


# Appendix

## Surviving the Quantum Cryptocalypse

Einstein famously described quantum mechanics as “spooky action at a distance.” Yet many quantum mysteries seem slightly less mysterious when understood through the perspective of computer science. The mathematical abstraction of quantum computing can be thought of as a generalization of probability theory and Shannon’s information theory.<sup>1</sup> Whereas classical probabilities are always positive, so-called probability amplitudes are complex numbers that can be positive, negative, or imaginary. The interference and cancellation of amplitudes is the basis for most of what gets described as quantum weirdness.<sup>2</sup> A brief review of some of the mathematics of quantum computing is offered here to help demystify some concepts.

A classical bit is a binary scalar  $b \in \{0,1\}$ . A qubit, by contrast, is a vector of complex-valued amplitudes representing the likelihood that the qubit is zero or one. By definition, a *qubit* is  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ , so it can be described as a *superposition* of  $|0\rangle$  and  $|1\rangle$ .<sup>3</sup> Just as a bit can be represented in countless different physical systems—transistors, vacuum tubes, semaphores, neurons—according to engineering convention, logical qubits can be physically realized in different ways. Functional qubits have been successfully represented as ions trapped by tuned lasers, quantum dots in silicon, loops of superconducting metals, and atomic vacancies in a diamond lattice.<sup>4</sup> Yet it is important to appreciate that a logical qubit is a mathematical abstraction. The quantum formalism has even been used to model some macroscale phenomena in social systems.<sup>5</sup>

A qubit can be visualized on the so-called Bloch sphere (fig. A.1) as a vector  $|\psi\rangle$  from the origin to any point along the surface of the sphere, and only along the surface.<sup>6</sup> The poles on the Z-axis represent the “computational basis” states of  $|0\rangle$  and  $|1\rangle$ , while the equator represents an infinite number of balanced superpositions of  $|0\rangle$  and  $|1\rangle$ . The X-axis states on the equator represent the “Hadamard basis” of balanced superposition states, known as  $|+\rangle$  and  $|-\rangle$ . Quantum computing on a single bit can be visualized as a rotation of this vector along the surface of the sphere. For example, the transformation of a qubit initialized to  $|0\rangle$  into a balanced superposition where  $|0\rangle$  and  $|1\rangle$  are equally likely can be visualized as a rotation of the qubit  $|\psi\rangle$  from the north pole  $|0\rangle$  to the point  $|+\rangle$  on the equator.<sup>7</sup>



**Figure A.1. Bloch sphere representation of a qubit as a complex-valued vector of probability amplitudes**

Things start to get interesting with multiple qubits. Unfortunately, the Bloch sphere does not work very well for visualizing more than one qubit. Instead, two qubits can be represented with a vector of four amplitudes and four computational bases for each of the possible combinations of zeros and ones:  $|\psi\rangle = \alpha|00\rangle + \beta|01\rangle + \gamma|10\rangle + \delta|11\rangle$ . The quantum state of a multi-qubit system is described as *entangled* if it cannot be expressed as the product of single qubit vectors.<sup>8</sup>

In general, the quantum state of a system of  $n$  qubits has  $2^n$  amplitudes. As a result, the simulation of quantum computers on classical machines quickly becomes intractable. Microsoft has developed a quantum programming language called Q# and a digital emulator that anyone can download.<sup>9</sup> Q# requires 16 gigabytes of memory to simulate 30 logical qubits, while 40 qubits require 16 terabytes and the services of the Azure cloud services. One hundred qubits would need a mind-boggling 19 million yottabytes and many times the age of the universe to do any useful computation. Computer scientists conjecture that there may also exist quantum algorithms that cannot be simulated on a classical machine regardless of resource constraints.<sup>10</sup>

In classical computing, strings of bits can be copied and manipulated with Boolean operators such as AND, OR, and NOT. Boolean operators can be combined into more complicated instructions, which in turn can be arranged into algorithms, which are logical recipes to compute the answers to problems given some input. The analog to a Boolean operator in quantum computing is a normalized linear transformation of the quantum vector known as a *quantum gate*. A quantum gate is a (normalized) linear transformation of the quantum vector. This gate is a logical concept distinct from any actual circuits in a real machine.<sup>11</sup>

To use a quantum algorithm to solve a practical problem, it must be possible to encode input data on a set of logical qubits, evolve the quantum state through a sequence of linear transformations (quantum gates) that manipulate superposition and entanglement, and measure the result. Measurement maps the overall quantum state onto just one of the system's  $2^n$  computational bases, which means that important quantum information may be lost if the problem is not set up the right way. Not every mathematical problem is so conveniently structured.

The way in which Shor's famous algorithm works is instructive.<sup>12</sup> Shor demonstrated that the problem of factoring an  $n$ -bit number, which can be reduced to the problem of finding a repeating period in a sequence of  $2^n$  numbers, can be solved with the Fourier transform on  $n$  qubits in such a way that the "wrong" answers cancel out while the "right" answer can be measured. Computer scientists use "big O" notation to describe the relative amount of computational resources that must be used to find an answer to a problem given an input of size  $n$  (generally assumed to be large given the data sets and bandwidth with which modern computers usually work). A polynomial time algorithm might run in  $O(n^2)$ , which means the resources required will scale with the square of the input. An exponential algorithm, however, might run in  $O(2^n)$ , and clearly  $2^n \gg n^2$  for large  $n$ . Classical factoring algorithms run in exponential time, which is why cracking a 2048-bit RSA key would take more time than the age of the universe. Yet Shor's algorithm runs in polynomial time,  $O(n^2)$ , meaning that it provides an exponential speedup over the fastest classical alternative (known as the general number sieve).<sup>13</sup>

Any usable quantum computer must be able to tolerate faults in the preparation, computation, and measurement of qubits. Errors occur in classical machines, too, but they are easier to correct. It is possible to check and repair classical errors using backup data and checksum bits with relatively low overhead.<sup>14</sup> In quantum computing, however, arbitrary data cannot simply be copied for inspection. The no-cloning theorem in quan-

tum mechanics states that it is not possible to perfectly copy an arbitrary qubit to another qubit without changing it.<sup>15</sup> This is a major contrast with classical computing.

To get around this problem, quantum error correction relies on some clever algorithmic tricks. For instance, it is possible to control for errors in a single qubit by using five additional (ancilla) qubits initialized to known values. Higher error rates in the physical implementation of logical qubits, moreover, require more overhead for error correction. A very large number of physical qubits may be needed to support small numbers of fully functional logical qubits.<sup>16</sup> Cracking a 2048-bit RSA key with Shor's algorithm requires just over 4,000 logical qubits.<sup>17</sup> Yet when we factor in error correction, the fastest known implementation would require 20 million qubits!<sup>18</sup> This recent innovation is a substantial improvement over the previous estimate of a billion qubits, but even 20 million is orders of magnitude more qubits than prototype machines are able to maintain in coherence today (e.g., Google's Sycamore entangled only 53 qubits).

Thermodynamic noise and electromagnetic effects can cause the decoherence of the quantum state, which disrupts superposition and entanglement and results in a loss of quantum information. Errors can be managed but not eliminated by operating in very cold, carefully shielded environments. One implication is that the ability to build and maintain advanced laboratories and precision machinery is a vital infrastructural enabler of innovation in quantum computing. Given the difficult engineering challenges associated with maintaining coherence at scale, moreover, not all qubits in machines that are described as "quantum computers" are fully functional qubits. More simply, "not all qubits are created equal."<sup>19</sup> The Canadian firm D-Wave boasts that its machines have thousands of qubits, but they are not fully functional. D-Wave implements a different approach to quantum computing known as quantum annealing that requires significant computational overhead to implement quantum gate arrays. For cryptology, therefore, "quantum simulators and quantum annealers are single-purpose devices, unable to run, for example, Shor's algorithm, and do not have any known applications to cryptanalysis."<sup>20</sup> Perhaps more importantly, the D-Wave implementation of qubits is so noisy that they can only support very limited forms of quantum functionality, if any.<sup>21</sup>

Quantum networking is a different type of technology from quantum computing, but it draws on similar quantum principles. Quantum mechanics can be used to provide an innovative solution to the cryptographic key distribution problem—a major headache for many secret communicators throughout history. The quantum computing version of the famous

Heisenberg Uncertainty Principle, which states that you cannot localize the position and momentum of a particle at the same time, can be leveraged for cryptography. The measurement of the value of a single qubit can be considered relative to an axis of the Bloch sphere. If you try to measure a qubit that is in a balanced superposition along the computational basis (i.e., the quantum state vector  $|\psi\rangle$  points to the equator and you try to determine the pole to which  $|\psi\rangle$  is closest), the result will necessarily be random:  $|0\rangle$  or  $|1\rangle$  is equally likely. Conversely, if a qubit in the state of  $|0\rangle$  or  $|1\rangle$  is measured along the Hadamard basis, the result will also be random:  $|+\rangle$  or  $|-\rangle$  is equally likely. A photon (or pulse of photons) can be used to represent a logical qubit. Polarization of photons with a rectilinear or diagonal filter can represent the computational and Hadamard basis states.<sup>22</sup> Measuring a horizontally polarized photon with a diagonal detector provides a random result.

This fact can be leveraged to detect an eavesdropper (known as Eve by explanatory convention) in a communication circuit (between Alice and Bob). Table A.1 provides a simplified illustration of the BB84 protocol for quantum key distribution (QKD). Alice generates a random string of bits and a random series of photon polarizations. Bob prepares his own random series of measurements in the diagonal or rectilinear basis. If Bob measures a qubit with the same polarization that Alice uses to prepare it, then he will get an accurate result, but if not he will get a random result. Bob then transmits his measurement scheme to Alice via a classical channel such as a telephone or internet connection, and Alice tells Bob which of his polarization choices match hers. By ignoring the mismatches and measuring the stream of photons from Alice in the correct bases, Alice and Bob can now share the same string of random bits. That is, they can construct a shared symmetric key that has not been revealed through the classical channel. Even though Eve can learn the measurement scheme by tapping the classical channel, she cannot deduce the private bit string unless she actually performs a measurement on the quantum channel. Covert measurement means that Eve will in effect be copying photons, but perfect copying of arbitrary quantum data is prohibited by the no-cloning theorem. If Eve tries to measure the qubits, she will change their value in the process. Bob will then end up measuring a different bit string than Alice transmits. When they then try to exchange symmetrically enciphered messages, Bob or Alice will decrypt gibberish; it will be obvious that there is a problem. Alice and Bob can also openly compare some random fraction of their bit string on the classical channel to check for errors, using the rest of the private portion if they learn that the error rate in the public portion is ac-

ceptably low. If they detect a high error rate, however, then Alice and Bob can restart again or cease communicating altogether.

**Table A.1. Example of quantum key distribution**

Action	Data															
Alice generates a secret random bit string.	1	0	0	1	1	0	1	0	0	1	1	1	0	0		
Alice generates a random polarization sequence.	+	x	+	x	+	+	+	x	x	x	+	+	x	x		
Alice transmits polarized photons.	-	/		\	-		-	/	/	\	-	-	/	/		
Bob generates a random measurement sequence and informs Alice via a classical communication channel.	x	x	x	+	+	+	x	+	x	+	+	+	x	+		
Alice informs Bob which polarizations match via the classical channel →Eve would not be detected here, but she cannot learn too much either.	N	Y	N	N	Y	Y	N	N	Y	N	Y	Y	Y	N		
Bob ignores the random bits created by his mismatches. →Alice and Bob now share the secret bit string 0100110.	?	0	?	?	1	0	?	?	0	?	1	1	0	?		
If Eve copies Alice's photons, Bob will measure random bits. →Bob's increased error rate reveals Eve's intervention.	?	?	?	?	?	?	?	?	?	?	?	?	?	?		

Quantum information science is an exciting growth area at the nexus of experimental physics and theoretical computer science. There are many other computational and cryptologic applications besides those discussed here, some of which could have important military or intelligence applications.<sup>23</sup> This ferment of research and discovery is simultaneously expanding the potential for *both* offensive and defensive advantage in global security affairs.

**Notes**

1. Mark M. Wilde, *Quantum Information Theory*, 2nd ed. (New York: Cambridge University Press, 2017).

2. Scott Aaronson, *Quantum Computing since Democritus* (Cambridge: Cambridge University Press, 2013), 132–49, doi:10.1017/CBO9780511979309.

3. Quantum states are represented as vectors in Dirac notation:  $|\psi\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix}$  with basis vectors  $|0\rangle \equiv \begin{bmatrix} 1 \\ 0 \end{bmatrix}$  and  $|1\rangle \equiv \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ .

4. Gabriel Popkin, “Scientists Are Close to Building a Quantum Computer That Can Beat a Conventional One,” *Science*, 1 December 2016, <https://www.sciencemag.org/>.

5. Jerome R. Busemeyer and Peter D. Bruza, *Quantum Models of Cognition and Decision* (Cambridge: Cambridge University Press, 2012), <https://pdfs.semanticscholar.org/>; and Emmanuel Haven and Andrei Khrennikov, *Quantum Social Science* (New York: Cambridge University Press, 2013).

6. This constraint arises because a qubit is defined in terms of probability amplitudes, and the sum of the probabilities of different outcomes of an event must sum to unity, that is,  $|\alpha|^2 + |\beta|^2 = 1$ .

7. That is, the Hadamard gate  $H \equiv \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$  can be used to rotate the vector  $|0\rangle$  into a balanced superposition of  $|+\rangle$  and  $|-\rangle$  through the linear transformation  $H|0\rangle = |+\rangle$ .

8. For example,  $|\psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$  is an entangled state (in which it is equally likely that both qubits are zero or both are one but completely unlikely that they have different values) because there do not exist separable qubits  $|\phi\rangle = \alpha|0\rangle + \beta|1\rangle$  and  $|\phi\rangle = \gamma|0\rangle + \delta|1\rangle$  such that  $|\phi\rangle \otimes |\phi\rangle = |\psi\rangle$ .

9. Microsoft, Microsoft Quantum Development Kit, accessed April 2020, <https://www.microsoft.com/>.

10. Ran Raz and Avishay Tal, “Oracle Separation of BQP and PH,” Electronic Colloquium on Computational Complexity, Report no. 107, 31 May 2018, <https://eccc.weizmann.ac.il/>.

11. For example, the controlled not gate  $CNOT \equiv \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$  flips the value of the second of two qubits contingent on the value of the first. A simple algorithm for placing two initialized qubits into an entangled superposition is given by the equation  $CNOT(H \otimes I)|00\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$  where the identity matrix  $I \equiv \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ .

12. For an accessible summary see Scott Aaronson, “Shor, I’ll Do It,” Shtetl-Optimized (blog), 24 February 2007, <https://www.scottaaronson.com/>. See also Emily Grumbling and Mark Horowitz, eds., *Quantum Computing: Progress and Prospects* (Washington, DC: The National Academies Press, 2019), 3–3–3–5, <https://doi.org/10.17226/25196>.

13. More precisely, Shor’s algorithm provides an exponential optimization for an important mathematical technique that has widespread application: “the [classical] fast Fourier transform (FFT), takes  $O(N \log N)$  time, which is only slightly longer than it takes to read the input data  $[O(N)]$ . While the classical FFT is quite efficient, quantum Fourier transform (QFT) is exponentially faster, requiring only  $O(\log^2 N) = O(n^2)$  time (where  $N = 2^n$ ) in its original formulation [by Shor], later improved to  $O(n \log n)$ ,” according to Grumbling and Horowitz, *Quantum Computing*, 3–3.

14. For example, a single bit (b) might be copied as three bits (bbb) with the value determined by a majority voting algorithm.

15. Stephen Wiesner, "Conjugate Coding," *SIGACT News* 15, no. 1 (January 1983): 78–88, <https://doi.org/10.1145/1008908.1008920>. The perfect copying of an arbitrary qubit  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$  would imply that a transformation exists such that  $|\psi\rangle|0\rangle \rightarrow |\psi\rangle|\psi\rangle = (\alpha|0\rangle + \beta|1\rangle)(\alpha|0\rangle + \beta|1\rangle) = \alpha^2|00\rangle + \alpha\beta(|01\rangle + |10\rangle) + \beta^2|11\rangle$ . This equation has quadratic terms, but all quantum computations must be linear transformations. This contradiction implies that copying an arbitrary qubit is not possible.

16. For example, the simulation of FeMoco in Nitrogenase would require only 111 logical qubits but 180 million physical qubits, assuming an error rate of 1 in a thousand per physical qubit; assuming higher quality physical qubits with error rates as low as 1 in a billion, the 111 logical qubits in this simulation would still require 230 thousand physical qubits. Grumbling and Horowitz, *Quantum Computing*, 3–15.

17. One implementation of Shor's algorithm requires  $2n+2$  logical qubits but requires more computational resources because it has a large number of T-gates ( $448n^3 \log 2n$ ) and high T-depth ( $480n^3$ ); a different implementation requires more qubits ( $4n - \log 2n$ ) but has better performance with  $420n^3$  T-gates and a T-depth of  $72n^2 \log 2n$ . The slight gains of Grover's algorithm against the Advanced Encryption Standard (AES) require even more resources: a speedup against AES-256 would require 6,681 logical qubits; an application of Grover's algorithm against Secure Hash Algorithm (SHA) 3-256, again for only a minor polynomial speedup, would require 3,200 logical qubits. See Martin Roetteler and Krysta M. Svore, "Quantum Computing: Codebreaking and Beyond," *IEEE Security Privacy* 16, no. 5 (September 2018): 22–36, <https://doi.org/10.1109/MSP.2018.3761710>.

18. Craig Gidney and Martin Eker, "How to Factor 2048 Bit RSA Integers in 8 Hours Using 20 Million Noisy Qubits" arXiv, 6 December 2019, <https://arxiv.org/>.

19. Swamit S. Tannu and Moinuddin K. Qureshi, "Not All Qubits Are Created Equal: A Case for Variability-Aware Policies for NISQ-Era Quantum Computers," preprint, arXiv:805.10224, May 2018, <https://arxiv.org/>. A fully functional logical qubit should support superposition, entanglement, and measurement with any valid quantum algorithm.

20. Stephen P. Jordan and Yi-Kai Liu, "Quantum Cryptanalysis: Shor, Grover, and Beyond," *IEEE Security and Privacy* 16, no. 5 (September 2018): 14–21, DOI: 10.1109/MSP.2018.3761719.

21. Troels F. Rønnow et al., "Defining and Detecting Quantum Speedup," *Science* 345, no. 6195 (25 July 2014): 420–24; and Philip Ball, "The Era of Quantum Computing Is Here. Outlook: Cloudy," *Quanta Magazine*, 24 January 2018, <https://www.quanta-magazine.org/>.

22. That is,  $0^\circ$  and  $90^\circ$  map to  $|0\rangle$  and  $|1\rangle$  while  $45^\circ$  and  $135^\circ$  map to  $|+\rangle$  and  $|-\rangle$ .

23. On cryptography, see Anne Broadbent and Christian Schaffner, "Quantum Cryptography beyond Quantum Key Distribution," *Designs, Codes and Cryptography* 78, no. 1 (January 2016): 351–82, <https://doi.org/10.1007/s10623-015-0157-4>. On national security applications, see Michael J. Biercuk and Richard Fontaine, "The Leap into Quantum Technology: A Primer for National Security Professionals," *War on the Rocks* (blog), 17 November 2017, <https://warontherocks.com/>.

#### Disclaimer and Copyright

The views and opinions in *SSQ* are those of the authors and are not officially sanctioned by any agency or department of the US government. This document and trademarks(s) contained herein are protected by law and provided for noncommercial use only. Any reproduction is subject to the Copyright Act of 1976 and applicable treaties of the United States. The authors retain all rights granted under 17 U.S.C. §106. Any reproduction requires author permission and a standard source credit line. Contact the *SSQ* editor for assistance: [strategicstudiesquarterly@us.af.mil](mailto:strategicstudiesquarterly@us.af.mil).