Post-quantum cryptography and the quantum future of cybersecurity

Yi-Kai Liu[1,2,*] and Dustin Moody[1]

[1] *National Institute of Standards and Technology (NIST), Gaithersburg, Maryland 20899, USA*
[2] *Joint Center for Quantum Information and Computer Science (QuICS), NIST/University of Maryland,*
College Park, Maryland 20742, USA

(Received 31 March 2023; revised 21 December 2023; published 9 April 2024)

We review the current status of efforts to develop and deploy post-quantum cryptography on the Internet. Then we suggest specific ways in which quantum technologies might be used to enhance cybersecurity in the near future and beyond. We focus on two goals: protecting the secret keys that are used in classical cryptography, and ensuring the trustworthiness of quantum computations. These goals may soon be within reach, thanks to recent progress in both theory and experiment. This progress includes interactive protocols for testing quantumness as well as for performing uncloneable cryptographic computations; and experimental demonstrations of device-independent random number generators, device-independent quantum key distribution, quantum memories, and analog quantum simulators.

DOI: 10.1103/PhysRevApplied.21.040501

I. INTRODUCTION

The Internet relies heavily on both public-key encryption schemes and digital signatures to ensure the confidentiality and authenticity of digital communications. However, many of these widely used cryptosystems could be broken by quantum algorithms, running on large-scale fault-tolerant quantum computers. Such machines do not yet exist, but could conceivably be built in the not-too-distant future.

To mitigate this potential threat, a large community of researchers is working to develop so-called *post-quantum cryptosystems*, to provide security against adversaries who have quantum computers. These post-quantum cryptosystems will be deployed on the Internet and in a wide range of devices. For some applications, this deployment needs to occur soon, well before the development of truly large-scale quantum computers, due to the concern that encrypted communications could be recorded today, and broken in the future.

A recent milestone in this process was the selection by the National Institute of Standards and Technology (NIST) of several schemes for public-key encryption and digital signatures, which will become US Government standards, and are expected to be widely adopted [1]. This marks the beginning of one of the most complex transitions in the history of the Internet, as these post-quantum cryptosystems are deployed in a myriad of different communications protocols, software applications, and hardware devices.

This transition to post-quantum cryptography will address most of the cybersecurity concerns that arise in a world where quantum technologies are widespread. However, there are a number of cybersecurity problems that are *not* solved by post-quantum cryptography. In this perspective, we will discuss some of these problems. We will then suggest some ways that *quantum* technologies may help to solve these problems, in the near future and beyond.

II. STRENGTHS AND WEAKNESSES OF POST-QUANTUM CRYPTOGRAPHY

We begin by describing the functionalities provided by post-quantum cryptosystems, how they are used on the Internet, and what are the weaknesses of these schemes, which quantum technologies might be able to address.

A. Securing the Internet

At the most basic level, the Internet requires some means to enable users to communicate privately, to protect their communications from being altered by unauthorized parties, and to enable users to verify all others' identities. These are provided by a combination of several cryptographic tools: public-key encryption schemes and block ciphers (which ensure privacy), and digital signatures and hash functions (which protect the authenticity and integrity of data).

Public-key encryption schemes and digital signatures play a special role on the Internet, because they can be used by parties that share only public information, i.e., parties that do not already possess a shared, private key. This feature is crucial for supporting communications among a large number of users. In these schemes, every user has two different keys, a *public key* and a *secret key*, which perform different functions, such as encryption and decryption, or

[*] yi-kai.liu@nist.gov

signing and verifying. As their names would suggest, the public key can be known by everyone, while the secret key should only be known by a single user.

Public-key encryption is often used for *key establishment*, that is, to set up encrypted communications sessions between users on the Internet. This is done in the following way: First, two users use each other's public keys to send encrypted messages to each other. By combining the contents of these messages, they can generate a shared secret key. They then use this shared secret key to encrypt their subsequent messages, using a block cipher. Here, public-key encryption is needed for the first step, but block ciphers are used to encrypt the subsequent messages, as block ciphers are much more efficient. This whole process is usually managed by a higher-level protocol, such as the Transport Layer Security (TLS) protocol.

A common use of digital signatures is to protect electronic files and software from being altered by an adversary. In addition, digital signatures are used to construct *certificates*, which are used to distribute public keys on the Internet in a trustworthy manner. Essentially, a certificate is a way of linking a public key pk to some entity E in the "real world," such as a person or a business with a physical address. The certificate is signed by a *certificate authority* (CA) – some organization that is able to verify that the public key pk indeed belongs to the entity E. Hence, a certificate provides a way for E to distribute its public key, which can be verified by anyone who already has the CA's public key.

Of course, this still leaves another problem: How does one get a trustworthy copy of the CA's public key? Sometimes this can come from another certificate, signed by some other "higher-level" CA. The highest-level CAs are called "root CAs," and their public keys reside on "root certificates," which are often built into operating systems and web browsers.

This system for managing public keys is sometimes called *public-key infrastructure* (PKI). It can scale up to large numbers of users, but it can also be vulnerable to attacks on certificate authorities (e.g., when an attacker gains access to a certificate authority's secret key, and uses it to generate fraudulent certificates). Such attacks have actually occurred, despite the elaborate safeguards used by certificate authorities to protect their secret keys [2]. As a result, the Internet community has deployed additional infrastructure to identify and revoke certificates that are compromised or fraudulent. This is sometimes called "certificate transparency" [3].

B. Post-quantum cryptography

Unfortunately, essentially all of the currently deployed public-key encryption schemes and digital signatures – such as RSA (Rivest-Shamir-Adleman), Diffie-Hellman, and elliptic curve cryptosystems – will be vulnerable to quantum attacks, using Shor's algorithm and its generalizations, if large-scale fault-tolerant quantum computers are built in the future. These attacks run in (quantum) polynomial time, and cannot be mitigated simply by using larger keys.

For comparison, quantum attacks on block ciphers and hash functions also exist, using techniques such as Grover's algorithm. However, in almost all cases, these attacks still take exponential time, and they can be mitigated by using larger keys or longer output, or adjusting other parameters of these cryptographic constructions.

The main goal of post-quantum cryptography, then, is to develop new schemes for public-key encryption and digital signatures, which are based on different computational problems that appear to be hard for quantum computers. A wide variety of post-quantum cryptosystems have been proposed in recent years, based on problems involving high-dimensional lattices, error-correcting codes, multivariate quadratic equations over finite fields, isogenies between elliptic curves, collisions in hash functions, and others. There are many technical problems in designing these cryptosystems, assessing their security, and deploying them on the Internet. Many of these problems have been studied intensely in recent years, leading up to the selection of the first post-quantum cryptosystems to be standardized by NIST [1].

At the same time, there has also been longstanding interest in *quantum key distribution* (QKD) [4]. Like public-key cryptography, QKD also enables two users to exchange keys, but its security is based on properties of quantum physics, rather than computational complexity. QKD cannot replace public-key cryptography, but it might nonetheless be useful in certain scenarios.

QKD has certain strengths and weaknesses. It is information-theoretically secure (at least in principle), meaning that its security does not rely on conjectures about the hardness of solving certain problems on quantum computers. However, QKD has difficulty operating over long distances (necessitating additional infrastructure, such as trusted repeater stations or quantum networks), and it has potential vulnerabilities to side-channel attacks. For more discussion of these issues, we refer the reader to Refs. [4–7].

In this paper, we will take a broader view. Assuming that post-quantum cryptosystems work as intended, what are the outstanding problems in cybersecurity that still need to be solved?

C. Outstanding problems

There are a number of problems in cybersecurity that are not solved by post-quantum cryptography. One such problem is that real software and hardware have a variety of imperfections, which can leak partial information about the secret key to an adversary.

Some of these imperfections are caused by human error (e.g., software and hardware bugs, or discrepancies between the implementation and the theoretical specification of the cryptosystem). To some extent, this can be addressed by testing and formal verification methods (e.g., software tools such as EasyCrypt and Cryptol) [8].

Other imperfections are caused by the physical properties of the hardware, such as data-dependent variations in the time or power needed to perform a cryptographic computation, and other kinds of “side channels.” To some extent, this can be addressed by “constant-time” or “masked” implementations of cryptographic algorithms, as well as other countermeasures [9–12].

Because of these concerns, it is common to use special-purpose hardware to perform especially sensitive cryptographic computations. There are many kinds of special-purpose cryptographic hardware, such as cryptographic coprocessors, hardware security modules, trusted platform modules, physical uncloneable functions, cryptographic ignition keys, and two-factor authentication devices [13–15]. Building such devices, and verifying their security, is a complex problem, particularly for users who require very high levels of security, such as certificate authorities.

A special case of this problem involves the handling of secret keys [16]. Here, hardware *random number generators* (RNGs) are often used to generate secret keys, and provide random bits for encryption and signing operations. The unpredictability of these random bits is critical for security. However, testing and certification of hardware RNGs is a difficult task, both in theory and in practice [17]. Additional complications arise when secret keys have to be backed up, transported from one physical location to another, or updated (e.g., for stateful hash-based signatures) [18].

In recent years, there has been spectacular progress in developing quantum technologies. If this progress continues, quantum hardware might someday provide novel solutions to the above problems, with theoretical guarantees of security that are not possible for hardware based on classical physics [19,20]. Instead, these theoretical guarantees rely on quantum phenomena, such as self-testing and uncloneability. We will describe these ideas in the following section.

III. QUANTUM TECHNOLOGIES FOR PROTECTING SECRET KEYS

In this section, we will describe some techniques in quantum cryptography that can be applied to the construction and verification of secure hardware for cryptography. We will briefly review some of the relevant theory literature, and then discuss the prospects for experimental realizations of these ideas.

A. Quantum mechanics, uncloneability, and verification

Here we will focus on cryptographic protocols that deal with *uncloneability* of quantum information, and *verification* of untrusted quantum devices.

The term “uncloneability” comes from the quantum no-cloning theorem, which states that it is impossible to make a perfect copy of an unknown quantum state. In the context of quantum cryptography, uncloneability refers to protocols that protect a piece of information that has a specific function – such as a secret key, a message, or a ciphertext – from being copied or forged by an adversary. Examples of such protocols include quantum key distribution, quantum money, uncloneable encryption, and quantum copy protection [4,21–33].

These protocols are useful for constructing secure hardware to store secret keys and perform cryptographic computations, because these protocols provide a mechanism for preventing the secret keys from being copied or extracted from the hardware. This behavior is an important requirement for secure cryptographic hardware, yet it is difficult to achieve in classical hardware, because there is no principle in classical physics that prevents the copying of information. But quantum hardware can achieve this kind of uncloneability, with very strong guarantees of security (at least in principle).

The precise meaning of “uncloneability” depends on the kind of cryptographic functionality that is being implemented. Recent research has expanded the range of uncloneable functionalities that can be achieved, to include secure software leasing, certified deletion, one-time programs, one-time memories, and one-time signatures [34–48]. A variety of techniques can be used to enforce uncloneability, including quantum communication with single-qubit states [23,24], storage of single-qubit states combined with secure hardware assumptions [26,40,41,47,48], and quantum computation combined with complexity-theoretic hardness assumptions [25,27,28,32].

Next, consider the problem of “verifying” an untrusted quantum device. This problem is related to basic questions about the foundations of quantum mechanics. For instance, how does one know that a quantum superposition state really exists, if the process of measuring the state causes the superposition to collapse? In the context of cryptography, one can ask an analogous question: How can one use classical communication to test whether an untrusted quantum device is operating correctly?

These questions can be answered in a variety of ways [49]. One approach is to separate the untrusted quantum device into two components, which are entangled but nonsignaling. Then one can use Bell tests and nonlocal games to construct cryptographic protocols for device-independent random number generation and quantum self-testing [50–58]. Some of these methods are closely related to device-independent quantum key distribution [59]. Another approach uses a verifier that can prepare single-qubit states, to construct interactive protocols for blind or delegated quantum computing [60,61]. A third approach uses a completely classical verifier, and quantum algorithms that run on the untrusted quantum device (using cryptographic techniques such as trapdoor claw-free functions), to construct computationally secure protocols for testing quantumness and delegating quantum computation [62–64].

These techniques are potentially useful for the testing and validation of cryptographic hardware, as well as quantum computers. As an example, consider how quantum experiments that violate Bell's inequality can be used to generate random numbers [51,52,56]. Here, one can view Bell's inequality as a *test* of the experimental apparatus, which *certifies* that the output of the experiment is truly random. More precisely, if the output of the experiment violates Bell's inequality, then it is incompatible with any "local realistic" model of the universe, and hence it contains entropy that is independent of all prior information in the universe [65].

The test described above is very strong: it holds for *any* experiment that consists of two nonsignaling devices, and it relies only on the input and output of the devices, and not on their internal functioning (which may be quantum, classical, or even adversarial). This is called *device-independent* security. Furthermore, if the output of the experiment violates Bell's inequality to a sufficiently large degree, this uniquely determines the behavior of the two nonsignaling devices (up to a small error, and internal degrees of freedom that do not affect the input and output of the devices). This property is called *rigidity*, or *self-testing*, and it provides a powerful mechanism for controlling the devices, and performing quantum computations [50,54].

These kinds of tests have numerous applications in cryptography. They can be used to certify hardware random number generators, which are used to generate secret keys, and to encrypt and sign messages. Compared to conventional methods [17], these tests provide a much stronger guarantee on the quality of a hardware RNG. In principle, these kinds of tests can even be used to certify arbitrary quantum computations performed on untrusted hardware. The main drawback is that quantum hardware is needed to pass these tests. The complexity of this hardware varies according to the task being performed, as we will discuss in the next section.

All of these topics are the subject of ongoing research, and our discussion here only covers a sample of the work in this area. In particular, while we have focused our discussion on quantum protocols involving unclonebility and verification, some of these techniques can also be used for other purposes, such as privacy and transparency [66]. More comprehensive surveys can be found in Refs. [19,20,67].

B. Outlook: quantum secure hardware

We now discuss the prospects for using these quantum techniques in the real world. As a first step, we focus on Bell tests and device-independent random number generators. Many of these techniques have already been demonstrated, at least as a proof of principle, on research-grade experimental hardware [51,68–70]. Recent demonstrations of device-independent QKD give a sense of the current state of the art [71–73].

Here, we will suggest some possible paths toward using these techniques in commercial-grade cryptographic hardware, through a series of incremental upgrades to present-day *hardware security modules* (HSMs). Some of these hardware upgrades will require further advances in quantum technologies, which have yet to be demonstrated. Nonetheless, the technologies we require are a subset of the technologies needed to build large-scale fault-tolerant quantum computers and networks. We are suggesting that this subset of technologies, by itself, can have interesting applications in cybersecurity.

The first step is to augment these HSMs with commercially available systems for quantum key distribution. These systems use photon detectors based on single-photon avalanche diodes (SPADs), which are readily available and easy to operate. These detectors should be placed inside the "security perimeter" of the HSMs.

Two such HSMs can be connected to a source of entangled photons, located outside the "security perimeter" of the HSMs, in order to perform a Bell test, which is the basic ingredient for a quantum entanglement-based RNG. In particular, this can be achieved using an entangled photon source based on spontaneous parametric down-conversion, which can be fairly inexpensive.

The main technical challenge, then, is to improve the security of this scheme, by patching various vulnerabilities that could be exploited by an adversary (while keeping the cost of the quantum hardware as low as possible). These security vulnerabilities are called *loopholes* in the context of Bell tests, and they have been studied intensively in recent years, leading up to the first demonstrations of "loophole-free" Bell tests [74–77]. In the context of device-independent RNGs, these loopholes are ways that an adversary can cause the HSMs to "pass" the Bell test, while producing an output that is predictable to the adversary (and hence is not truly random).

There are two main loopholes. One is the "signaling loophole," where an adversary causes the HSMs to communicate with each other, in order to generate correlations that can "cheat" the Bell test. In practice, this is prevented by electromagnetic shielding on the HSMs.

The other is the "detection loophole," where an adversary exploits the fact that a significant fraction of the entangled photons are not registered by the single-photon detectors in the HSMs. By causing the single-photon detectors to fail in a biased way, the adversary can "cheat" the Bell test.

The detection loophole can be closed by using matter qubits, such as trapped ions, nitrogen vacancy centers, or trapped neutral atoms [51,74,77], or by using high-efficiency superconducting nanowire single-photon detectors (SNSPDs) [68,69]. In most situations, the latter approach is the least costly. However, SNSPDs are still significantly more expensive than commonly available SPAD detectors. Is there any way to reduce this cost further?

One intriguing possibility is to implement device-independent RNG using an *asymmetric* or *hybrid* Bell test [78,79], which allows one HSM to have lower detection efficiency, provided that the other HSM has near-perfect detection efficiency. This approach has been studied in the context of loophole-free Bell tests [80,81], and would be easier to implement in our setting, where we are concerned only with the detection loophole, and not with the signaling loophole. In concrete terms, this scheme might be implemented using *atom-photon* entanglement, i.e., using one HSM that contains a matter qubit (which is expensive), and another HSM that uses an inexpensive SPAD detector combined with continuous-variable homodyne measurement (which is also inexpensive).

Using atom-photon entanglement has certain drawbacks, as one must deal with two sets of technical difficulties: those related to entangling atoms (e.g., needing to cool and stabilize an atom, and needing to collect the single photon emitted by the atom); and those related to working with photons (e.g., the problems of photon loss, and low photon detection efficiency). In addition, some of the techniques that are useful for entangling atoms, such as heralding and postselection, are no longer feasible when working with atom-photon entanglement. Thus, realizing these hybrid Bell tests will require improvements on the current state of the art [82].

However, these hybrid Bell tests may have an important practical advantage: they may allow one party to use a low-cost HSM that is outfitted with equipment that resembles present-day QKD hardware. Thus, one can deploy a large fleet of these low-cost HSMs now, and use them for QKD. At a future date, one may be able to upgrade these low-cost HSMs to perform device-independent random number generation (or device-independent QKD), by augmenting them with a few more-expensive HSMs that are equipped with matter qubits capable of generating matter-photon entanglement.

This scenario is reminiscent of the public-key infrastructure of the Internet, where a few root CAs generate certificates for a larger number of lower-level CAs, who generate certificates for an even larger number of users. This is an example of a familiar trick in systems engineering: amortizing the cost of the most expensive part of a system over many uses of the system.

The ideas described above could provide a highly secure way to generate secret keys within (a pair of) HSMs. What might be the next step along this path?

A plausible next step would be to equip an HSM with *unentangled* quantum memories, which are not capable of performing entangling operations, but can nonetheless be used to implement certain "uncloneable" cryptographic functionalities, such as quantum money, one-time signatures, or one-time programs. This takes advantage of the fact that, while many of these cryptographic protocols require the ability to perform quantum computation, there exist simple versions of these protocols that only require unentangled quantum memories or isolated qubits, and can be highly noise-tolerant [26,40,41,43,44,47,48].

These single-qubit quantum memories would require significant improvements on existing quantum technologies. But they would avoid the even greater technical challenges of building a true quantum computer.

In particular, these kinds of primitive quantum memories can be built using a variety of physical systems, including solid-state systems [83]. Some of these systems have the potential to achieve hour-long coherence times, even at room temperature [84–87]. Finally, these systems are relatively easy to scale up, compared to true quantum computing hardware [88], because they do not require two-qubit couplings or interactions, and they can tolerate substantial variability in the physical properties of the individual qubits.

IV. TRUSTWORTHY QUANTUM COMPUTATIONS

This brings us to a second problem in cybersecurity that is not solved by post-quantum cryptography alone: How does one ensure the trustworthiness of quantum computations? As we mentioned earlier, many theoretical solutions to this problem have been proposed [49], including quantum self-testing [54], blind or delegated quantum computing [60,61], and computationally secure interactive protocols [62–64]. Here, we suggest some ways that these ideas might be applied to the kinds of quantum computations that will be useful for science and engineering, in the near future and beyond.

A. Random samples, and analog quantum simulators

As a first example, consider the recent experimental demonstrations of "quantum computational advantage" [89,90]. In these experiments, complex quantum superposition states were prepared, and then measured, producing *random samples* from a complicated probability distribution. It is difficult to verify the correctness of these samples, and, in some cases, it is actually possible for a classical adversary to “spoof” the distribution, without being detected by commonly used verification tests [91]. Although there are stronger interactive protocols that can in principle prevent such spoofing [92], this remains a significant concern for these types of experiments [93].

Perhaps the next milestone in the development of these technologies will be to perform the first quantum simulations that give useful insight into condensed matter physics or quantum chemistry [94]. Here, we use the term *quantum simulation* to mean a simulation of the time evolution of some quantum system that appears in nature, where this simulation is performed on a *quantum simulator* – a quantum device that can be programmed to simulate many different quantum systems.

Different kinds of quantum simulators have different degrees of programmability. At one extreme, *analog* quantum simulators typically implement a many-body Hamiltonian with many quantum particles, but relatively few tuneable degrees of freedom. At the opposite extreme, highly programmable *digital* quantum simulations can be performed by encoding the state of a quantum system into a quantum computer, and running a quantum circuit that simulates the system’s time evolution.

Recent progress, particularly in analog quantum simulators [95], raises an interesting question: How can one verify the correctness of these simulations, as they continue to grow in complexity? This question is hard, for two reasons. First, these results will likely be difficult to reproduce using simulations performed on a classical computer, or empirical measurements of a quantum system occurring in nature. Second, analog quantum simulators have too few degrees of freedom to perform arbitrary quantum computations, and thus are unable to run many of the verification protocols described above, such as blind quantum computing, or protocols involving trapdoor claw-free functions.

One possible solution could be to construct “trapdoor Hamiltonians” that appear to have complicated dynamics, but can be simulated on classical computers if one knows some special information. We call this information a trapdoor, by analogy with trapdoor one-way functions in cryptography. For instance, a trapdoor could consist of a unitary change of basis that reveals some special structure in the Hamiltonian, such as “stoquasticity,” which may make the Hamiltonian easier to simulate [96,97]. Thus, the trapdoor could provide an efficient means of checking the correctness of analog quantum simulations that were performed without knowledge of the trapdoor. This is reminiscent of proposals for trapdoor simulation of quantum circuits [98].

B. Protecting quantum data

In the more distant future, digital quantum simulations using large, fault-tolerant quantum computers might be able to reproduce some of the complex quantum states that exist in nature, but in an artificial form (i.e., encoded into qubits) that can be measured in far greater detail than a natural quantum system. Such quantum states could be the first real example of *quantum data*.

These quantum states might someday be used as scientific reference materials, in the same way that classical databases of chemical and material properties are used today. In particular, these quantum states could be more useful than any classical information derived from them, such as expectation values of observables. This is because, given such a quantum state ρ, one can use a quantum computer to perform *any* efficiently computable measurement on ρ.

Providing access to these quantum states is an important motivation for developing quantum networks and quantum memories [99]. Perhaps, new fully quantum methods will be used to verify these quantum states, and protect their integrity. These might include “swap tests” that (nondestructively) compare one quantum state to another, via an entangled measurement [100], as well as encryption and signatures of arbitrary quantum states [101].

These ideas demonstrate the rich interaction between classical and quantum cryptography, and the needs of cybersecurity involving both classical and quantum computers. This is the quantum future of cybersecurity, and we believe that we are already on the path toward this future.

ACKNOWLEDGMENTS

We thank Dominik Hangleiter, Atul Mantri, the NIST PQC team, and an anonymous reviewer for helpful discussions and comments. Any opinions expressed here are those of the authors, and not of NIST or the University of Maryland. Any mention of commercial products is for information only, and does not indicate endorsement.

[1] G. Alagic, D. Apon, D. Cooper, Q. Dang, T. Dang, J. Kelsey, J. Lichtinger, C. Miller, D. Moody, R. Peralta, R. Perlner, A. Robinson, D. Smith-Tone, and Y.-K. Liu, Status report on the third round of the NIST post-quantum cryptography standardization process, NISTIR 8413, National Institute of Standards and Technology, 2022. https://doi.org/10.6028/NIST.IR.8413-upd1.

[2] J. Wolff, How a 2011 hack you’ve never heard of changed the internet’s infrastructure. *Slate* (website), https://slate.com/technology/2016/12/how-the-2011-hack-of-diginotar-changed-the-internets-infrastructure.html, 2016.

[3] “Certificate transparency.” Community website, https://certificate.transparency.dev/.

[4] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev, The security of practical quantum key distribution, Rev. Mod. Phys. **81**, 1301 (2009).

[5] C. Portmann and R. Renner, Security in quantum cryptography, Rev. Mod. Phys. **94**, 025008 (2022).

[6] "NSA guidance on quantum key distribution (QKD) and quantum cryptography (QC)." https://www.nsa.gov/Cybersecurity/Quantum-Key-Distribution-QKD-and-Quantum-Cryptography-QC/.

[7] R. Renner and R. Wolf, The debate over QKD: A rebuttal to the NSA's objections, ArXiv eprint arxiv:2307.15116 (2023).

[8] G. Barthe, F. Dupressoir, B. Grégoire, C. Kunz, B. Schmidt, and P.-Y. Strub, in *Foundations of Security Analysis and Design VII* (2013), p. 146.

[9] P. C. Kocher, in *Annual International Cryptology Conference (CRYPTO)* (Springer Berlin, Heidelberg, 1996), p. 104.

[10] S. Chari, C. S. Jutla, J. R. Rao, and P. Rohatgi, in *Annual International Cryptology Conference (CRYPTO)* (Springer Berlin, Heidelberg, 1999), p. 398.

[11] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, J. Horn, S. Mangard, P. Kocher, D. Genkin, and Y. Yarom *et al.*, Meltdown: Reading kernel memory from user space, Commun. ACM **63**, 46 (2020).

[12] P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, and T. Prescher *et al.*, Spectre attacks: Exploiting speculative execution, Commun. ACM **63**, 93 (2020).

[13] S. Pearson and B. Balacheff, *Trusted Computing Platforms: TCPA Technology in Context*, Hewlett-Packard professional books (Prentice Hall PTR, 2003).

[14] S. Smith, *Trusted Computing Platforms: Design and Applications* (Springer, Berlin, Heidelberg, 2013).

[15] "OpenTitan: Open source silicon root of trust." Community website, https://opentitan.org/.

[16] E. Barker, "Recommendation for key management: Part 1 – general," SP 800-57 Part 1 Rev. 5, National Institute of Standards and Technology, (2020). https://doi.org/10.6028/NIST.SP.800-57pt1r5.

[17] M. S. Turan, E. Barker, J. Kelsey, K. McKay, M. Baish, and M. Boyle, Recommendation for the entropy sources used for random bit generation, SP 800-90B, National Institute of Standards and Technology (2018). https://doi.org/10.6028/NIST.SP.800-90B.

[18] D. Cooper, D. Apon, Q. Dang, M. Davidson, M. Dworkin, and C. Miller, Recommendation for stateful hash-based signature schemes, SP 800-208, National Institute of Standards and Technology (2020). https://doi.org/10.6028/NIST.SP.800-208.

[19] A. Broadbent and C. Schaffner, Quantum cryptography beyond quantum key distribution, Designs, Codes and Cryptography **78**, 351 (2016).

[20] P. Wallden and E. Kashefi, Cyber security in the quantum era, Commun. ACM **62**, 120 (2019).

[21] S. Wiesner, Conjugate coding, ACM Sigact News **15**, 78 (1983).

[22] C. H. Bennett, G. Brassard, S. Breidbart, and S. Wiesner, in *Advances in Cryptology (CRYPTO)* (Springer Berlin, Heidelberg, 1983), p. 267.

[23] C. Bennett and G. Brassard, in *Proc. IEEE Int. Conf. Computers, Systems, and Signal Processing, Bangalore, India, 1984* (IEEE Computer Society, IEEE Circuits and Systems Society, and Indian Institute of Science, 1984), p. 175.

[24] D. Gottesman, Uncloneable encryption, arXiv preprint arXiv:quant-ph/0210062, (2002).

[25] S. Aaronson, in *2009 24th Annual IEEE Conference on Computational Complexity* (IEEE Computer Society, Los Alamitos, CA, 2009), p. 229.

[26] F. Pastawski, N. Y. Yao, L. Jiang, M. D. Lukin, and J. I. Cirac, Unforgeable noise-tolerant quantum tokens, Proc. Natl. Acad. Sci. **109**, 16079 (2012).

[27] E. Farhi, D. Gosset, A. Hassidim, A. Lutomirski, and P. Shor, in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (Association for Computing Machinery, New York, 2012), p. 276.

[28] S. Aaronson and P. Christiano, in *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2012), p. 41.

[29] D. Gavinsky, in *2012 IEEE 27th Conference on Computational Complexity* (IEEE Computer Society, Los Alamitos, CA, 2012), p. 42.

[30] E. Diamanti, H.-K. Lo, B. Qi, and Z. Yuan, Practical challenges in quantum key distribution, Npj Quantum Inf. **2**, 1 (2016).

[31] R. Radian and O. Sattath, in *Proceedings of the 1st ACM Conference on Advances in Financial Technologies* (Association for Computing Machinery, New York, 2019), p. 132.

[32] M. Zhandry, in *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)* (Springer, 2019), p. 408.

[33] O. Shmueli, in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2022), p. 790.

[34] P. Ananth and R. L. L. Placa, in *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)* (Springer Cham, 2021), p. 501.

[35] A. Broadbent and R. Islam, in *Theory of Cryptography Conference* (Springer Cham, 2020), p. 92.

[36] T. Hiroka, T. Morimae, R. Nishimaki, and T. Yamakawa, in *Annual International Cryptology Conference (CRYPTO)* (Springer Cham, 2022), p. 239.

[37] J. Bartusek and D. Khurana, in *Annual International Cryptology Conference (CRYPTO)* (Springer Cham, 2023), p. 192.

[38] A. Poremba, in *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)* (Schloss Dagstuhl-Leibniz-Zentrum für Informatik, Wadern, Germany, 2023).

[39] A. Broadbent, G. Gutoski, and D. Stebila, in *Annual Cryptology Conference (CRYPTO)* (Springer Berlin, Heidelberg, 2013), p. 344.

[40] Y.-K. Liu, in *Annual Cryptology Conference (CRYPTO)* (Springer Berlin, Heidelberg, 2014), p. 19.

[41] Y.-K. Liu, in *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)* (Springer Berlin, Heidelberg, 2015), p. 785.

[42] S. Ben-David and O. Sattath, Quantum tokens for digital signatures, arXiv preprint arXiv:1609.09047, (2016).

[43] M.-C. Roehsner, J. A. Kettlewell, T. B. Batalhão, J. F. Fitzsimons, and P. Walther, Quantum advantage for probabilistic one-time programs, Nat. Commun. **9**, 1 (2018).

[44] K.-M. Chung, M. Georgiou, C.-Y. Lai, and V. Zikas, Cryptography with disposable backdoors, Cryptography **3**, 22 (2019).

[45] R. Amos, M. Georgiou, A. Kiayias, and M. Zhandry, in *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2020), p. 255.

[46] A. Coladangelo, C. Majenz, and A. Poremba, Quantum copy-protection of compute-and-compare programs in the quantum random oracle model, arXiv preprint arXiv:2009.13865, (2020).

[47] A. Broadbent, S. Gharibian, and H.-S. Zhou, Towards quantum one-time memories from stateless hardware, Quantum **5**, 429 (2021).

[48] Q. Liu, in *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)* (Schloss Dagstuhl-Leibniz-Zentrum für Informatik, Wadern, Germany, 2023).

[49] A. Gheorghiu, T. Kapourniotis, and E. Kashefi, Verification of quantum computation: An overview of existing approaches, Theory Computing Syst. **63**, 715 (2019).

[50] D. Mayers and A. Yao, Self testing quantum apparatus, Quantum Inf. Comput. **4**, 273 (2004).

[51] S. Pironio, A. Acín, S. Massar, A. B. de La Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, and C. Monroe, Random numbers certified by Bell's theorem, Nature **464**, 1021 (2010).

[52] R. Colbeck and A. Kent, Private randomness expansion with untrusted devices, J. Phys. A: Math. Theor. **44**, 095305 (2011).

[53] R. Colbeck and R. Renner, Free randomness can be amplified, Nat. Phys. **8**, 450 (2012).

[54] B. W. Reichardt, F. Unger, and U. Vazirani, Classical command of quantum systems, Nature **496**, 456 (2013).

[55] C. A. Miller and Y. Shi, Robust protocols for securely expanding randomness and distributing keys using untrusted quantum devices, J. ACM (JACM) **63**, 1 (2016).

[56] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner, Bell nonlocality, Rev. Mod. Phys. **86**, 419 (2014).

[57] A. Acín and L. Masanes, Certified randomness in quantum physics, Nature **540**, 213 (2016).

[58] I. Šupić and J. Bowles, Self-testing of quantum systems: A review, Quantum **4**, 337 (2020).

[59] U. Vazirani and T. Vidick, Fully device independent quantum key distribution, Commun. ACM **62**, 133 (2019).

[60] A. Broadbent, J. Fitzsimons, and E. Kashefi, in *2009 50th Annual IEEE Symposium on Foundations of Computer Science* (IEEE Computer Society, Los Alamitos, CA, 2009), p. 517.

[61] D. Aharonov, M. Ben-Or, E. Eban, and U. Mahadev, Interactive proofs for quantum computations, arXiv preprint arXiv:1704.04487 (2017).

[62] Z. Brakerski, P. Christiano, U. Mahadev, U. Vazirani, and T. Vidick, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society, Los Alamitos, CA, 2018), p. 320.

[63] U. Mahadev, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society, Los Alamitos, CA, 2018), p. 259.

[64] A. Cojocaru, L. Colisson, E. Kashefi, and P. Wallden, in *International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)* (Springer Cham, 2019), p. 615.

[65] There is a slight caveat: In order to run this test, one needs to provide a random input. Hence this scheme cannot generate randomness in a world that is completely deterministic, but it can generate new randomness from old randomness.

[66] A. Coladangelo, S. Goldwasser, and U. Vazirani, in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2022), p. 1378.

[67] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, J. L. Pereira, M. Razavi, J. Shamsul Shaari, M. Tomamichel, V. C. Usenko, G. Vallone, P. Villoresi, and P. Wallden, Advances in quantum cryptography, Adv. Opt. Photonics **12**, 1012 (2020).

[68] M. Giustina, A. Mech, S. Ramelow, B. Wittmann, J. Kofler, J. Beyer, A. Lita, B. Calkins, T. Gerrits, S. W. Nam, Rupert Ursin, and Anton Zeilinger, Bell violation using entangled photons without the fair-sampling assumption, Nature **497**, 227 (2013).

[69] B. G. Christensen, K. T. McCusker, J. B. Altepeter, B. Calkins, T. Gerrits, A. E. Lita, A. Miller, L. K. Shalm, Y. Zhang, S. W. Nam, N. Brunner, C. C. W. Lim, N. Gisin, and P. G. Kwiat, Detection-Loophole-Free Test of Quantum Nonlocality, and Applications, Phys. Rev. Lett. **111**, 130406 (2013).

[70] W.-Z. Liu, M.-H. Li, S. Ragy, S.-R. Zhao, B. Bai, Y. Liu, P. J. Brown, J. Zhang, R. Colbeck, J. Fan, Qiang Zhang, and Jian-Wei Pan, Device-independent randomness expansion against quantum side information, Nat. Phys. **17**, 448 (2021).

[71] D. P. Nadlinger, P. Drmota, B. C. Nichol, G. Araneda, D. Main, R. Srinivas, D. M. Lucas, C. J. Ballance, K. Ivanov, E.-Z. Tan, P. Sekatski, R. L. Urbanke, R. Renner, N. Sangouard, and J.-D. Bancal, Experimental quantum key distribution certified by Bell's theorem, Nature **607**, 682 (2022).

[72] W. Zhang, T. van Leent, K. Redeker, R. Garthoff, R. Schwonnek, F. Fertig, S. Eppelt, W. Rosenfeld, V. Scarani, C. C.-W. Lim, and Harald Weinfurter, A device-independent quantum key distribution system for distant users, Nature **607**, 687 (2022).

[73] V. Zapatero, T. van Leent, R. Arnon-Friedman, W.-Z. Liu, Q. Zhang, H. Weinfurter, and M. Curty, Advances in device-independent quantum key distribution, Npj Quantum Inf. **9**, 10 (2023).

[74] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, Loophole-free

Bell inequality violation using electron spins separated by 1.3 kilometres, Nature **526**, 682 (2015).

[75] L. K. Shalm, E. Meyer-Scott, B. G. Christensen, P. Bierhorst, M. A. Wayne, M. J. Stevens, T. Gerrits, S. Glancy, D. R. Hamel, and M. S. Allman *et al.*, Strong Loophole-Free Test of Local Realism, Phys. Rev. Lett. **115**, 250402 (2015).

[76] M. Giustina, M. A. Versteegh, S. Wengerowsky, J. Handsteiner, A. Hochrainer, K. Phelan, F. Steinlechner, J. Kofler, J.-Å. Larsson, and C. Abellán *et al.*, Significant-Loophole-Free Test of Bell's Theorem with Entangled Photons, Phys. Rev. Lett. **115**, 250401 (2015).

[77] W. Rosenfeld, D. Burchardt, R. Garthoff, K. Redeker, N. Ortegel, M. Rau, and H. Weinfurter, Event-Ready Bell Test Using Entangled Atoms Simultaneously Closing Detection and Locality Loopholes, Phys. Rev. Lett. **119**, 010402 (2017).

[78] N. Brunner, N. Gisin, V. Scarani, and C. Simon, Detection Loophole in Asymmetric Bell Experiments, Phys. Rev. Lett. **98**, 220403 (2007).

[79] A. Cabello and J.-Å. Larsson, Minimum Detection Efficiency for a Loophole-Free Atom-Photon Bell Experiment, Phys. Rev. Lett. **98**, 220402 (2007).

[80] N. Sangouard, J.-D. Bancal, N. Gisin, W. Rosenfeld, P. Sekatski, M. Weber, and H. Weinfurter, Loophole-free Bell test with one atom and less than one photon on average, Phys. Rev. A **84**, 052122 (2011).

[81] C. Teo, M. Araújo, M. T. Quintino, J. Minář, D. Cavalcanti, V. Scarani, M. Terra Cunha, and M. França Santos, Realistic loophole-free Bell test with atom–photon entanglement, Nat. Commun. **4**, 2104 (2013).

[82] T. van Leent, M. Bock, F. Fertig, R. Garthoff, S. Eppelt, Y. Zhou, P. Malik, M. Seubert, T. Bauer, W. Rosenfeld, Wei Zhang, Christoph Becher, and Harald Weinfurter, Entangling single atoms over 33 km telecom fibre, Nature **607**, 69 (2022).

[83] D. D. Awschalom, R. Hanson, J. Wrachtrup, and B. B. Zhou, Quantum technologies with optically interfaced solid-state spins, Nat. Photonics **12**, 516 (2018).

[84] P. C. Maurer, G. Kucsko, C. Latta, L. Jiang, N. Y. Yao, S. D. Bennett, F. Pastawski, D. Hunger, N. Chisholm, M. Markham, D. J. Twitchen, J. I. Cirac, and M. D. Lukin, Room-temperature quantum bit memory exceeding one second, Science **336**, 1283 (2012).

[85] K. Saeedi, S. Simmons, J. Z. Salvail, P. Dluhy, H. Riemann, N. V. Abrosimov, P. Becker, H.-J. Pohl, J. J. Morton, and M. L. Thewalt, Room-temperature quantum bit storage exceeding 39 minutes using ionized donors in silicon-28, Science **342**, 830 (2013).

[86] M. Zhong, M. P. Hedges, R. L. Ahlefeldt, J. G. Bartholomew, S. E. Beavan, S. M. Wittig, J. J. Longdell, and M. J. Sellars, Optically addressable nuclear spins in a solid with a six-hour coherence time, Nature **517**, 177 (2015).

[87] P. Wang, C.-Y. Luan, M. Qiao, M. Um, J. Zhang, Y. Wang, X. Yuan, M. Gu, J. Zhang, and K. Kim, Single ion qubit with estimated coherence time exceeding one hour, Nat. Commun. **12**, 233 (2021).

[88] N. P. de Leon, K. M. Itoh, D. Kim, K. K. Mehta, T. E. Northup, H. Paik, B. Palmer, N. Samarth, S. Sangtawesin, and D. W. Steuerman, Materials challenges and opportunities for quantum computing hardware, Science **372**, eabb2823 (2021).

[89] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).

[90] H.-S. Zhong *et al.*, Quantum computational advantage using photons, Science **370**, 1460 (2020).

[91] X. Gao, M. Kalinowski, C.-N. Chou, M. D. Lukin, B. Barak, and S. Choi, Limitations of linear cross-entropy as a measure for quantum advantage, PRX Quantum **5**, 010334 (2024).

[92] K.-M. Chung, Y. Lee, H.-H. Lin, and X. Wu, in *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)* (Springer Cham, 2022), p. 707.

[93] D. Hangleiter and J. Eisert, Computational advantage of quantum random sampling, Rev. Mod. Phys. **95**, 035001 (2023).

[94] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Practical quantum advantage in quantum simulation, Nature **607**, 667 (2022).

[95] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletic, and M. D. Lukin, Probing topological spin liquids on a programmable quantum simulator, Science **374**, 1242 (2021).

[96] J. Klassen, M. Marvian, S. Piddock, M. Ioannou, I. Hen, and B. M. Terhal, Hardness and ease of curing the sign problem for two-local qubit hamiltonians, SIAM J. Comput. **49**, 1332 (2020).

[97] D. Hangleiter, I. Roth, D. Nagaj, and J. Eisert, Easing the Monte Carlo sign problem, Sci. Adv. **6**, eabb8341 (2020).

[98] D. J. Bernstein, "Trapdoor simulation of quantum algorithms." Slides presented at NIST Workshop on Cybersecurity in a Post-Quantum World, Gaithersburg, MD, April 2–3, 2015, https://csrc.nist.gov/csrc/media/events/workshop-on-cybersecurity-in-a-post-quantum-world/documents/presentations/session10-bernstein-dan.pdf (2015).

[99] S. Wehner, D. Elkouss, and R. Hanson, Quantum internet: A vision for the road ahead, Science **362**, eaam9288 (2018).

[100] H. Buhrman, R. Cleve, J. Watrous, and R. De Wolf, Quantum Fingerprinting, Phys. Rev. Lett. **87**, 167902 (2001).

[101] G. Alagic, T. Gagliardoni, and C. Majenz, Can you sign a quantum state?, Quantum **5**, 603 (2021).