoT devices will be small and inexpensive, so much so that even high-end brands will cut corners on security, thus the chances that the market will produce QRNG for IoT devices is particularly unlikely. NIST is in effect substituting the market with security fundamentals for anyone to use. NIST is also upgrading its Randomness Beacon to use QRNG, as currently it uses two classical generators to prevent guile.

Higher levels of assurance require implementing the QRNG locally, so that the high-quality random bits are generated where they are needed, and not by some third party. For instance, ID Quantique has long sold QRNG hardware that plugs into a standard personal computer or server. In 2020, the company announced a QRNG chip that could fit into mobile phone handsets.[36] This device uses the random "shot noise" from a light-emitting diode (LED) to generate numbers. Every time the LED fires, the number of photons emitted fluctuates randomly. A CMOS sensor array sensitive to single-photon events detects the number emitted and their positions (see Figure 7.1).

7.4 Quantum Key Distribution

When Rivest, Shamir, and Adleman wrote their article introducing the RSA encryption system, they explained it with a woman, "Alice," who wanted to send a secret message to a man named "Bob."[37] Since then, Alice, Bob, and a whole cast of other characters have been used to help scientists analyze and explain security protocols. There is Eve, the eavesdropper, who attempts to "intercept" (a strained metaphor)

[36] Quantique, "Quantis QRNG Chip" (2020).

[37] Ronald L. Rivest, Adi Shamir, and Len Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems" (1978).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

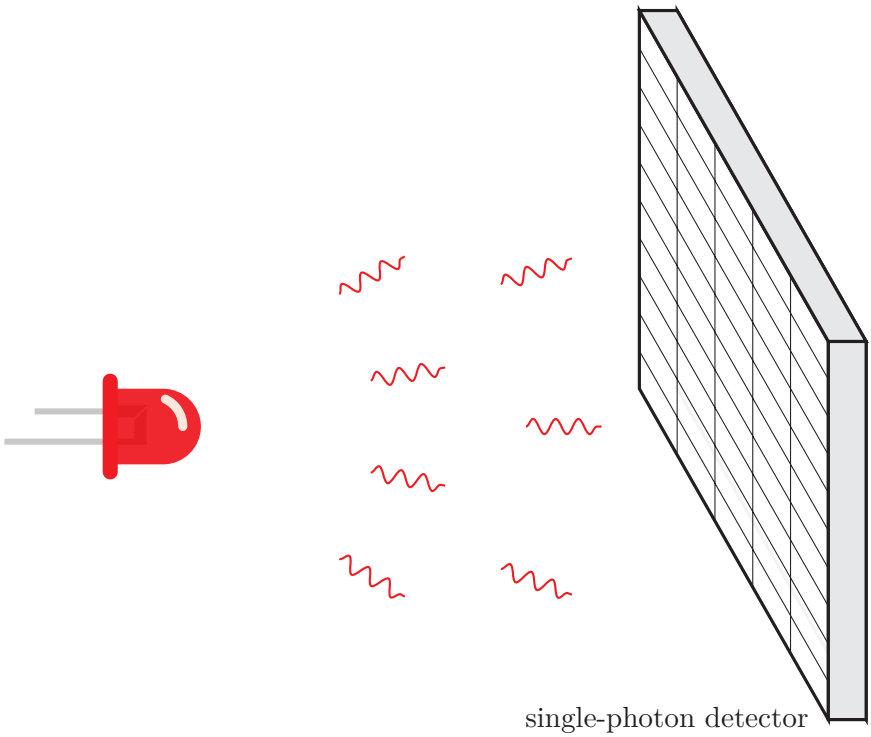

Figure 7.1. A mechanism for QRNG designed by ID Quantique fits into a mobile phone handset and pairs an LED and single-photon sensor array to derive randomness from photonic noise.

this conversation. And there is Mallory, a malicious attacker, who can modify the message or inject new ones.

Quantum Key Distribution (QKD) describes an approach where Alice and Bob can exchange an encryption key guaranteed to enjoy *unconditional* security. No computer available today or in the future can compromise this system, because the attacker does not have enough information to make sense of the ciphertext.

7.4.1 BB84

In 1984, Charles Bennett and Giles Brassard published the BB84 protocol, demonstrating how Alice and Bob could exchange encryp-

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

tion keys using quantum states.[38] Using the protocol, Alice and Bob get the same stream of `0` and `1` bits that they can use for any purpose. For example, they can use the sequence in 8-bit chunks as a *one-time pad* (see Figure 7.2), using each group of 8 bits to encrypt the next byte of the message. Alternatively, they can use the sequence in 256-bit chunks as AES-256 encryption keys.

The one-time pad is the gold standard for communications security because it is information-theoretic secure.[39] Even if the attacker tries every possible key, there is not enough information in the encrypted message to distinguish a correctly decrypted message from an incorrectly decrypted message. The reason is that the key is as long as the message, so every possible key makes the message decrypt a different way. This means that trying every possible key makes the encrypted message decrypt to every possible message.

One-time pads are the stuff of spy thrillers and history books, but they are not used much today because it is too difficult to distribute the pads in advance and then assure that each is used just once. The Soviet Union attempted to use one-time pads for its diplomatic communications after World War II and it failed; the NSA revealed its success in cracking the Soviet codes in 1995 (see Figure 7.6).[40]

BB84 is revolutionary, because Bennett and Brassard's approach deals with two central challenges in communication: how to generate a secure, shared secret, and how to distribute it at a distance. Two other key challenges – usability and the time it takes to generate and transmit the key securely – are up to the companies that create applications using QKD protocols.

However, modern QKD systems cannot generate a stream of bits fast enough to encrypt modern data links. For this reason, QKD systems typically operate in a slightly less secure mode in which BB84 is used to exchange 256-bit encryption keys which are then used with conventional encryption algorithms such as AES-256. With a 256-bit key, each encrypted message will have only 2^{256} possible decryptions, and the likelihood is that all but one of them will be gibberish. As we discussed in Chapter 5, it isn't possible to try all 2^{256} keys, so using BB84 to exchange AES-256 keys is considered secure. However, it is only computationally secure, not information-theoretic

[38] C. H. Bennett and G. Brassard, "Quantum Cryptography: Public Key Distribution and Coin Tossing" (1984).

[39] Shannon, *Communication Theory of Secrecy Systems* (1949).

[40] National Security Agency and Central Security Service, "VENONA" (2021).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

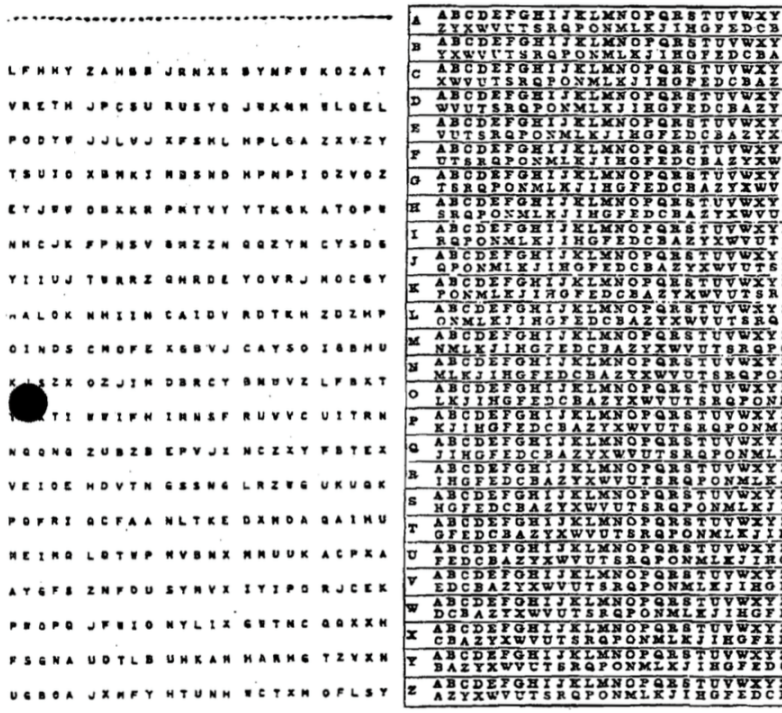

Figure 7.2. This table from the NSA's DIANA program illustrates how one-time pads produce messages with keys the same length of ciphertext. The key is on the left-hand side. The right-hand side is the table used to convert plain text to ciphertext (and vice versa). This key starts with the letter "L," so the user encrypting a message would use the L row on the table to choose the first letter of ciphertext. Assume that Alice wants to say "The Magic Words Are Squeamish Ossifrage" to Bob. To encrypt, Alice notes the first letter from the key, left-hand pane, which is L. Turning to the table, row L, and then to the letter T, the corresponding ciphertext underneath the T is a V. To encrypt the next letter, Alice would then use F from the key to locate the letter H and choose the ciphertext N, and so on. Alice and Bob must have identical cards and must destroy them after the process.

secure. As a compromise, these systems might change their AES-256 keys every few seconds, to minimize the amount of ciphertext that has been encrypted with any given AES-256 key.

7.4.2 How QKD Works

Most QKD systems are based on the idea of sending a stream of photons from a sender (Alice) to a recipient (Bob). For more background on polarized light, see Appendix B.3, "Quantum Effects 2: Polarization".

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

Here we provide a simplified explanation for how BB84 operates. The first thing to know is that actually using BB84 in a production system requires considerable mastery of the quantum realm and engineering cleverness not explained here.

In modern QKD systems, the photons either travel down a fiber-optic strand, or they are created in pairs in a satellite and sent to two independent ground stations.[41] In the first case, Alice prepares a stream of photons by sending each through a polarizing filter that is either polarized horizontally (H), vertically (V), at a 45° angle, or at a 135° angle. Alice makes this choice at random, recording both the number of the photon and the orientation of her polarizing filter. Sending with a H or a 45° is tentatively sending a `0`, while sending with a V or a 135° is tentatively sending a `1`. (Alice can't actually number each photon, so instead she will encode each photon's value in the light stream itself.)

Let's say Alice sends 10 photons:

Photon #	Alice Filter orientation	Tentative bit
0	45°	`0`
1	45°	`0`
2	45°	`0`
3	H	`0`
4	V	`1`
5	135°	`1`
6	45°	`0`
7	45°	`0`
8	H	`0`
9	135°	`1`

When Bob receives the photons, he also passes them through a filter that is also randomly oriented at either V or at 135°. He then measures the presence or absence of the photon with a single photon detector:

[41] The protocol involving a pair of entangled photons is called E91, after its inventor Artur Ekert (Ekert, "Quantum Cryptography Based on Bell's Theorem" (1991)).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

Photon #	Bob Filter orientation	Photon detected?	tentative bit
0	135°	NO	0
1	135°	NO	0
2	V	YES	1
3	V	NO	0
4	V	YES	1
5	V	YES	1
6	135°	NO	0
7	V	NO	0
8	135°	NO	0
9	V	YES	1

Now Alice and Bob need to compare notes to see if the measurement that Bob made of the photon was compatible with the photon that Alice prepared and sent. If Bob measured with his V filter, then he will detect light if Alice sent the light with her V filter, but not if she used her H filter. But if Alice sent it with her 45° or 135° filters, the measurement that Bob made is meaningless: there's a 50–50 chance that a photon polarized with the 45° filter will pass through a V filter.

To compare notes, Bob can reveal which filter he used to measure each photon. Alice then tells Bob which of his measurements he should keep and which he should throw out:

Photon #	Bob to Alice	Alice to Bob
0	135°	KEEP
1	135°	KEEP
2	V	–
3	V	KEEP
4	V	KEEP
5	V	–
6	135°	KEEP
7	V	–
8	135°	–
9	V	–

At this point, Alice and Bob know that photons 0, 1, 3, 4, and 6 were sent and received with compatible polarizing filters. Alice looks

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

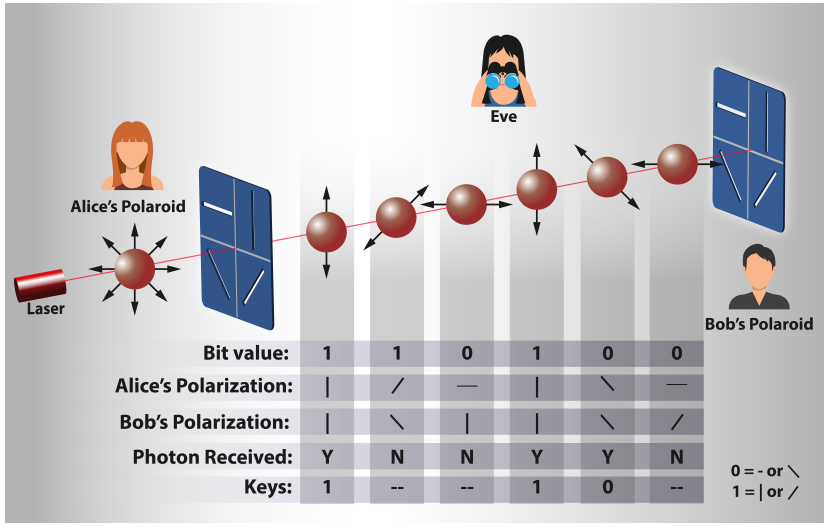

Figure 7.3. The BB84 protocol illustrated. Adapted from Aliberti and Bruen by Twitter user farooqumer89.

at her table and discovers that the tentative bits corresponding to those numbers are `0` `0` `0` `1` `0`. Bob looks at his table and gets the same sequence of bits.

To determine that the system is operating properly, Alice and Bob can now decide to reveal every even bit of the resulting sequence. Alice says that even bits are `0`, `0`, and `0`. Bob notes that his are the same. Alice and Bob then use the remaining bits (`0`, `1`) as their secret key.

If Alice and Bob do not reveal to each other the same bits, then either the system is not operating properly, or else an attacker is intercepting the beam and injecting a photon sequence of their own. In either case, Alice and Bob know not to use that key.

Because of measurement error, the sequence of bits that Alice and Bob recover are not exactly the same. A variety of error correction techniques exist that can be used to account for these errors, at the cost of using even more bits.

The two-photon system is similar, except that a pair of entangled photons are sent from the satellite to both Alice and Bob, who then both measure the polarization and compare notes. In this design, the satellite cannot determine the key that Alice and Bob agree upon, nor can anything else in the universe: each photon can only be measured

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

once. Of course, once Alice and Bob agree upon a key, a suitably skillful attacker might be able to steal it from either Alice or Bob if their QKD device does not properly protect the key after it has been created.

7.4.3 Why QKD Is Secure

What makes QKD secure is the fact that the actions of Alice and Bob measuring the photon are independent, but the measurements are correlated *if and only if Alice and Bob choose compatible measurements*. If Alice measures the photon with a horizontal polarizing filter and Bob uses a filter that is polarized vertically, their measured results are linked and they have now agreed on a common bit. But if Bob uses a filter at 45°, the measures are incompatible and there is no correlation between them. This is the essence of Einstein's "spooky action at a distance," the paradox of entanglement. Because Alice and Bob chose their measurements at random, only 50 percent of them will be compatible: the remaining measurements will be thrown out.

Now let's say an attacker, Eve, tries to crash the party. Eve attempts the well-known "man-in-the-middle" attack: she catches the photons headed for Bob, measures them, and then prepares a new photon and sends it to Bob. Can Eve get away with this deception? In a properly implemented QKD system, the answer is no. That's because when Eve receives, measures, and retransmits the photon, she doesn't know how Bob is going to measure it. By chance, she will only measure the photon in a compatible manner 50 percent of the time. The other 50 percent of the time, she will measure the photon in a way that is incompatible. When she sends each of those incorrectly measured photons to Bob, Eve has a 50 percent chance of sending them in the correct state, and 50 percent chance of sending them in the wrong state.

When Bob compares notes with Alice, they first reveal how the photons were measured and throw out the photons for which Alice's and Bob's measurements were incompatible. But after this step, they intentionally reveal a certain percentage of the remaining photons. When Bob and Alice discuss these intentionally revealed photons, they will discover that their measurements disagree roughly half of the time. This indicates either that their equipment is not working properly, or that Eve is attempting to perform a man-in-the-middle attack.

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

Quantum Computing and Bitcoin

Cryptocurrencies such as Bitcoin are speculative investment and value transfer mechanisms that are based on a *distributed ledger*, a kind of shared database, that is difficult to corrupt. Bitcoin, the first cryptocurrency, relies on SHA-256 to build its ledger.

The Bitcoin ledger consists of many transactions, each of which is basically an electronic check that is signed with a private key. The check transfers some amount of Bitcoin from the user's corresponding public key (a Bitcoin "address") to another public key. These transactions are grouped into blocks. In addition to these electronic checks, each block contains the hash of the previous block, a signature by the block's "miner," and a block of random values placed there by the miner. The random values are manipulated such that the SHA-256 hash of the new block begins with a large number of zeros. To do so, the Bitcoin "miner" takes the block of transactions and makes systematic changes to that random block until the hash has enough zeroes.

Because the hashes generated by SHA-256 appear random, with each bit having an equal chance of being a `0` or a `1`, finding hashes with a large number of leading zeros is computationally intensive. In March 2020, Bitcoin blocks had 76 leading binary `0`s, followed by 180 bits of `0`s and `1`s; the number of leading `0`s is automatically adjusted to be longer and longer as more and faster Bitcoin miners join the network; each additional leading `0` requires roughly twice as much computational power to find.

In 2019, the National Academies estimated that a large quantum computer could attack Bitcoin's ledger system but the attack requires 2403 qubits and 180 000 years. Given that the ledger gets a new block every 10 minutes, attacking the ledger itself in order to obtain free Bitcoin appears unlikely.

Bitcoin holders may still be vulnerable because a quantum computer could be tasked with cracking the public key of an individual Bitcoin user's wallet and then stealing that user's money. Alas, the victim would have little recourse owing to the social contract underlying cryptocurrencies.

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

Quantum Money

Stephen Wiesner's idea of using the entanglement of two particles to create unforgeable banknotes (see p. 137) led Bennett and Brassard to come up with the idea of quantum cryptography in the first place. Since then, many scientists have proposed systems that rely on quantum effects to store and transmit value, now broadly called *quantum money*. These schemes vary in their implementation. Some provide information-theoretic security while others rely on public key systems.[a] But given current constraints in quantum memory, computing, and networking, hopes for quantum money systems are far off.

If they ever do arrive, some of the affordances promised will be contested by parties with interests in transactions. Cryptocurrencies like Bitcoin and most if not all envisioned quantum currencies contain mechanisms to ensure that a purchaser actually has sufficient funds and to prevent "double spending." Beyond that, however, most of these mathematical monies are quite spartan.

Conventional value transfer mechanisms such as checks and credit cards are complex for many reasons. For instance, policy decisions must be made to reconcile the different, conflicting interests held by ordinary consumers, merchants, banks, and governments in payments. A consumer might want the ability to repudiate a value transfer, in case of fraud, coercion, or because of poor-quality goods received, while merchants might want to block repudiation. Governments typically want the ability to unmask all parties in a transaction. Such mechanisms are missing – intentionally – from cryptocurrencies.

Yet, as Bitcoin has become more mainstream, the original vision of a bank-free, anonymous, peer-to-peer payment system has ceded to something more akin to a commodities market, one mediated by exchanges that are regulated by governments and that follow taxation and anti-money-laundering rules to identify market participants.

[a] Hull et al., "Quantum Technology for Economists" (2020).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

Of course, Eve could go further, and pretend to be Bob to Alice and to be Alice to Bob. To prevent this, Alice and Bob need to have a way of authenticating the open messages that they send to each other. Today the easiest way to do this authentication is with public key cryptography. This use of public key cryptography is considered acceptable in QKD systems, because even if an attacker records the authentication messages and cracks the private keys behind them at some point in the future, that won't change the fact that the messages were properly authenticated when they were sent. No secret information is revealed if the authentication keys are cracked in the future.

Eve can prevent Alice and Bob from communicating securely by using electronic warfare approaches. Eve could inject noise to deny or degrade the quantum channel and cause Alice and Bob to have to revert to other, less secure communication, but she can't decipher the messages sent. (Indeed, risks of denial of service are among the reasons the NSA has spurned QKD in favor of quantum-resistant (or post-quantum) cryptography.[42]) And once the key is exchanged between Alice and Bob, the duo do not need a "quantum internet" or quantum states to talk securely. Alice and Bob can use the quantum key to communicate on existing classical channels, encrypting their communications with a conventional quantum-resistant symmetric algorithm such as AES-256.

7.4.4 QKD Gains Momentum

Since BB84 was proposed, new protocols and even implementations have emerged. For instance, in 1991, Arthur Ekert proposed the satellite entanglement protocol described above.[43] Recall that Alice and Bob receive correlated photons from a split-beam laser. Using Bell tests (see Section B.4, p. 513), Alice and Bob compare the correlations of their photons to ensure that Eve has not intercepted them. Under Ekert's proposal, even if Eve is operating the laser, she cannot determine the states of Alice and Bob's photons without interfering with the Bell correlations, thus revealing her attack. Ekert's proposal anticipates the possibility of a QKD-as-a-service approach – a satellite delivering entangled photons from space to the ground, allowing

[42] National Security Agency, "Quantum Key Distribution (QKD) and Quantum Cryptography (QC)" (2020).

[43] Ekert, "Quantum Cryptography Based on Bell's Theorem" (1991).

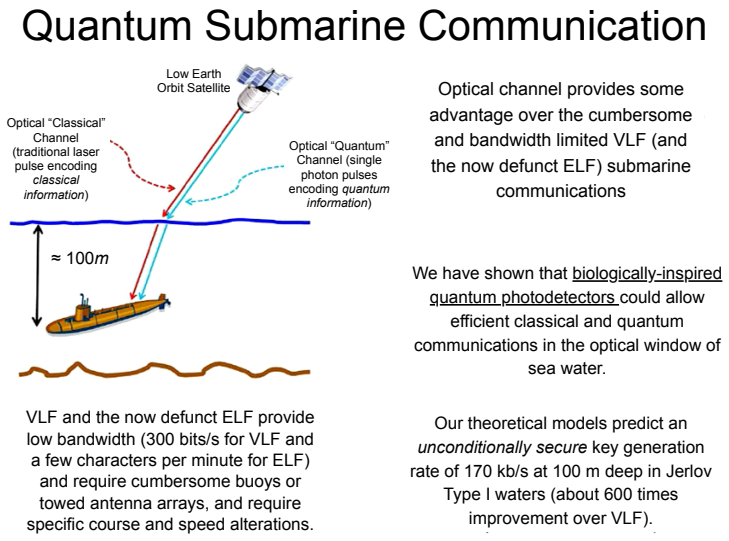

Figure 7.4. In a 2018 address to the National Academies, Dr. Marco Lanzagorta, explained how quantum communications might enable new forms of secure, satellite-to-submarine communication. Image courtesy US Naval Research Laboratory.

any two parties to communicate securely, and not even the satellite can decipher their shared key.

Scientists have also proposed BB84 protocols to improve communications with satellites directly. In one scheme, a submarine equipped with a photosensor or towing a small buoy can exchange photons with a satellite, even while submerged (see Figure 7.4). The submarine would have to make speed versus depth tradeoffs, that is, at a depth of about 60 meters, data could be exchanged at 170 kilobits per second, but this rate drops in murky waters and at deeper levels. Nonetheless, the approach is stealthy and has advantages over existing submarine communication approaches.[44]

Long-distance quantum channels for key distribution require special ingenuity to overcome a variety of technical challenges. Chinese scientists, led by that nation's "father of quantum," Jian-Wei Pan, demonstrated entanglement at 1200 kilometers by using a satellite

[44] Marco Lanzagorta, "Envisioning The Future of Quantum Sensing and Communications" (2018); Marco Lanzagorta, *Underwater Communications* (2013).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

nicknamed Micius.[45] The satellite beamed photons between distant base stations that were in the coverage area of the Micius for just five minutes.[46] Pan's team pointed to the use of the entangled photons for an Ekert-protocol secure exchange, at a distance currently impossible to achieve with terrestrial, fiber-optic connections (the quantum states degrade in the glass fiber after a distance of around 100 km without taking special measures). Yet, the approach still faces many challenges as revealed in the paper's methods. Pan's team had to beam millions of photons a second to maintain the link, and only a handful reached the base stations because of atmospheric and other interference.

Pan's achievement is part of a $100 million project in China, the Quantum Experiments at Space Scale program (QuESS). The entangled distribution over such a great distance demonstrated a substantial goal of the program. Key exchange was realized later the same year, using a mixed fiber-optic/satellite path of over 7000 km.[47] Pan's team demonstrated the key exchange by holding a videoconference between Beijing and Austria. However, this demonstration did not use end-to-end entanglement between Alice and Bob, as described by Ekert. In this initial experiment, Pan's team used the BB84 protocol, and the satellite operated as a trusted relay. Micius exchanged separate keys with each of the different ground stations.

With a relay, the implementation is not fully quantum – it's not a quantum internet – and the parties must trust the satellite's security. That's a concern. Governments will probably trust their own satellites, but this trust should not be absolute, as the computers in satellites are vulnerable to cyber attack just like computers down here on the ground.

In 2020, Pan's team announced a satellite-terrestrial quantum network covering 4600 km. The network has over 150 users, and achieved a transfer rate of 47 kilobytes a second, more than sufficient for exchanging 256-bit AES keys.[48]

[45] Launched in 2016 at the low-earth orbit of 500 km, Micius travels in a Sun-synchronous path. Micius is named for the fifth-century BCE Chinese philosopher Mozi, founder of Moism, who wrote original works on optics.

[46] Yin et al., "Satellite-Based Entanglement Distribution Over 1200 Kilometers" (2017).

[47] Liao et al., "Satellite-Relayed Intercontinental Quantum Network" (2018).

[48] Y.-A. Chen et al., "An Integrated Space-To-Ground Quantum Communication Network Over 4,600 Kilometres" (2021).

In the US, fewer than ten QKD networks have been implemented in recent years. The first, DARPA's QKD network, was implemented by Raytheon BBN, at Harvard and Boston Universities in 2003.[49] The team used dark fiber (unused fiber-optic cables) in Cambridge, Massachusetts to connect the almost 30 km long network. The network, which had trusted optical point-to-point systems and untrusted, relaying infrastructure, operated for four years. Here "untrusted" means that the relaying infrastructure could not impact the security of the data sent over the fiber.

At Los Alamos National Laboratory, scientists created a hub-and-spoke quantum network.[50] In the implementation, a central, trusted server performs the key exchange, which then enables nodes in the spokes to communicate among each other with authenticated quantum encryption. This sort of trust model works when all of the networks have a some reason to trust the central node; in the LANL demonstration, their model was a power distribution network.

Major challenges still exist for QKD implementation. The point-to-point nature required to preserve quantum states between Alice and Bob makes QKD networks more like the early telegraph than the telephone or Internet. Quantum states decohere in long fiber runs, so some networks require repeating, which, like the Micius satellite demonstration, requires trusting the repeater. Alice and Bob also need sophisticated equipment: lasers, single-photon detectors, interferometers and the like. These are now packaged in commodity QKD systems that communicate over fiber-optics, although systems that communicate in free space or using satellites are still basic science endeavors. Even so, QKD is among the most mature quantum technologies, and solving these limitations is receiving significant attention. The next section turns to such commercialization.

7.4.5 QKD Commercialized, Miniaturized

As early as 2009, three companies (ID Quantique, Switzerland; MagiQ Technologies, US; and Smartquantum, France) offered working QKD devices.[51] According to the Quantum Computing Report, at least a

[49] Elliott and Yeh, *DARPA Quantum Network Testbed* (2007).

[50] Hughes et al., "Network-Centric Quantum Communications with Application to Critical Infrastructure Protection" (2013).

[51] Scarani, Bechmann-Pasquinucci, et al., "The Security of Practical Quantum Key Distribution" (2009).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

Figure 7.5. In 2019, Air Force Research Laboratory scientists demonstrated daylight QKD using this rig at the Starfire Optical Range, located at Kirtland Air Force Base in Albuquerque, New Mexico. This is important because stray daylight entering the collector causes substantial noise that interferes with the measurement, limiting long-distance QKD during the daytime. (The Air Force's Directed Energy Directorate, which developers lasers and optics, was identified for transfer to the US Space Force in 2020.) Image by US Air Force photographer Todd Berenger.

dozen private firms are working on QKD offerings, along with a few large public companies.[52]

Despite the growing competition in QKD, adoption of QKD has been weak. For starters, without large, encryption-breaking quantum computers, there is no demonstrated need for the technology. In 2015, an unclassified summary of the US Air Force advisory board report threw cold water on QKD, apparently finding that QKD significantly increases system complexity while providing "little advantage over the best classical alternatives."[53] The USAF's full report is not publicly available, but perhaps the board meant that as system com-

[52] ArQit, InfiniQuant, KETS Quantum Security, Phase Space Computing, QEYnet, Qrate Quantum Communications, Quantropi, Quantum Xchange, Qubit Reset LLC, Quintessence Labs, QuNu Labs, SeQureNet, and VeriQloud; larger firms include Nippon Telegraph and Telephone Corporation (NTT), Raytheon BBN Technologies, and Toshiba.

[53] US Air Force Scientific Advisory Board, *Utility of Quantum Systems for The Air Force Study Abstract* (2016).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

4. New York - Moscow 1340 [753], 21 September [20 September] 1944:

--149F-- detained VOLOK (?who is?) working at the ENORMOZ plant. He is a fellow countryman [U.S. Communist]. --1U-- (?recognition?) (?of? ?from?) his work they dismissed (?him?). The cause of the dismissal was his active work in the past in progressive organizations.

According to --1U-- of the fellow countrymen [U.S. Communists], LIBERAL (?is in touch with CHESTER he --2F-- cutter **ERCESE? [this part very dubious]) once a month. CHESTER is interested in whether we are satisfied with the cooperation and whether there are not any misunderstandings. About concrete details of the work he does not inquire. Inasmuch as CHESTER knows about the role of LIBERAL's group we beg consent to inquire of CH. through LIBERAL about (?sketches (drafts)?) from (?the milieu?) of persons working on ENORMOZ and other spheres of technical science.

Here the subject changes; in the new section, there is some mention of a person named LARIN, but the text is unintelligible. The signature is MAY.

5. New York - Moscow 1699 [conclusion of 940], 2 December 1944 (the preceding part or parts of this message cannot be located):

Conclusion of telegram no. 940

Stated to be (?participants?) --1G-- (?research?) on the problem are HANS BETHE, NIELS BOHR, ENRICO FERMI, JOHN NEUMANN, BRUNO ROSSI, GEORGE KISTIAKOWSKI, EMILIO SEGRE, G.I. TAYLOR, WILLIAM PENNEY, ARTHUR COMPTON, ERNEST LAWRENCE, HAROLD UREY, HANS (?STAN? ?STROGN?) AR(?K? ?L? ?M?), EDWARD TELLER, PERCY BRIDGEMAN, WERNER EISENBERG[a], --1F-- AS --4F-- [There follows a repetition of all these names.] --5F-- (?of?) our country turned [or "applied"] to NAPOLI the latter (?did not?) --2F-- him [or "his"] --2F-- BEK [Beck?] --7F--. When he tried to see RULEV, he was not admitted to see him by the latter's secretary.

(?ANTON?)

a. Mistake for WERNER HEISENBERG? It has been known for some time that Heisenberg was working for the German Reich throughout the war.

Figure 7.6. Richard Hallock, an analyst at the US Army's Signal Intelligence Service, discovered that Soviet spies were reusing portions of one-time pads. The revelation allowed the Service, a forerunner of the National Security Agency, to decrypt them. This summary of intercepted communications shows that the Soviets had identified the main scientists involved in the Manhattan Project (Soviet cryptonym "ENORMOZ"; "LIBERAL" is Julius Rosenberg). The American analysts also ponder whether the Russians thought that Werner Heisenberg was working on the American fission project; alas he was working for the Germans. The decryption project, code name VENONA, ran from 1943 through 1980. (National Security Agency and Central Security Service, "VENONA" (2021))

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

plexity increases, so do attack surfaces. A more complex system gives attackers more opportunities to interfere with communications, and perhaps the side channel attacks possible on quantum devices will be more difficult for network operators to understand. Aside from device problems, there remains the old problem that users can be fooled into granting access. Perhaps the USAF report's skepticism reflects that the US government has a decades-old system of using trusted human couriers to transport high-value key material.

In October 2020, the NSA released a statement clarifying that it would not use QKD to secure the classified and sensitive-level networks it is responsible for protecting, and this NSA statement articulated the likely reasons why QKD has not been more commercially successful. Calling out the hype, the NSA statement recognized that QKD advocates "occasionally state bold claims based on theory" but that in reality, the technology is "highly implementation-dependent rather than assured by laws of physics." The NSA's specific objections related to the need to install new, more complex and expensive infrastructure that itself may have vulnerabilities.[54] Indeed, Russian scientist Vadim Marakov has elucidated a series of attacks on QKD *systems* (but not the underlying BB84 protocol).[55] The NSA concluded that whatever confidentiality QKD offers "can be provided by quantum-resistant cryptography, which is typically less expensive with a better understood risk profile."[56] As with the NSA, many companies probably see little reason to adopt a technology that will require infrastructure changes, require more training, introduce new complexities, and all for limited benefits against attackers many years in the future.

Nevertheless, QKD vendors are trying to overcome the skepticism. Four recent developments paint a path for greater QKD adoption in both the private sector and in governments. First, QKD devices have been miniaturized. ID Quantique and MagiQ both market rack-mounted QKD systems. Second, the general upset caused by the Snowden documents caused policymakers in other regions to make stronger communications security a priority and to make large

[54]Scarani and Kurtsiefer, "The Black Paper of Quantum Cryptography: Real Implementation Problems" (2014).

[55]Anqi et al., "Implementation Vulnerabilities in General Quantum Cryptography" (2018).

[56]National Security Agency, "Quantum Key Distribution (QKD) and Quantum Cryptography (QC)" (2020).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

vertical industrial policy investments in quantum technologies. This policy commitment may overcome the natural resistance to a switch to QKD. For instance, the European Union's quantum technologies strategy makes wide dispersal of QKD (and QRNG) a priority, even for consumer devices. The European Union's OpenQKD project, a three-year €15 million program (2019–2022), explicitly seeks standardization and other objectives to kick start a Continental QKD industry. Third, progress is being made on technical challenges, such as increasing the length of fiber over which QKD can operate: in 2018 scientists demonstrated QKD over a 400 km fiber run.[57] These ultra-long runs cause signal attenuation, and key acquisition slows to a crawl (as much as 24 hours for a key block), but improvements are steady. Finally, concerns about the privacy and security of 5G telecommunications networks is driving international concern and an unprecedented search for technical security measures.

On this last point, the security of 5G, consider the activity of South Korea Telecom (SK Telecom). Operating in the shadow of North Korea, with its active, audacious intelligence activities, SK Telecom officials must contemplate that their own employees might be forced into revealing telecommunications data to North Korea. In 2016, SK Telecom started implementing QKD in some back-haul operations of their LTE network. This effort expanded in later years to 5G infrastructure. As QKD is implemented in SK Telecom's stack, the number of employees who could be coerced into revealing information to North Korea presumably winnows.

QKD or quantum networking to a consumer handset will probably never be a reality, but QRNG may be on the threshold of widespread adoption: In May 2020, ID Quantique announced that its system-on-a-chip QRNG had been implemented in a handset offered by SK Telecom. In September 2020, as part of South Korea's $133 billion "digital new deal" program, the country will pilot QKD implementations in several critical infrastructures.

7.5 Quantum Internet

What's colloquially called "quantum internet" could be thought of as the attempt to bring quantum computing to an infrastructure reminiscent of the Internet. With a quantum internet, any two parties on a large network could communicate over some kind of quantum

[57] Boaron et al., "Secure Quantum Key Distribution Over 421 Km of Optical Fiber" (2018).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

circuit made up of flying qubits, just as the conventional Internet allows two parties to communicate using a virtual circuit built using packet switching. With a quantum network, Alice and Bob could communicate using quantum states, allowing them to enjoy the protection of quantum cryptography, and also giving them the ability to engage in quantum protocols or compute with quantum algorithms.

There are three non-obvious advances that follow from the resilient management of quantum states across distance and devices: first, mastery of quantum networking would make it possible to assemble a quantum computing cluster. Thus quantum networking could change the strategy by which organizations plan to build large quantum computers. Instead of mastering the management of a single device with many qubits, a quantum network would allow organizations to connect together several smaller, perhaps less expensive and easier-to-manage devices into a cluster that has more qubits and volume than any competitor. Such a quantum network might reside within a single building. But while companies such as IBM, with its research lab full of quantum devices, seems well poised to do this, there is (as of yet) no public evidence that IBM or others are taking this tack.

Second, a quantum network could enable *blind* quantum computing. Recall that quantum computing, because of its expense and complexity, is likely to be available as a cloud service rather than as on-premises devices. Currently, users of cloud-based quantum computers offered by Amazon and its competitors access those devices through classical communication-and-control computers. In a world with a functioning quantum internet, that cloud access could become end-to-end quantum intermediated. At that point, the owner of the cloud-based quantum computer would be blind to the user's action. Being blinded would limit policy options because the quantum computing owner might not be able to detect and deter unwanted uses of the device, such as cryptanalysis or currently unimagined noisome behavior.

Depending on how it is implemented, a quantum internet might deny adversaries the ability to spy on metadata. Currently metadata, the data about data in the communications network, such as who calls whom and when, is a key tool of intelligence agencies. Metadata is well structured and relatively easy to analyze. Most people can be identified by their metadata (because most people do not constantly obtain new, clean communications devices) and even though

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

metadata lacks information about the content of communications, metadata often hints at individuals' activities. If a quantum internet is used to set up quantum circuits between the endpoints so that the flying qubits properly travel from Alice to Bob, then such a setup might be susceptible to surveillance. But if the quantum internet is itself controlled *inband* with its own quantum signaling, then it will be difficult to track who is talking to whom. Although this would be a real "going dark" problem that might have intelligence agencies and advertising agencies alike worried, such a possible network seems decades in the future.

Indeed, the challenge of realizing a large-scale quantum network is related to the very attributes that give quantum communications so much privacy: the no-cloning property. Jian-Wei Pan's team demonstrated quantum communication over short distances, extending networks on optical fiber over a distance of about 100 kilometers in 2008.[58] In traditional fiber-optic networks, light becomes diffused from the twists and turns of the fiber and needs to be periodically "repeated," or boosted, to travel to its final destination.[59] But the act of repeating requires copying, which is something that quantum networks can't do. Thus, a repeater on a quantum network breaks the end-to-end guarantees that users of a quantum network would want the network to provide. Although an approach may be developed to address this problem, in the near term, quantum networks will likely involve some sort of trusted repeater that catches the flying qubit, performs a classical computation, and then transmits a brand-new flying qubit down the fiber.

Repeater node trust could be seen as a blessing or a curse: depending on one's perspective, it either can enable lawful access to otherwise unbreakable key exchange, or it represents a problematic security loophole. Still, even a classically relayed quantum network is advantageous, in that if one controls the relay points, one could detect interception and still enjoy lawful access when needed.[60] For instance, the political attributes of China probably fit neatly with

[58] Yuan et al., "Experimental Demonstration of a BDCZ Quantum Repeater Node" (2008).

[59] Briegel et al., "Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication" (1998).

[60] Farrell and Newman, "Weaponized Interdependence: How Global Economic Networks Shape State Coercion" (2019). Consider the rise of "Weaponized Interdependence," state use of networked infrastructures to leverage panoptic capabilities and use chokepoints for control.

the limits of classical repeaters. Those nodes could be operated by state-controlled companies, and surveilled when desired by domestic law enforcement and intelligence, while denying that same ability to foreign adversaries. Jian-Wei Pan himself boasted, "China is completely capable of making full use of quantum communications in a regional war … The direction of development in the future calls for using relay satellites to realize quantum communications and control that covers the entire army."

A *quantum repeater* or *quantum memory router* can overcome the trust problem. The first re-transmits the flying qubit, and the second allows the flying qubit to fly off in one of several possible directions. Such devices are still in their infancy.[61] Quantum internet routers are in effect small quantum computers. One approach uses atomic vapor technologies, specifically Electromagnetically Induced Transparency (EIT), introduced in Section 2.2, "Atomic vapor technologies" (p. 41). Scientists are working on the fidelity of copying and storage time; as of 2019, EIT memory loses fidelity in just microseconds.[62]

Quantum "teleportation" is a mechanism being explored to build quantum networks. Teleportation in science fiction is as unexplained as it is exciting. What exactly do teleporters do? How they work seems to change from season to season and among different series. The most well-developed fictional teleportation system appears in *Star Trek*, but the fictional "transporter" was originally created by the series writers to save the cost (in terms of special effects and screen time) of needing to use the ship's shuttle craft to send the crew down to the planet.[63] Over time, the transporter became a useful plot device for creating and then exploring psychological situations, but similar to the show's "warp drive," the underlying physics were never satisfactorily explained.[64]

[61] Yan and Fan, "Single-Photon Quantum Router with Multiple Output Ports" (2014); Pant et al., "Routing Entanglement in The Quantum Internet" (2019); Korzeczek and Braun, "Quantum-Router: Storing and Redirecting Light at The Photon Level" (2020).

[62] Yunfei Wang et al., "Efficient Quantum Memory for Single-Photon Polarization Qubits" (2019).

[63] Whitfield and Roddenberry, *The Making of Star Trek* (1968).

[64] In both the original and Next Generation Star Trek series, transporters caused accidents and created doppelgangers: a good and evil Captain Kirk, and a copy of Commander Riker. In Star Trek Voyager, a teleporter accident fused a Vulcan (Tuvok) with a Talaxian (Neelix), creating the unfortunate Tuvix. In Spaceballs (1987), President Skroob's head materialized backwards, so that he faced his pos-

In contrast to mythical teleportation devices, *quantum teleportation* is an effect that is well understood and has even been demonstrated. Quantum teleportation moves the *quantum state* from one particle to a second, irrevocably changing the state of the first particle in the process. Because the state is moved and not copied, quantum teleportation violates neither the Heisenberg uncertainty principle nor the "No Cloning" theorem, which holds that quantum states cannot be precisely copied.

One possible way to construct a quantum router is to use quantum teleportation to transmit data to some point in the distance, in effect creating a point-to-point communication between Alice and Bob. Teams at TU-Delft led by Stephanie Wehner and Ronald Hanson have impressive accomplishments in advancing entanglement and in teleportation. In a TU-Delft demonstration of quantum teleportation, Alice and Bob share a classical communication channel and an entangled particle. The entangled particle is a nitrogen-14 spin inside a diamond. Known as a "nitrogen-vacancy" chamber, this imperfection in a synthetic diamond isolates and insulates the nitrogen atom from the outside environment (see Chapter 2, Section 2.2, "Nitrogen vacancy" (p. 41)). That isolation makes the nitrogen spin more resilient to unwanted interference. With the nitrogen atoms entangled over a distance, Alice takes a second atom, the information bit, and performs a so-called "Bell measurement" between her entangled atom and the second atom. The measurement causes a corresponding change to Bob's entangled qubit. Bob can then extract the information – the state that Alice sent – by communicating with Alice over a classical channel. Alice tells Bob the transformations she made; by performing these same steps, Bob can extract the value of the original state.[65] Because this process uses both quantum entanglement and classical channels as a medium, teleportation protocols do not support faster-than-light communication, as is sometimes claimed.[66]

terior, to the delight of the crew. An earlier transporter appeared in the movie "The Fly" (1958), in which a teleporter affixed a fly's head atop a smart scientist's body. The scientist kept his mind, but was under siege from the fly's entomic instincts. See Rzetenly, "Is Beaming Down in Star Trek a Death Sentence?" (2017) for contemporary examination regarding the philosophical implications of creating a perfect copy of a person while destroying the original.

[65] Pfaff et al., "Unconditional Quantum Teleportation between Distant Solid-State Quantum Bits" (2014).

[66] J. G. Ren et al., "Ground-To-Satellite Quantum Teleportation" (2017).

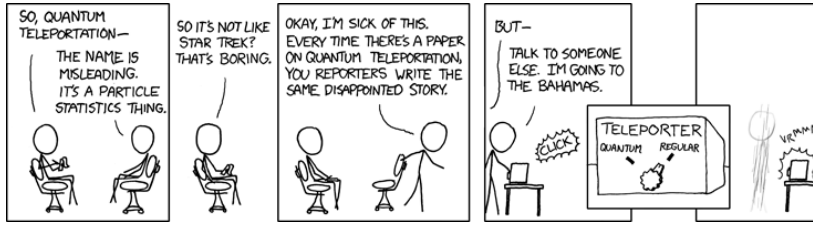

Figure 7.7. xkcd #465: Quantum Teleportation. Used with permission. xkcd.com/465/

(See the sidebar "Alas, Faster-than-light Communication Is Not Possible" on page 301.)

Quantum teleportation was first conceived by an international team that included Charles Bennett and Gilles Brassard.[67] In 1997, scientists at the Austrian Institut für Experimentalphysik demonstrated teleportation in a laboratory setting using photons and their spins. Jian-Wei Pan was part of that team, then training under Austrian physicist Anton Zeilinger. Since then, teleportation has been demonstrated at greater distances. The TU-Delft team demonstrated teleportation at 3 meters in 2014 and by 2017, Jian-Wei Pan's team demonstrated teleportation at 1400 km using entangled photons between a base station in Ngari, Tibet (elevation 4500 m) and the Micius satellite.

To enable teleportation over greater distances, and indeed in a quantum internet, scientists are experimenting with entanglement *swapping*. In entanglement swapping, communication between Alice and Bob is made possible even if they lack a point-to-point path. The process works with a device, operated by a third party (here called Faythe), close enough to Alice and Bob to receive an entangled photon separately from each of them.[68]

The European Union has identified a quantum internet as a central goal in its €1 billion investment in quantum technologies,[69] and scientists there have already achieved several key steps towards the creation of a quantum internet. The most synoptic expression of this vision, written by the German physicist Stephanie Wehner, makes it

[67] Charles H. Bennett et al., "Teleporting an Unknown Quantum State via Dual Classical and Einstein-Podolsky-Rosen Channels" (1993).
[68] Halder et al., "Entangling Independent Photons by Time Measurement" (2007).
[69] European Commission, High Level Steering Committee, DG Connect, "Quantum Technologies Flagship Intermediate Report" (2017b).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

clear that a quantum internet is seen as a special purpose network to exist alongside the conventional Internet.[70] The quantum internet is intended to maintain a channel capable of special functions, such as quantum key distribution, secure identification, and others.

If nations decided to invest in creating a quantum internet, network paths would become a key focus. From a technical perspective, all paths would have to be fully quantum mechanical, or the quantum state would collapse and the technology would fail. Strategically, adversaries along those paths could easily interfere with the quantum state, causing it to collapse. These attacks on availability need not be at the router or even that sophisticated. Anything that degrades the light will work, meaning that these attacks might be easily deniable, and attributable to accident and so on.

Going back to the time of the telegraph, communications find their way along wires on specified routes. If a telegraph pole fell in a storm, that path would be interrupted, and the pole would have to be replaced or a new path set into place. One major advance of the Internet was packet switching, the conversion of communications into datagrams that could take multiple routes. The sender and recipient need not specify these routes. But this lack of specificity comes with a downside: because the communications' paths change dynamically, an attack can intentionally interfere with one route and force the communications to travel over another route with lower legal or technical protections.[71] Recently, the risk that internet communications take unnecessarily circuitous routes through other legal jurisdictions has become a concern of some nations. A 2019 study focusing on path-based risks evaluated tens of thousands of likely paths a user's browser might take when visiting popular sites. The group found that 33 percent "unnecessarily expose network traffic to at least one nation state, often more."[72] Some nations are building local internet exchange points to keep more communications domestic, and out of paths that traverse China, Russia, the US, or its "five-eyes" allies.

A quantum internet would almost certainly require that nations and sophisticated companies create dedicated fiber links for a quan-

[70] Wehner, Elkouss, and Hanson, "Quantum Internet: A Vision for The Road Ahead" (2018).

[71] Woo, Swire, and Desai, "The Important, Justifiable, and Constrained Role of Nationality in Foreign Intelligence Surveillance" (2019).

[72] Holland, J. M. Smith, and Schuchard, "Measuring Irregular Geographic Exposure on The Internet" (2019).

tum network, making it more like a separate, dedicated private network. The infrastructure for communication is likely to become much more state-specific. Already, sophisticated users are able to choose the paths that their conventional internet communications travel; the same will likely be true of quantum networks, if they are ever created. Already the Dutch telecom provider KPN has built a fiber-optic, quantum channel network backbone between Leiden, Delft, Amsterdam, and The Hague. (The KPN network does not require repeating, because of the short distances among these cities.[73])

Another option comes from satellites. It seems less likely that a satellite could be manipulated by an adversary than an underwater repeater. At least a half a dozen countries are pursuing satellite-based QKD programs.[74] Either physical or cyber manipulations could be impactful. Thus, initiatives such as Elon Musk's SpaceX/Starlink satellite network, which intends to populate the sky with internet-providing satellites, could also form the backbone of a tamper-resistant network that is mostly classical but could include quantum elements: perhaps two quantum-enabled ground-stations on opposite sides of the planet would communicate with a message passed from satellite to satellite.

Similarly, one might imagine businesses that place point-to-point servers connected by quantum channels in physically inaccessible places, for instance submerged in containers that if opened would fail.

7.6 Conclusion

Quantum communications can be binned into two categories: first, the related applications of quantum random number generation and key distribution, and second, technologies that enable a quantum network or quantum internet. While quantum random number generation and key distribution are both maturing technologies, early systems have been commercialized and are in use today. These technologies meet two central requirements for secure communications technologies: they are information-theoretically secure and enable distribution of keys at a distance. Those who adopt QKD will never have to worry that the keys they use today in encryption systems based on the RSA or Elliptic Curve public key cryptography systems might be cracked by some powerful quantum computer in the future

[73]Baloo, "KPN's Quantum Journey, Cyberweek 2019, Tel Aviv, Israel" (2019).
[74]Khan et al., "Satellite-Based QKD" (2018).

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

Alas, Faster-Than-Light Communication Is Not Possible

Experiments in entanglement show that entangled particles somehow "know" the quantum state of their twin. One might think of entangled particles as parts of a connected system. Scientists do not know how they are connected, but scientists can show through Bell tests (see Section B.4 (p. 513)) that they are.

Quantum teleportation takes advantage of the linkage between distant particles to teleport a state from Alice's entangled particle to Bob's. Because Bob's particle reacts instantly, even when separated by great distances, some have speculated that teleportation could somehow enable faster-than-light (superluminal) communication. Alas, quantum teleportation does not enable faster-than-light communication.

Superluminal communication is impossible because quantum teleportation protocols depend on classical channels to extract the meaning from the entangled qubits. After teleporting a state to Bob, Alice and Bob communicate over a classical channel. Bob determines the teleported state by applying transformations that correspond to Alice's instructions.[a] This is the basis of the BB84 and E91 protocols.

So as one can see, the reversion to a classical channel, and the complexity of the information exchange and discovery, makes it impossible to communicate faster than light speed.

[a] Pfaff et al., "Unconditional Quantum Teleportation between Distant Solid-State Quantum Bits" (2014).

– although adopters of today's QKD systems still need to verify that the QKD systems themselves are still secure against traditional vulnerabilities, such as electromagnetic radiation or cyberattack.

Yet, if experience with other privacy-enhancing technologies holds, only entities with the most to lose will affirmatively adopt them. Banks, militaries, intelligence agencies, and other entities with the awareness and budget are likely adopters. But for everyone else, three other requirements must be met: the system has to be fast, it has to be usable by anyone, and it has to be on by default. The coming availability of classical encryption that is quantum resistant will be satisfactory for many actors. Unless some economic interest arises and militates strongly in favor of quantum encryption, most consumers and businesses will rely on classical alternatives.

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press

The quantum internet's best use in the future – aside from its ability to procure funding for prestigious science projects – seems to be the interconnection of existing, small quantum computers into a cluster of unprecedented power. The other benefits, relating to time synchronization and astronomy, seem so tethered to scientific and technical users that it is difficult to see how they would inspire a commitment to outlay the money to make a quantum internet happen. In the nearer term, the quantum internet's potential to make communications end-to-end secure and eliminate metadata surveillance may be the driving factor for nation states to invest in the technology.

https://doi.org/10.1017/9781108883719.011 Published online by Cambridge University Press