A quantum leap in security

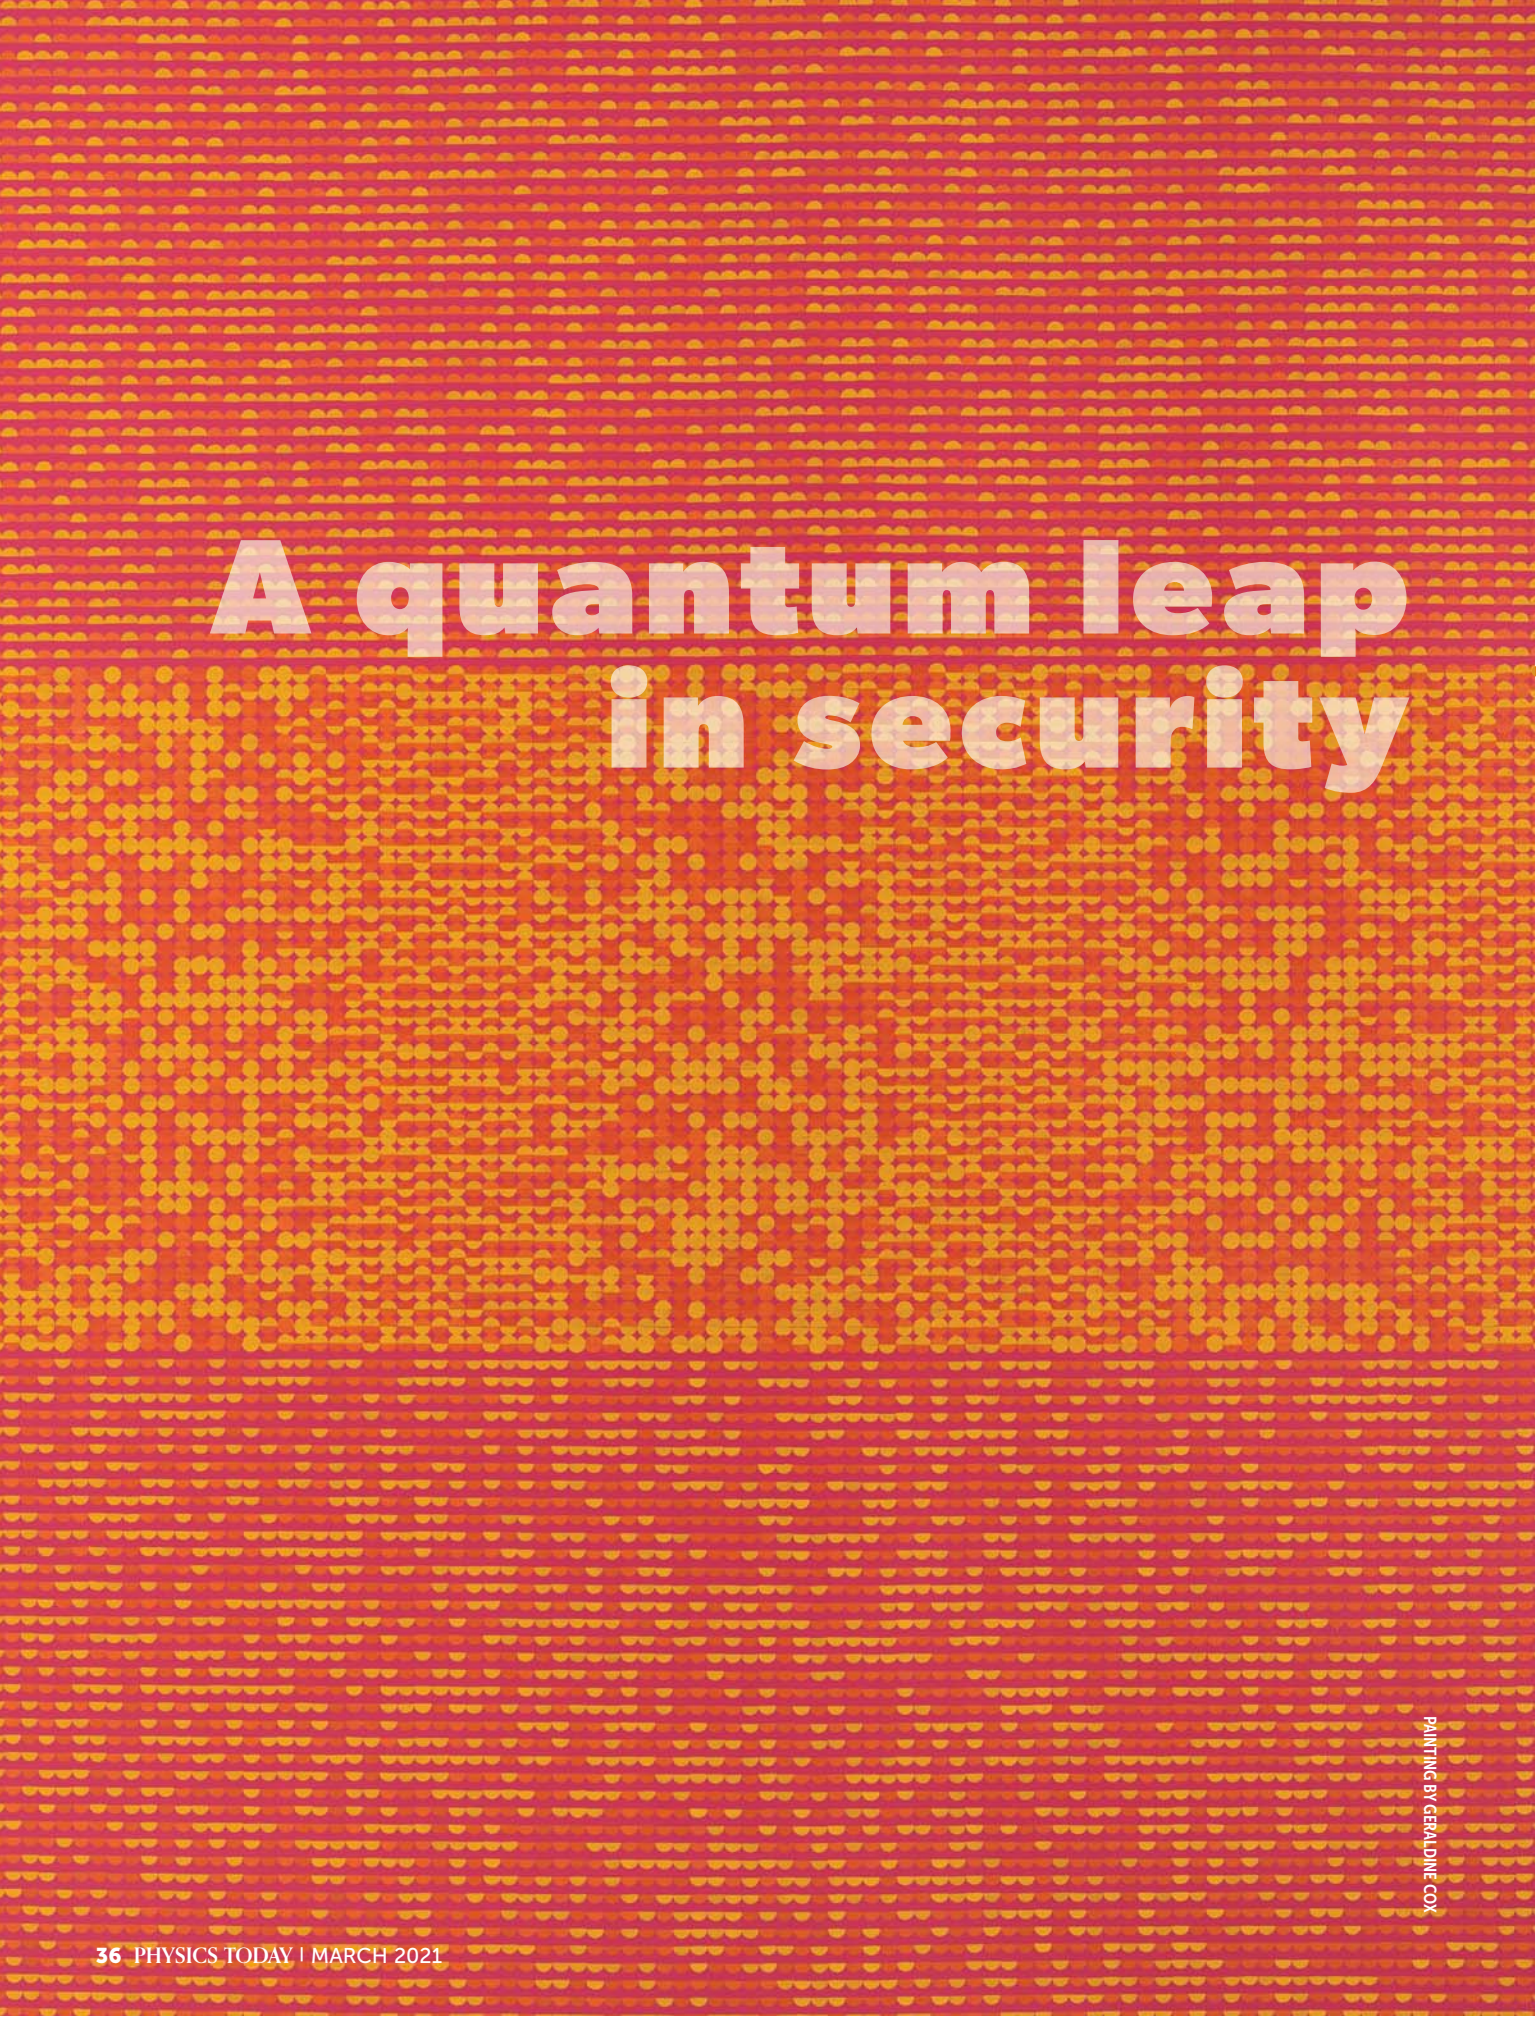

PAINTING BY GERALDINE COX

Marcos Curty is an associate professor of telecommunication engineering at the University of Vigo in Spain. **Koji Azuma** is a distinguished researcher at NTT Basic Research Laboratories in Atsugi, Japan. **Hoi-Kwong Lo** is a professor of physics at the University of Toronto and is the director of physics and astronomy at the University of Hong Kong.



Marcos Curty, Koji Azuma, and Hoi-Kwong Lo

One-photon and two-photon interferences have recently led researchers to develop new classes of quantum cryptographic protocols.

We all send sensitive data such as credit card information over the internet daily. Internet security currently relies on several computational assumptions. For example, the security of a well-known public-key encryption scheme—the so-called RSA cryptosystem—hinges on the belief that no efficient algorithm for performing prime factorization of large integers will appear in the next decade on conventional computers. But a quantum computer could efficiently factor large integers and thus break the most widely used public-key encryption schemes, including the RSA and elliptic curve cryptosystems.[1] Put simply, when a fully functioning quantum computer is built, much of conventional cryptography will fall apart.

Motivated by that eventuality—and by the many potential future applications of quantum computers in biomedicine, chemistry, artificial intelligence, and other fields—researchers have recently made tremendous progress toward constructing a large-scale, universal quantum computer. (To learn about how quantum hardware is becoming increasingly accessible, see the article by Harrison Ball, Michael Biercuk, and Michael Hush on page 28 of this issue.) Technology giants Alibaba, Google, IBM, and Microsoft are in the race. In 2019 Google claimed to have achieved the first experimental demonstration of quantum supremacy—a quantum computer capable of solving a problem unfeasible for a conventional computer.[2] Rapid developments have spurred the US National Security Agency, which spearheads code-making and code-breaking in the country, to start planning for a transition to quantum-safe cryptosystems over the next decade or so.

There are two main approaches to quantum-safe cryptography. The first one, post-quantum cryptography, relies on conventional public-key cryptosystems that experts believe are resistant to existing quantum algorithms. Its security against future advances in classical or quantum algorithms, however, has yet to be established. The second approach, quantum key distribution (QKD),[3] relies on the quantum no-cloning theorem, which states that any attempt to copy an unknown quantum state, or even try to obtain information about it, disturbs the original state. With that theorem, QKD securely distributes a common string of secret bits, called a cryptographic key, between two distant parties, typically named Alice and Bob.

The security of QKD holds even if a potential eavesdropper—say, Eve—has computational capabilities that reach the limit allowed by quantum mechanics. The security is achieved by sending nonorthogonal quantum signals through an open channel, such as an optical fiber or a free-space link. Any eavesdropping attempt to access the transmitted information can be caught because it introduces detectable errors.

If the established secret key is combined with a one-time-pad cryptosystem, Alice and Bob can communicate in absolute privacy through an untrusted channel (see the article by Daniel

Gottesman and Hoi-Kwong Lo, PHYSICS TODAY, November 2000, page 22). In the one-time-pad cryptosystem, to create the ciphertext that she sends to Bob, Alice applies bitwise XOR operations between her message and the key. The XOR operation outputs a bit value of 1 only if the two input bits differ from each other. On the receiving side, Bob decrypts the ciphertext by using bitwise XOR operations with his copy of the key. The length of the key needs to coincide with that of the message, and the key must be discarded once used.

To generate a secret key using QKD, Alice and Bob must first distribute a (possibly virtual) bipartite quantum-entangled state through a quantum channel. A bipartite quantum state is entangled precisely if it exhibits stronger than classical correlations (see the article by Reinhold Bertlmann, PHYSICS TODAY, July 2015, page 40). For example, suppose that Alice's and Bob's systems are prepared in the entangled state $|\psi^-\rangle_{AB} = 1/\sqrt{2}\,(|0\rangle_A\,|1\rangle_B - |1\rangle_A\,|0\rangle_B)$, called a Bell state. Here, $|0\rangle$ and $|1\rangle$ form an orthonormal basis termed *Z*. If Alice and Bob measure their individual systems in the *Z* basis, the measurements will produce opposite results: If Alice obtains $|0\rangle$, Bob generates $|1\rangle$.

That property holds for any common measurement basis selected by Alice and Bob. Most importantly, their results are totally random and unpredictable for Eve. If Alice and Bob associate the bit value 0 to the result $|0\rangle$ and the bit value 1 to the result $|1\rangle$, they obtain a secret key, and Bob needs to flip only his bit values to match those of Alice. Therefore, if they share many Bell states, they can perform secure communication by means of the one-time-pad cryptosystem.

In practice, channel loss, channel noise, device imperfections, and a possible attack by Eve might prevent Alice and Bob from sharing perfect Bell states. Still, quantum mechanics allows them to verify if the shared states are sufficiently close to Bell states. If they are, Alice and Bob can distill a smaller fraction of perfect Bell states from the original states using local operations and public, classical communication. That fraction determines the length of the secret key that Alice and Bob can extract from their shared systems.

Progress and challenges

Researchers have developed high-speed QKD systems with repetition rates of up to 10 GHz; implemented long-distance, fiber-based, point-to-point QKD links as far as 421 km apart; and enabled the multiplexing of quantum and classical signals in the same fiber, which is necessary for QKD to be compatible with conventional optical communication systems. QKD networks are now being deployed worldwide for secure communication in metropolitan and suburban areas.

Quantum-repeater technology would enable entanglement distribution over arbitrarily long distances, but it has yet to be developed. Currently, QKD networks typically rely on a trusted-node architecture to overcome the distance limitation imposed by channel loss.[4] For that setup, QKD only protects the communication between adjacent nodes in the network, and a copy of the key is available at all trusted nodes. In China, a 2000 km QKD backbone with about 30 trusted nodes connects Beijing and Shanghai, and a ground-to-satellite QKD network has recently enabled a secure video conference between Beijing and Vienna, 7600 km apart. Likewise, Europe and the US are developing blueprints for building continental-scale QKD networks this decade.

Despite such tremendous achievements, some fundamental challenges remain. The most pressing one is to guarantee the security of real-life QKD implementations. Because of device imperfections, real devices could, for example, leak electro-

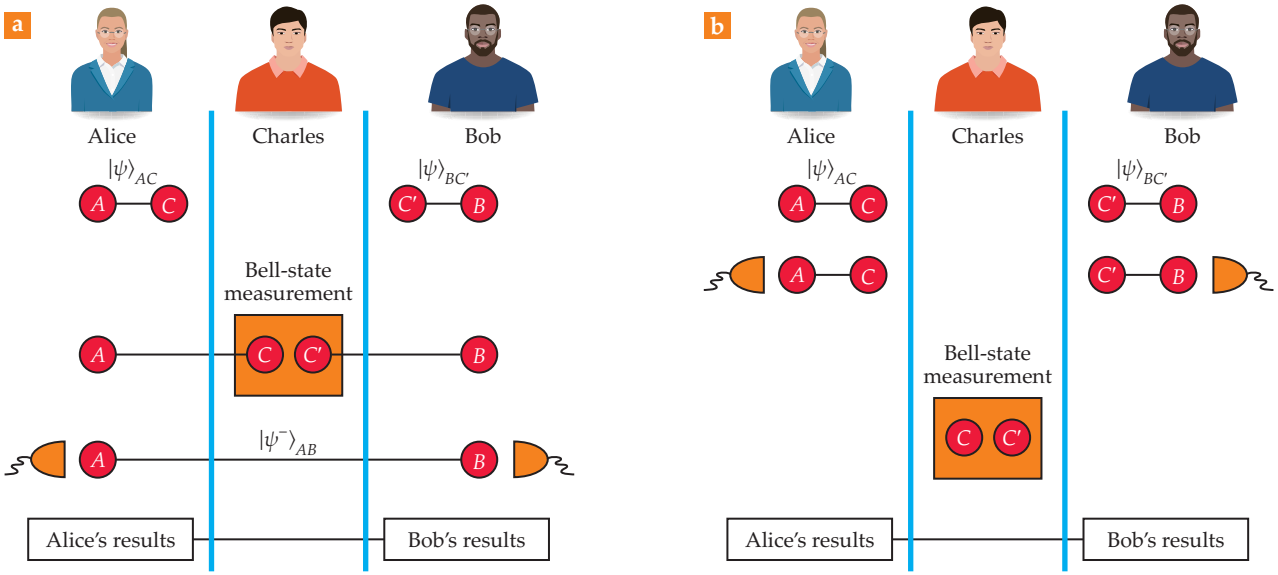

FIGURE 1. QUANTUM KEYS can be distributed via entanglement swapping. **(a)** Two parties, Alice and Bob, who wish to communicate information, prepare entangled states $|\psi\rangle_{AC}$ and $|\psi\rangle_{BC'}$ and send the particles *C* and *C*′, respectively, to a third person, Charles. If Charles's Bell-state measurement is successful, particles *A* and *B* become entangled in a Bell state, such as $|\psi^-\rangle_{AB}$. Alice and Bob can verify the entanglement by measuring their particles *A* and *B* in the *Z* and the *X* bases at random and then comparing their results. **(b)** Because Alice's and Bob's measurements commute with those of Charles, they could measure the particles *A* and *B* before they send him *C* and *C*′. That approach is equivalent to a prepare-and-measure scheme in which they send Charles the states that would result from such a process without actually preparing entangled states. (Adapted from F. Xu et al., *IEEE J. Sel. Top. Quant. Electron.* **21**, 6601111, 2015.)

magnetic or acoustic radiation, or behave differently from what is typically assumed in security analyses. The deviations could open security loopholes, or so-called side channels, which Eve might exploit. Single-photon detectors at a receiver are particularly sensitive to quantum hacking attacks. Researchers have demonstrated that by injecting strong light into the receiver, Eve could control which of Bob's detectors observes a signal each time and thus obtain the secret key.[5]

One possible approach to bridge the gap between theory and practice is QKD that is device-independent (DI),[6] though no experimental implementation has been realized so far. The solution uses a Bell inequality to verify if Alice and Bob share an entangled state, and it thus does not require them to characterize the internal functioning of the apparatuses. Despite its theoretical beauty, DI QKD is impractical with current technology: It demands a nearly perfect single-photon detection efficiency and yet would provide only a low key rate—the number of secret bits obtained per transmitted signal—at short distances of about 40 km. In addition, ensuring that the measurement apparatuses do not leak any information to Eve might be challenging for uncharacterized devices. For example, certain single-photon detectors emit backflash light that reveals which detector observes a signal each time.

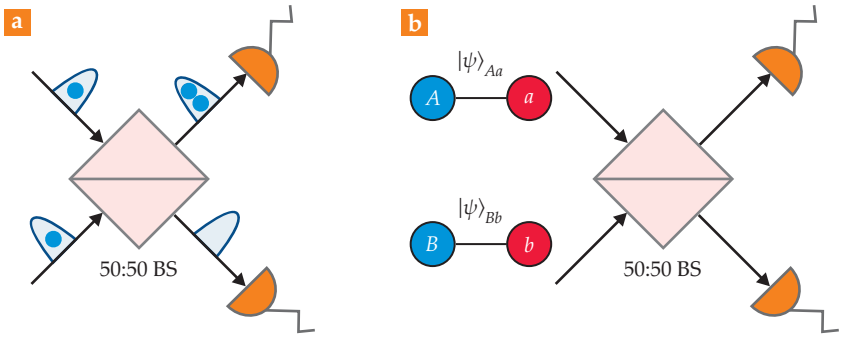

FIGURE 2. GENERATING INTERFERENCE with a two-photon and one-photon approach. **(a)** Due to the Hong-Ou-Mandel effect, if two indistinguishable photons (left) enter through different input ports of a 50:50 beamsplitter (BS), both photons exit the BS through the same randomly chosen output port (orange, right). If the photons do not overlap perfectly in time, the probability that they exit the BS through different ports increases with the time delay between them, up to a value of 0.5. **(b)** Alice prepares the entangled state $|\psi\rangle_{Aa} = \sqrt{p}|0\rangle_A\,|0\rangle_a + \sqrt{1-p}\,|1\rangle_A|1\rangle_a$ where p is an arbitrary nonzero probability, $|0\rangle_A$ and $|1\rangle_A$ are an orthonormal basis, and $|0\rangle_a$ and $|1\rangle_a$ represent a vacuum and a one-photon state, respectively. Bob prepares an analogous state $|\psi\rangle_{Bb}$. If the photonic systems a and b interfere at the BS and the detectors at its output ports observe precisely one photon, then the particles A and B become a Bell state because of one-photon interference. (Figure by Marcos Curty, Koji Azuma, and Hoi-Kwong Lo.)

The performance of QKD in terms of key rate versus distance still needs improvement. In point-to-point QKD configurations, the key rate scales at most linearly with the transmittance of the quantum channel, which is the probability that a one-photon pulse emitted by Alice reaches Bob.[4] For typical optical-fiber channels, the transmittance decreases exponentially with the distance and thus so does the key rate. Whereas quantum repeaters are the ideal solution, researchers can increase secret-key rates via multiplexing techniques.

Network distribution

To distribute an entangled state and protect the QKD setups from side-channel attacks, researchers have developed an alternative solution called entanglement swapping. It relies on a third party called Charles, who holds the measurement unit and forms a small quantum network with Alice and Bob. Each of them prepares an entangled state locally and sends one half of the entangled pair of particles to Charles and keeps the other half in their respective lab. At the receiving side, Charles detects the arriving signals with a measurement that projects them into a Bell state. Remarkably, if Charles's measurement is successful in the ideal noiseless scenario, the local particles at Alice's and Bob's labs become entangled in a Bell state even though they have not interacted with each other. Figure 1a represents the process.

Alice and Bob can verify that they actually share Bell states, or states sufficiently close to them, independently of the method used to distribute the states. To complete the verification, they measure their local systems in two conjugate bases Z and X and then compare a randomly chosen subset of the results. Here, the X basis is defined by two orthonormal states $|+\rangle = 1/(\sqrt{2})(|0\rangle + |1\rangle)$ and $|-\rangle = 1/(\sqrt{2})(|0\rangle - |1\rangle)$. If they share states sufficiently close to Bell states, they proceed with the key generation phase; otherwise, they abort. To generate a key, Alice and Bob then process their data by performing error-correction and privacy-amplification steps, the second of which removes any information that Eve could have learned about the data. Privacy amplification requires Alice and Bob to apply a particular hash function to the corrected data, which maps a bit string to a shorter bit string. The result is an almost perfectly secure key.

Alice and Bob can then verify that Charles behaved honestly by confirming that they share entangled states. Because all of the detectors are within Charles's station, Alice and Bob are protected from all possible side-channel attacks that target the measurement unit.

From a practical point of view, the setup in figure 1a can be simplified further as illustrated in figure 1b. Because Alice's and Bob's local measurements commute with Charles's, the measurement order is irrelevant. Alice and Bob could each measure one half of the entangled pair before they send the other half to Charles. More importantly, the procedure is equivalent to a so-called prepare-and-measure scenario in which Alice and Bob directly prepare the states of the signals that are sent to Charles without first generating entangled states. In practice, the arrangement means that Alice and Bob do not need to distribute real entanglement between them. They merely need to share virtual entanglement or, to be more precise, to confirm that they would have shared real entanglement if they had prepared and sent real entangled states.

The essential ingredient that enables entanglement swapping is the Bell-state measurement performed at Charles's station. Remarkably, researchers can make such a measurement by using a simple interferometric setup with standard, linear

optical components based on two-photon or one-photon interference. Both effects, each resulting in a different QKD protocol with its own merits, are illustrated in figure 2.

The scenarios that have been considered thus far assume that Alice and Bob can prepare perfect entangled states. In a prepare-and-measure setup, that assumption corresponds to generating photon-number states, which have a well-defined number of photons. But those states are challenging to prepare with the experimental capabilities currently available. Instead, researchers typically prefer to implement QKD using attenuated laser sources that emit weak coherent pulses (WCPs). Although Alice and Bob could still run QKD protocols with that experimental setup, they would need to estimate certain quantities related to the photon-number states. Fortunately, they can estimate those quantities with the decoy-state method.[7] It requires Alice and Bob to randomly vary the photon-number statistics of the respective signals they each generate. If Alice and Bob generate phase-randomized WCPs, the decoy-state method requires them to simply change the laser's intensity setting.

Measurement-device-independent QKD

An important, recent research direction builds QKD networks with untrusted relays using QKD that is measurement-device-independent (MDI).[8] It's currently the most popular and effective solution to counter quantum hacking because of its practicality and high key-generation rate at long distances. Indeed, numerous experimental demonstrations of MDI QKD have been reported in recent years that have achieved 1 Mb/s secret-key rates[9] and transmission distances of 404 km with telecommunication fibers.[10] The successes would be enough, for example, to encrypt a high-quality video call with the one-time-pad cryptosystem or to distribute secret keys between the Canadian cities of Toronto and Ottawa.

MDI QKD builds on the entanglement-swapping protocol that uses a Bell-state measurement based on two-photon interference and is implemented in a prepare-and-measure fashion using WCPs and decoy states. Secret bits are distilled from the one-photon contributions emitted by Alice and Bob and successfully detected by Charles, who could be the QKD network provider.

A schematic diagram of MDI QKD is shown in figure 3a. Alice and Bob each send Charles phase-randomized WCPs independently prepared in one of the four polarization states employed in 1984 by Charles Bennett and Gilles Brassard.[3] The previously introduced orthonormal states $|0\rangle$ and $|1\rangle$ may now be defined as the horizontally polarized one-photon state $|H\rangle$ and the vertically polarized one-photon state $|V\rangle$. Generating the QKD signals is then equivalent to having Alice and Bob each prepare the Bell state $|\psi^-\rangle = 1/\sqrt{2}\,(|H\rangle|V\rangle - |V\rangle|H\rangle)$ and then measure the first particle in either the Z or X basis selected at random.

Charles is supposed to measure the incoming signals with a Bell-state measurement and then announce his results. The setup exploits the Hong-Ou-Mandel effect to identify two of the four Bell states, which is enough to achieve secure QKD. Depending on the Bell states announced and the polarization bases used, Bob might need to bit-flip part of his polarization

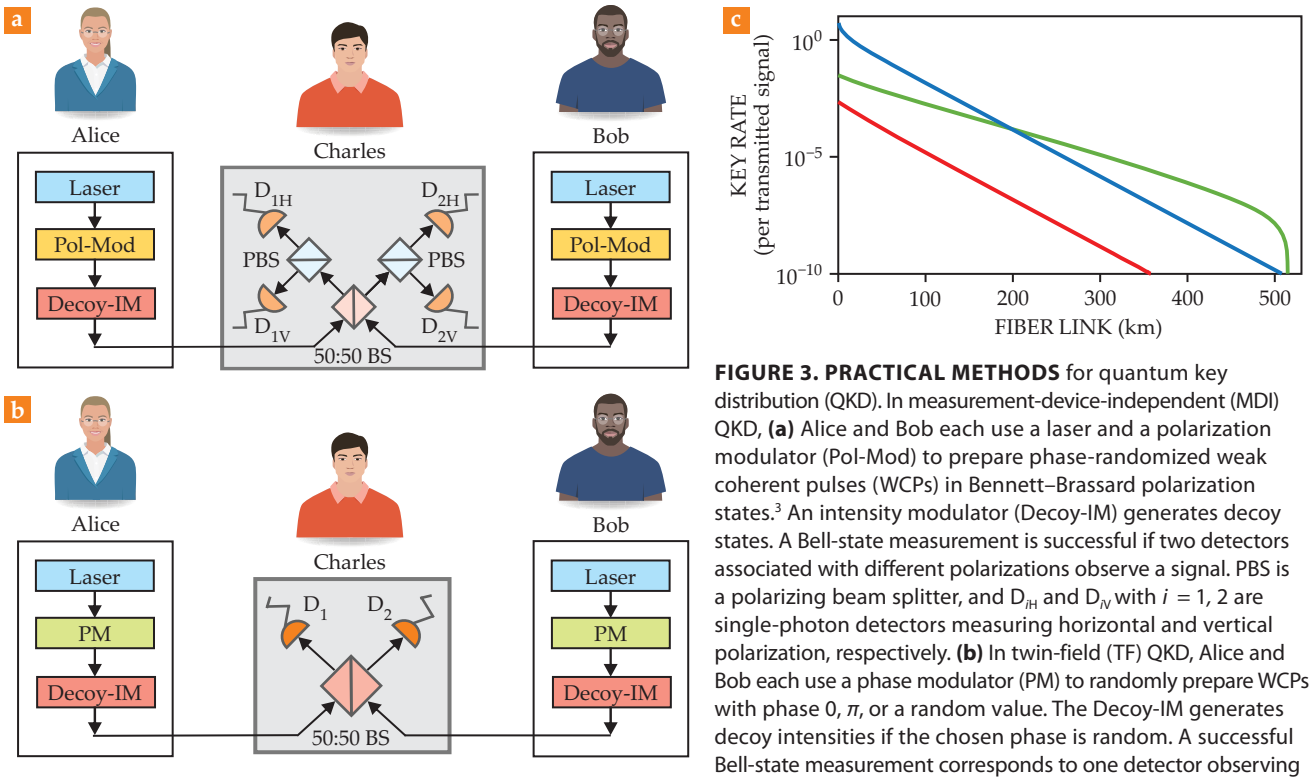

FIGURE 3. PRACTICAL METHODS for quantum key distribution (QKD). In measurement-device-independent (MDI) QKD, **(a)** Alice and Bob each use a laser and a polarization modulator (Pol-Mod) to prepare phase-randomized weak coherent pulses (WCPs) in Bennett–Brassard polarization states.[3] An intensity modulator (Decoy-IM) generates decoy states. A Bell-state measurement is successful if two detectors associated with different polarizations observe a signal. PBS is a polarizing beam splitter, and D_{iH} and D_{iV} with $i = 1, 2$ are single-photon detectors measuring horizontal and vertical polarization, respectively. **(b)** In twin-field (TF) QKD, Alice and Bob each use a phase modulator (PM) to randomly prepare WCPs with phase 0, π, or a random value. The Decoy-IM generates decoy intensities if the chosen phase is random. A successful Bell-state measurement corresponds to one detector observing a signal. **(c)** The asymptotic rate at which secret keys are generated for MDI QKD (red line) and TF QKD (green line) depends on the distance between Alice and Bob. The blue line is the private capacity of point-to-point QKD.[4] (Panels a and b by Donna Padian; panel c by Marcos Curty, Koji Azuma, and Hoi-Kwong Lo.)

data to match Alice's. The raw key is formed by the polarization data in which both Alice and Bob employ the Z basis, and Charles declares a successful result. Then they use the decoy-state method to estimate the number of bits in the raw key that have been obtained from their one-photon emissions. Alice and Bob use the X basis events and the decoy-state method to estimate how much information Eve could have learned about the raw key. A secure key is then made by applying error correction and privacy amplification to the raw key.

When compared with conventional QKD schemes that suffer from detection side-channels, the key advantage of MDI QKD is that Charles need not be trustworthy. He can only learn if Alice's and Bob's raw key bits are the same or different but not their particular values. The secret-key rate still scales linearly with the quantum channel transmittance because MDI QKD requires that two photons—one from Alice and the other from Bob—reach Charles. To overcome that limit, one can furnish MDI QKD with quantum memories[11] or quantum nondemolition measurements[12] or use another approach known as twin-field (TF) QKD,[13] shown schematically in figure 3b.

An alternative approach

The elegant idea of TF QKD replaces the entanglement-swapping operation based on two-photon interference with one based on one-photon interference.[9] The method, if successful, projects the incoming states into a Bell state $|\psi^{\pm}\rangle = 1/\sqrt{2}(|0\rangle|1\rangle \pm |1\rangle|0\rangle)$, whose orthonormal states $|0\rangle$ and $|1\rangle$ refer to the vacuum and a one-photon state, respectively. With that projection, only one photon from Alice or Bob sent to Charles is sufficient to generate a secret key. TF QKD doubles the transmission distance compared with MDI QKD and is robust to any possible side-channel attack because Charles can be untrusted. Figure 3c compares the key rates of the two QKD protocols.

Several recently introduced variants of TF QKD offer security against general attacks.[14,15] Most importantly, the ideal setup[15] can be well approximated with a prepare-and-measure scheme in which Alice and Bob each send Charles WCPs whose phase is randomly and independently selected as 0, π, or a random value. A phase value of 0 encodes a bit value of 0; π encodes a bit value of 1; and a random phase value corresponds to a decoy state. If a random phase is selected, the pulse intensity is also randomly chosen, usually from among three settings. With decoy states, the privacy amplification that needs to be applied to the raw key can be tightly estimated. The raw key is obtained from those instances that encode a bit value and result in a detection at only one of Charles's detectors.

The main experimental challenge of TF QKD is maintaining the phase stability between Alice's and Bob's signals, which is not required in MDI QKD. That demand means that TF QKD needs an auto-compensating technique, such as a Sagnac loop, or phase locking of the remote laser pulses. But despite the experimental difficulties, various research groups have already performed proof-of-principle demonstrations[16] and have achieved transmission distances longer than 500 km.[17]

Closing the gap

Quantum interference enables a family of novel protocols that offer unprecedented levels of security and performance for QKD. The protocols are particularly suited for an untrusted network setting with multiple users. Each user holds a low-cost, compact, chip-based QKD transmitter, and they all share the measurement unit that contains the single-photon detectors.

However, in real-life network settings, the symmetric scenario, in which the channel loss between Alice and Charles is the same as or similar to that between Bob and Charles, is not always true. Some researchers introduced efficient variants of MDI QKD and TF QKD for asymmetric configurations that allow Alice and Bob to use different intensity settings for their signals.[18] The protocols could prove useful in a general quantum network with vastly different channel losses. In such a network, users are dynamically added or deleted at any time without compromising network performance.

A fundamental question that remains unanswered is how to protect QKD transmitter hardware against quantum hacking. Protection will require the development of security proofs that can handle device imperfections in the transmitters and hardware countermeasures that prevent the manipulation of devices. Fortunately, those tasks are, in general, much simpler than protecting the measurement unit. Alice and Bob could use optical isolation, spectral filters, and monitor detectors to physically protect their transmitters from Eve. In addition, security proofs that include most transmitters' imperfections have been developed in recent years. When combined with the setups introduced in this article, the security proofs can close the gap between QKD theory and practice.

We are grateful for financial support from the European Union's Horizon 2020 research and innovation programme under a QCALL Marie Sklodowska-Curie grant, Spain's Ministry of Economy and Competitiveness, the European Regional Development Fund through grant TEC2017-88243-R, the PRESTO program of the Japan Science and Technology Agency and grant JP-MJPR1861, the Natural Sciences and Engineering Research Council of Canada, the Canada Foundation for Innovation, the Ontario Research Fund, Huawei Technologies Canada, Mitacs, the US Office of Naval Research, the Royal Bank of Canada, and the University of Hong Kong startup grant.

REFERENCES

1. P. W. Shor, in *Proceedings 35th Annual Symposium on Foundations of Computer Science*, IEEE (1994), p. 124.
2. F. Arute et al., *Nature* **574**, 505 (2019).
3. C. H. Bennett, G. Brassard, *Proceedings of the International Conference on Computers, Systems and Signal Processing*, IEEE (1984), p. 175; A. K. Ekert, *Phys. Rev. Lett.* **67**, 661 (1991).
4. S. Pirandola et al., *Nat. Commun.* **8**, 15043 (2017).
5. L. Lydersen et al., *Nat. Photonics* **4**, 686 (2010).
6. A. Acín et al., *Phys. Rev. Lett.* **98**, 230501 (2007).
7. W.-Y. Hwang, *Phys. Rev. Lett.* **91**, 057901 (2003); H.-K. Lo, X. Ma, K. Chen, *Phys. Rev. Lett.* **94**, 230504 (2005); X.-B. Wang, *Phys. Rev. Lett.* **94**, 230503 (2005).
8. H.-K. Lo, M. Curty, B. Qi, *Phys. Rev. Lett.* **108**, 130503 (2012).
9. M. Lucamarini et al., *Nature* **557**, 400 (2018).
10. L. C. Comandar et al., *Nat. Photonics* **10**, 312 (2016).
11. H.-L. Yin et al., *Phys. Rev. Lett.* **117**, 190501 (2016).
12. C. Panayi et al., *New J. Phys.* **16**, 043005 (2014).
13. K. Azuma, K. Tamaki, W. J. Munro, *Nat. Commun.* **6**, 10171 (2015).
14. X.-B. Wang, Z.-W. Yu, X.-L. Hu, *Phys. Rev. A* **98**, 062323 (2018).
15. M. Curty, K. Azuma, H.-K. Lo, *npj Quant. Inf.* **5**, 64 (2019).
16. M. Minder et al., *Nat. Photonics* **13**, 334 (2019).
17. J.-P. Chen et al., *Phys. Rev. Lett.* **124**, 070501 (2020).
18. W. Wang, F. Xu, H.-K. Lo, *Phys. Rev. X* **9**, 041012 (2019); F. Grasselli, Á. Navarette, M. Curty, *New J. Phys.* **21**, 113032 (2019). PT