QUANTUM COMMUNICATION · SATELLITE QKD · SPACE ENGINEERING

Keys From *Orbit*

Forty years after two researchers proposed encoding secrets in single photons, spacecraft are distributing cryptographic keys between continents. The physics is settled. The engineering — pointing a metre-wide beam of single photons at a telescope from 800 kilometres up, at night, in clear weather, while moving at 7.5 kilometres a second — is where the real story lives.

Feature · Quantum Technologies · 2026 · approx. 26 minute read

By Sara Malik and Grieg Greger, PakCrypt.Org

Quantum Key Distribution | BB84 | Decoy States | Entanglement | Optical Ground Station | Trusted Node | SNSPD | Post-Quantum Cryptography

IN BRIEF

Quantum key distribution lets two parties share a secret whose security rests on the laws of physics rather than on the difficulty of a mathematical problem. Optical fibre cannot carry it more than a few hundred kilometres, because a quantum signal cannot be amplified. Satellites can — and since 2016 they have. This article traces how the idea travelled from a physicist's unpublishable manuscript to spacecraft linking Beijing and Cape Town, and then examines, honestly, the seven engineering problems that still make satellite QKD one of the hardest communication technologies humans have ever attempted to operate routinely.

1,203	~40 dB	0.12	1.07 Mbit
KM ENTANGLEMENT FROM ORBIT, 2017	TYPICAL DOWNLINK PHOTON LOSS	BITS PER SECOND, BEST ENTANGLED LINK	SECURE KEY IN ONE MICROSATELLITE PASS

1. Why Put Cryptography in Orbit

Every secure connection you make today — to a bank, a government portal, a messaging app — begins with a handshake in which two strangers who have never met agree on a secret number. That handshake is the most quietly astonishing trick in computing, and it rests on a mathematical wager: that certain problems, such as factoring a large integer or reversing an elliptic-curve multiplication, are too hard to solve in any reasonable time.

A sufficiently large quantum computer would win that wager. Peter Shor showed in 1994 that such a machine could factor integers and compute discrete logarithms efficiently, which retires essentially every public-key algorithm now in service. Nobody knows when such a machine will exist. What everyone in the field does know is that the threat has a peculiar time signature: an adversary can record encrypted traffic today and decrypt it whenever the machine arrives. Cryptographers call this *harvest now, decrypt later*. For anything that must stay secret for thirty or fifty years — diplomatic cables, medical genomics, industrial designs, state archives — the exposure is not in the future. It is already running.

There are two serious responses. The first, and by far the more important, is **post-quantum cryptography**: replace the vulnerable mathematics with new mathematics believed hard for quantum computers too. It is software, it works over the existing internet, and international standards were finalised in 2024 [FIPS2024]. For almost every purpose it is the right answer, and the rest of this article should not be read as an argument against it.

The second response abandons mathematical assumptions altogether. In **quantum key distribution** (QKD), the secret is not computed; it is *measured*. Two parties derive shared random bits from the outcomes of measurements on individual quantum particles, and the laws of quantum mechanics place a hard, quantifiable ceiling on how much an eavesdropper can have learned. The key that comes out is not "hard to break". It is *information-theoretically secure*: no future computer, however powerful, retroactively reveals it.

> **The property that makes QKD different**
>
> A key established by post-quantum cryptography in 2026 and recorded by an adversary can, in principle, be broken in 2050 if the underlying mathematics falls. A key established by QKD in 2026 cannot. To learn it, the

> eavesdropper had to be physically present, interfering with individual photons, on the night the key was made — and that interference is detectable. Cryptographers call this *everlasting secrecy*.

So why go to space for it? Because on the ground, quantum signals die. A single photon in optical fibre has perhaps a 95 % chance of surviving each kilometre at telecom wavelengths — a loss of about 0.2 decibels per kilometre. That sounds gentle until you compound it. Over 100 km, roughly one photon in a hundred survives. Over 500 km, about one in *ten billion*. Classical networks solve this with amplifiers and repeaters every few tens of kilometres. Quantum networks cannot: the **no-cloning theorem**, the same principle that makes eavesdropping detectable, forbids copying an unknown quantum state. You cannot amplify what you are not allowed to duplicate.

That leaves two options for long distances. Either you chain together *trusted relay stations* that decrypt and re-encrypt the key at each hop — which means every hut along the route knows your secret — or you go up. Above about 10 kilometres of altitude the atmosphere effectively ends, and the remaining 780 kilometres to a low-orbit satellite are vacuum, where a photon travels essentially for free. The loss on a space link is dominated not by absorption but by geometry: the beam simply spreads out. And geometric spreading scales with the *square* of distance, not exponentially.

Fibre loss is exponential in distance. Free-space loss is quadratic. Somewhere past a few hundred kilometres, the sky becomes the cheaper medium — which is the entire argument for satellite QKD in one sentence.

· ORIGINS ·

2. Fifty Years in Five Acts

The history of quantum cryptography is unusually well documented, partly because so much of it happened between a handful of people who kept in touch, and partly because its origin story is genuinely strange.

Act I: A paper nobody would publish

Around 1970, Stephen Wiesner — then a graduate student at Columbia — wrote up an idea he called *conjugate coding*. Because certain pairs of quantum properties cannot be measured simultaneously with arbitrary precision, he argued, one could make banknotes carrying quantum “watermarks” that a counterfeiter could not copy, and could encode two messages such that reading one destroyed the other. Journals rejected it. The manuscript circulated informally for over a decade before finally appearing in a computing newsletter in 1983 [Wiesner1983].

One of the few people who took it seriously was Charles Bennett, who had met Wiesner as a student. In the early 1980s Bennett, at IBM, and Gilles Brassard, at the Université de Montréal, turned the idea from a curiosity about money into a protocol for keys. Their 1984 conference paper [BB84] — four pages, in a proceedings volume that was for years difficult to obtain — described what is still the workhorse of the field. Everyone calls it **BB84**.

In 1991 Artur Ekert, then in Oxford, published an independent route to the same goal [Ekert1991]: instead of sending prepared states, distribute *entangled pairs* and let the correlations themselves certify security, using a violation of Bell's inequality as the proof that nobody has tampered. That distinction — prepare-and-

measure versus entanglement — still divides the field's architectures today, and it turns out to matter enormously for satellites.

Act II: Thirty-two centimetres

In 1989 Bennett, Brassard and colleagues built the first working demonstration. It transmitted a key across 32.5 centimetres of open air on an optical bench [Bennett1992]. The apparatus made an audible noise when it switched polarisation states, which meant that the machine, as Brassard later liked to observe, was perfectly secure against any eavesdropper who happened to be deaf. It is a useful reminder of a theme this article returns to repeatedly: the mathematics of QKD is airtight, and the hardware leaks.

Fibre demonstrations followed through the 1990s, first over kilometres, then tens of kilometres. But the exponential wall was obvious from the start, and attention turned to open air.

Act III: Mountain to mountain

Free-space experiments escalated quickly. A 1998 Los Alamos experiment ran a link of about a kilometre in daylight. In 2002 a European team pushed a key across 23.4 km between two Alpine peaks, at night [Kurtsiefer2002]. Then in 2007 came the result that convinced people space was feasible: a 144-kilometre link between the Canary Islands of La Palma and Tenerife, using an optical ground station originally built by the European Space Agency for laser communication tests [Schmitt-Manderbach2007]. Both prepare-and-measure and entanglement-based keys were demonstrated over that path [Ursin2007].

The significance was not the distance itself but the *loss*. A 144-kilometre horizontal path through dense low-altitude air is optically about as punishing as a slant path from a low-orbit satellite to a mountain-top telescope. The Canary link showed the photon budget could be closed. What remained was to put one end on a moving platform.

That happened in stages: keys exchanged with a transmitter on an aircraft in 2013, and in the same year with a payload on a hot-air balloon and a floating platform, deliberately testing what happens when the sender is in motion and the pointing loop must work in real time [Nauerth2013, WangJY2013].

Act IV: Micius

On 16 August 2016, China launched a 635-kilogram spacecraft named after the ancient philosopher Micius. It carried a 300-millimetre transmitting telescope, an entangled-photon source, and a decoy-state BB84 transmitter operating at 850 nanometres. Within a year it had produced three results that between them defined the field [Lu2022].

2017 — **Satellite-to-ground key distribution**

Decoy-state BB84 downlinks to metre-class ground telescopes produced kilobit-per-second secret keys at ranges up to 1,200 km — around 20 orders of magnitude more efficient than sending the same photons through fibre over the same distance [Liao2017].

2017 — **Entanglement across a continent**

The satellite beamed one photon of an entangled pair to each of two ground stations 1,203 km apart, and the pairs still violated a Bell inequality on arrival — the first demonstration that entanglement survives a trip from orbit through the atmosphere at that scale [Yin2017].

2018 — **Intercontinental video call**

Using the satellite as a relay, keys were shared between stations near Beijing and Vienna — about 7,600 km apart — and used to encrypt a video conference. The satellite made a separate key with each side and combined them, a mechanism explained below [Liao2018].

2020 **Entanglement-based keys, at last**

Not merely distributing entanglement but distilling a finite-key-secure cryptographic key from it, over 1,120 km. The rate: **0.12 bits per second**. That number, four orders of magnitude below the prepare-and-measure result, is the single most important figure in the entire architecture debate [Yin2020].

2022 **Miniaturisation**

A microsatellite of about 100 kg, carrying a 23-kg quantum payload, launched into low orbit — a tenth the mass of Micius, built to be manufactured in numbers rather than as a one-off science mission.

2025 **Real time, and a portable ground station**

The microsatellite demonstrated real-time key distillation on board, delivering up to **1.07 megabits** of secure key in a single pass to a ground terminal weighing under 100 kg rather than the multi-tonne observatories used previously — and relayed keys between stations in China and South Africa, about 12,900 km apart [Li2025].

Act V: Everyone else

What began as one country's science programme is now a crowded field, and the diversity of approaches is itself informative.

- **CubeSats.** A Singapore-led team flew an entangled-photon source in a nanosatellite, demonstrating in 2020 that the delicate business of generating correlated photon pairs survives launch and orbit in a shoebox-sized package [Villar2020]. Successor missions with UK partners launched in 2025 and 2026, aimed squarely at driving cost down.
- **Europe.** A dedicated QKD satellite developed by an industrial consortium under space-agency and Commission funding is scheduled for launch in the late-2026 window, intended as the space segment of a continent-wide quantum communication infrastructure. Its stated purpose is sovereignty: not depending on cryptographic infrastructure controlled elsewhere.
- **Canada.** A mission taking the opposite architectural bet — putting the *receiver* in orbit and the sources on the ground. Uplinks suffer worse turbulence, but they let you keep the complicated, upgradeable hardware where engineers can reach it.
- **Fibre integration.** Meanwhile the terrestrial backbone kept growing. A national network exceeding 10,000 km of fibre with well over a hundred backbone nodes now ingests satellite-derived keys alongside fibre-derived ones [Chen2025].

> **The pattern worth noticing**
>
> Every satellite QKD system that is *operational* today uses the same architecture: prepare-and-measure decoy-state BB84, on a single downlink, at night, in clear weather, serving a small number of large ground stations. Every entanglement payload in orbit is a demonstrator. That is not a coincidence or a failure of ambition — the rest of this article explains why the physics pushes every programme to the same answer.

3. How a Satellite Actually Makes a Key

Strip away the vocabulary and BB84 is a game about mismatched rulers.

The sender — conventionally Alice, here aboard the spacecraft — emits a photon and encodes a random bit in its polarisation. But she also picks, at random, one of two *bases*: she might use the rectilinear basis, where horizontal means 0 and vertical means 1, or the diagonal basis, where 45° means 0 and 135° means 1. The receiver, Bob, on the ground, has no idea which she chose, so he too picks a basis at random for each arriving photon.

When their choices match, Bob's measurement reproduces Alice's bit. When they differ, quantum mechanics guarantees the result is pure coin-flip noise. Afterwards, over an ordinary public radio or internet channel, they compare which *bases* they used — never the bits — and discard every event where they disagreed. Roughly half the data survives. That surviving string is the raw key.

The security comes from what an eavesdropper must do. To learn anything, Eve has to measure the photons in transit. She does not know the basis either. Half the time she picks wrong, and her measurement irreversibly scrambles the state she then forwards. Her interference shows up as errors in the small sample of key bits Alice and Bob publicly compare. Below a threshold error rate they can mathematically squeeze out a shorter key about which Eve provably knows almost nothing. Above it, they abort. There is no middle ground where Eve reads the traffic undetected.

> **Weak lasers and the decoy trick**
>
> Nobody has a perfect single-photon source that works in orbit. Real systems fire heavily attenuated laser pulses containing, on average, about half a photon. Occasionally a pulse contains two — and Eve could then split one off, keep it, and let the other through undisturbed.
>
> The fix, developed in the 2000s and now universal, is *decoy states*: Alice randomly varies the brightness of her pulses among a few preset levels. Eve cannot tell which level she is attacking, so any photon-splitting strategy distorts the statistics of the different levels differently — and the distortion is visible. Decoy-state BB84 is what every operational satellite flies.

What comes out of the sky is not yet a key. It is a stream of detection events riddled with errors, most of them from noise rather than attackers. Four classical steps follow, and they matter more to real systems than most popular accounts admit: *sifting* (discarding mismatched bases and non-detections), *parameter estimation* (measuring the error rate and bounding Eve's information), *error correction* (reconciling the two strings while leaking as little as possible), and *privacy amplification* (hashing the reconciled string down to a shorter one about which the leaked information is negligible).

There is a fifth step that satellite links cannot skip. Security proofs are cleanest when you have infinitely many photons; a satellite pass lasts perhaps five to eight minutes. **Finite-key analysis** applies statistical corrections for the short block, and those corrections can eat tens of percent of the apparent yield. When a vendor quotes a key rate without saying whether finite-key effects were included, they have quoted a number that does not exist.

The relay trick — and its price

A satellite in low orbit cannot see two distant continents at once. So how did keys get from Beijing to Vienna?

STEP 1 **Make a key with A**

On a night pass over the first station, the satellite and station A run BB84 and end up sharing a secret string K_a. The satellite stores it.

STEP 2 **Make a different key with B**

Hours later, over the second station, it independently establishes K^b.

STEP 3 **Broadcast the combination**

It computes the bitwise exclusive-or $K_a \oplus K^b$ and transmits that publicly over ordinary radio. Anyone may listen; the combination reveals nothing about either key on its own.

STEP 4 **Station B recovers A's key**

Because B knows K^b, it computes $(K_a \oplus K^b) \oplus K^b = K_a$. The two ground stations now share a key, and they never needed to see the satellite at the same moment.

This is elegant, it works, and it has one consequence that shapes every national programme on Earth: **the satellite knew both keys.** It is what the field calls a *trusted node*. Whoever built it, wrote its firmware, generated its random numbers and commands it can, in principle, recover every key it ever relayed. We will come back to this.

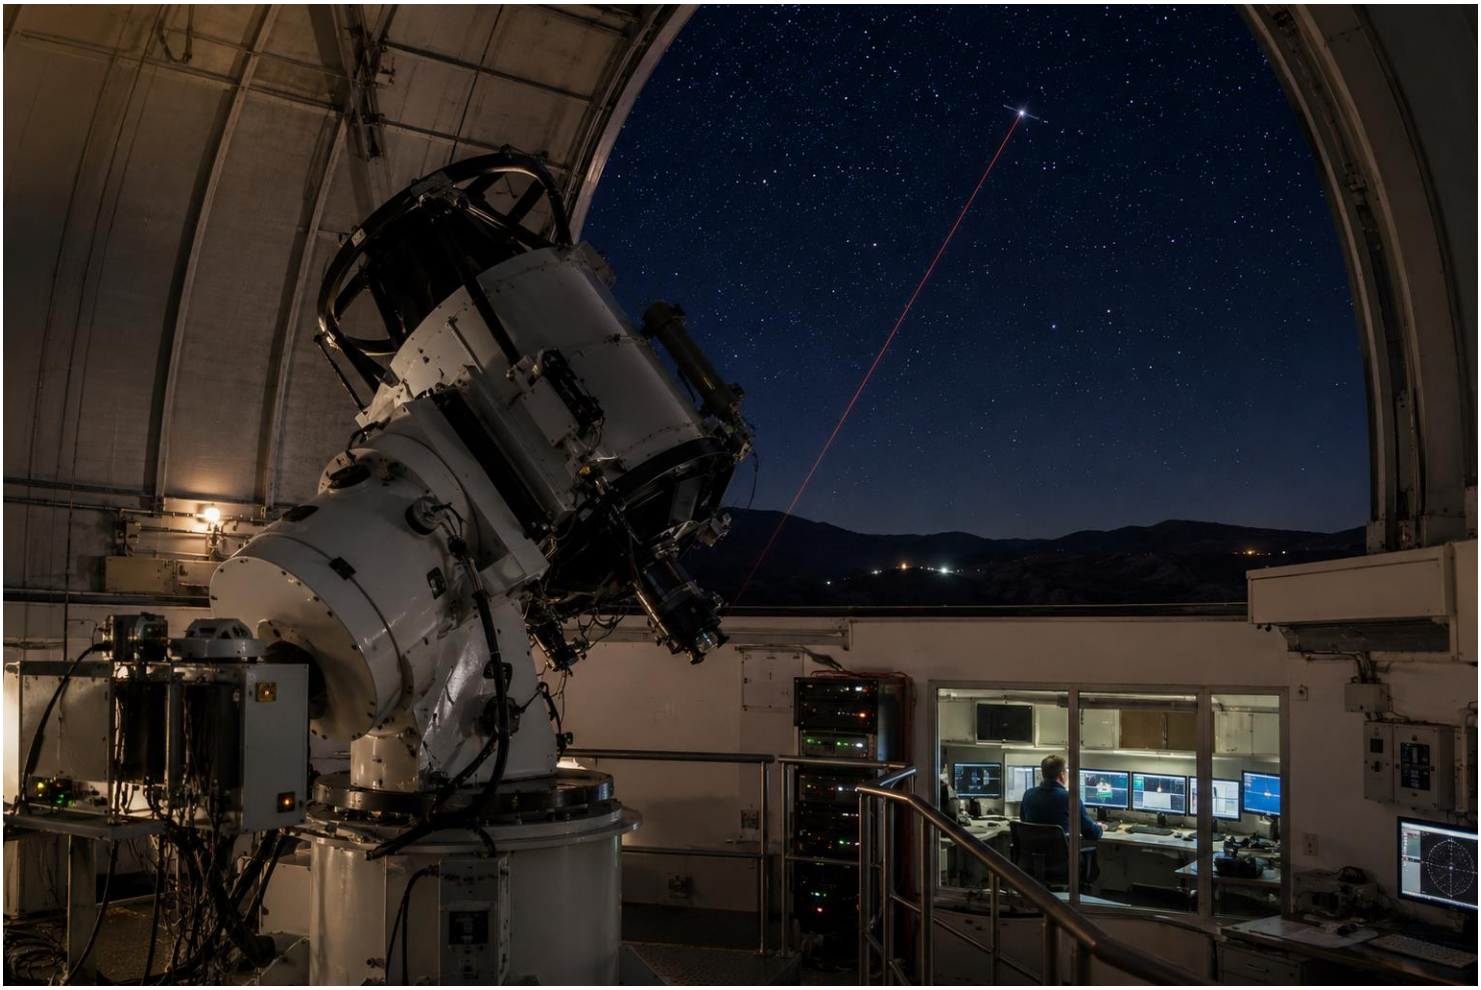

4. Nine Photons Out of a Billion

Ask an engineer what makes satellite QKD hard and you will get a different answer depending on which subsystem they lost sleep over. But every answer eventually reduces to the same root cause: you are trying to deliver individual photons across hundreds of kilometres, and you are not allowed to make more of them.

The tyranny of diffraction

Light will not travel in a perfectly straight line. Push a beam through an aperture of diameter D and it spreads with an angle set by the wavelength:

$$\theta \approx 2.44\,\lambda / D$$

DIFFRACTION-LIMITED DIVERGENCE — EQ. (1)

The numbers are unforgiving. A 200-millimetre telescope sending 850-nanometre light has a divergence of roughly ten microradians — about 0.0006 degrees, which sounds superb until you multiply by distance. At 800 kilometres, that beam has swelled into a patch on the ground some **ten to fifteen metres across**. A ground telescope of 0.8 metres captures the fraction of the patch it happens to cover:

$$\eta_{geo} = \left(D_{rx} / \theta R \right)^2$$

GEOMETRIC COLLECTION EFFICIENCY — EQ. (2)

Plug the numbers in: about 0.5 %, or a loss of roughly 23 decibels, before anything else has gone wrong. Add atmospheric absorption and scattering (2–3 dB straight up from a high, dry site, considerably more at low elevation angles), imperfect pointing, losses in the receiving optics, the narrow filters, and the detector's own inefficiency, and the total for a realistic link sits somewhere around **35 to 45 decibels**.

What 40 decibels means

It means one part in ten thousand. Fire a hundred million photons per second at the sky and you may collect ten thousand per second on the ground — from which, after sifting, error correction and privacy amplification, you might keep one or two thousand bits of actual secret key per second. A radio engineer would simply turn up the power. Here you cannot: turning up the power means putting more photons in each pulse, which is precisely what hands an eavesdropper an attack. The link budget is not a starting point for negotiation. It is the whole game.

Hitting a moving coin

A satellite in an 800-kilometre orbit travels at about 7.5 kilometres per second and crosses the sky in five to eight minutes. Seen from the ground it sweeps across at roughly half a degree per second — slow enough to look leisurely, catastrophically fast for an optical link whose beam is ten microradians wide. To stay locked, both ends must hold their aim to within a few microradians. That is roughly the angle a one-euro coin subtends at a distance of five kilometres, and it must be held continuously, on a moving target, for minutes at a time.

Nothing about a conventional satellite dish helps here. The pointing chain is closer to astronomy than to telecommunications, and it runs in stages:

- **Prediction.** Orbital ephemeris propagated forward gets both ends pointing within a tenth of a degree — good enough to be in the right part of the sky, hopeless for the link itself.
- **Beacon acquisition.** The ground station fires a bright laser beacon, deliberately spread over about half a milliradian so that it will actually land on the spacecraft despite the pointing error. The satellite sees it on a camera or quadrant detector and knows precisely where the ground station is.
- **Closing the loop.** A *fast steering mirror* — a small mirror on piezoelectric actuators, correcting a thousand times a second — takes over the fine work, holding the beam to one or two microradians while the whole spacecraft slews. The ground station closes its own loop on a beacon coming down.
- **Only then, the quantum signal.** Photons are sent only while both loops report lock, and every microradian of residual jitter shows up directly as lost key.

Engineers who have built these systems will tell you, almost without exception, that this control loop — not the exotic quantum hardware — is where programmes fail. A telescope that cannot point is an expensive mirror.

Ground stations, meanwhile, are not telecom equipment. A working optical ground station is an astronomical observatory with a cryptographic annexe: a 0.6- to 1-metre diffraction-limited telescope on a direct-drive mount, a dome, a concrete pier isolated from the surrounding ground so that a passing truck does not shake the link, adaptive or tip–tilt optics to fight atmospheric blurring, and a control room. Nobody puts one on a rooftop in a city centre, for reasons the next sections make painfully clear.

· CHALLENGE TWO ·

5. Catching, Counting and Clocking

Detecting one photon at a time

At the bottom of the link sits a device that must reliably register the arrival of a single quantum of light, and — just as importantly — must almost never claim to have seen one when it has not. Two technologies compete.

Property	Silicon avalanche photodiode (SPAD)	Superconducting nanowire (SNSPD)
Efficiency	50–65 % at 850 nm	85–95 %
False counts	A few hundred per second	Tens per second
Timing precision	350–500 picoseconds	Under 50 picoseconds
Cooling	Thermoelectric, −30 °C	Cryogenic, 1–2.5 kelvin
Practical consequence	Cheap, rugged, deployable anywhere	Best performance, needs a closed-cycle cryocooler, kilowatts of power and regular servicing

The superconducting detector is a nanometre-thin wire of niobium nitride held just below its critical current. A single absorbed photon deposits enough energy to break superconductivity in a tiny hotspot, producing a measurable voltage pulse. It is a beautiful device and its performance is not really matched by anything else.

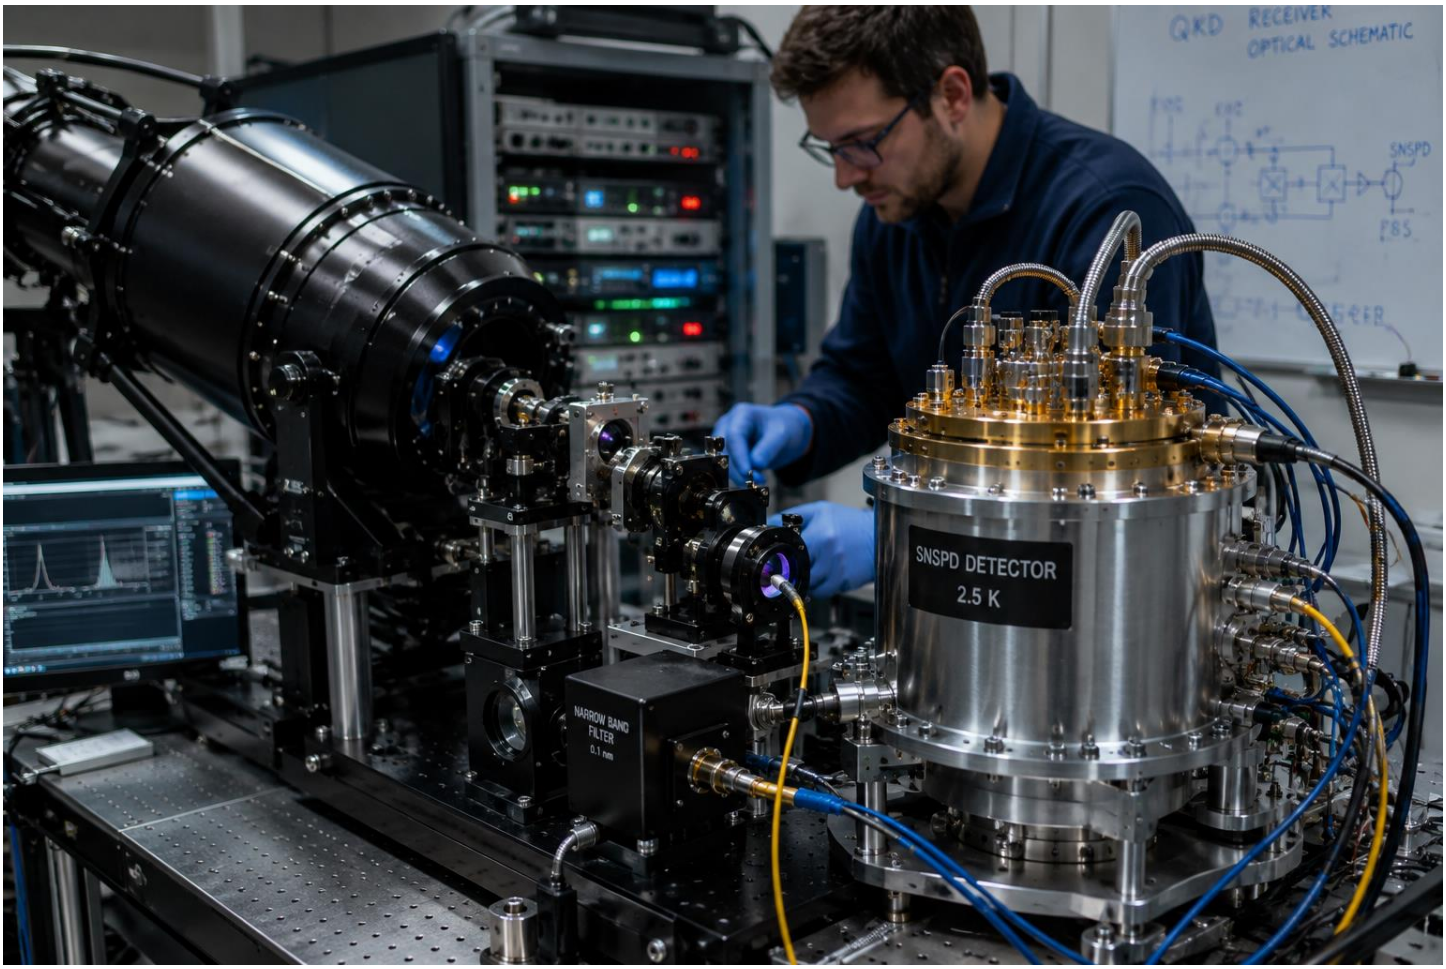

Its cost, though, is dominated not by the wire but by the refrigerator wrapped around it — which is why several programmes deliberately begin with the humbler semiconductor detector, get the rest of the system working, and add cryogenics later.

Knowing exactly when to look

Here is the trick that makes the whole thing possible at all. Background light — starlight, scattered city glow, airglow, moonlight — arrives at random times. The signal photons do not: they leave the spacecraft at precisely known instants, dictated by a clock. So the receiver simply refuses to listen except during a window about a *nanosecond* wide, timed to when the photon is expected. That single measure throws away well over 99.9 % of the background.

Making it work requires a timing chain most people never think about: an atomic-grade frequency reference (typically a rubidium oscillator disciplined by satellite navigation signals) at each end, time-tagging electronics resolving tens of picoseconds, and bright synchronisation pulses interleaved into the quantum stream — one every thousand pulses or so — to keep the two clocks locked to each other.

Doppler, in two flavours

A source hurtling towards you and then away from you shifts the wavelength of its light:

$$\Delta\lambda / \lambda = v_{los} / c$$

RELATIVISTIC DOPPLER SHIFT, FIRST ORDER — EQ. (3)

With a line-of-sight velocity of a few kilometres per second, that fraction is around one part in a hundred thousand — a shift of roughly ten *picometres* at 850 nm. Against a typical filter 0.1 nanometres wide, that is a tenth of the passband: noticeable, but harmless. It only becomes an engineering headache if you narrow the filter drastically, which is exactly what daylight operation demands. Then the filter must be actively temperature-tuned to chase the shift through the pass.

The more troublesome effect is temporal. As the range changes at kilometres per second, the flight time of the photons drifts — by around ten *microseconds* for every second that elapses. Against a one-nanosecond detection window, that is enormous. The receiver must continuously predict and compensate the drift from the orbital model, correcting itself against the arriving synchronisation pulses. Get it wrong by a nanosecond and the gate closes on empty sky.

> **Three beams, one telescope**
>
> A working link is not one optical channel but three, sharing the same aperture and separated by wavelength and brightness: a strong *beacon* for tracking, bright *synchronisation* pulses for the clock, and the almost-invisible *quantum* signal carrying the key. Designing them so the loud ones do not blind the receiver looking for the quiet one is its own small art.

· CHALLENGE THREE ·

6. The Sky Gets a Vote

Every operational satellite QKD system in the world today works at night, in clear weather. This is not conservatism or a lack of engineering ambition. It follows from the physics of trying to detect one photon in the presence of a star.

The sun is a very loud transmitter

A receiver cannot distinguish a signal photon from a background photon that happens to arrive in the same direction, at the same wavelength, at the same instant. So rejection has to happen in all three domains at once: a narrow spectral filter, a tight field of view, and the nanosecond time gate. In darkness, at a good site, those measures reduce background to a few hundred counts per second against a signal of tens of thousands — an error rate around 1 %, comfortably below the threshold at which a key can still be extracted.

In daylight the sky delivers something like a million counts per second into the same filter. The error rate saturates at 50 % — the value you would get from pure noise — and no key exists at any rate. Daylight quantum links *have* been demonstrated, using extreme filtering, longer wavelengths where the sky is dimmer, and single-mode fibre as an ultra-tight spatial filter. They are impressive. They are also slower, far more complex, and not yet how any national network schedules its service.

The moon is a milder version of the same problem. Around full moon the background can climb tenfold, and operators plan around it much as astronomers do.

Cloud is not attenuation. It is a wall

> **Why weather is binary**
>
> A radio link degrades gracefully through cloud and rain. An optical quantum link does not degrade — it stops. Even modest cloud imposes tens of decibels of loss on a budget that has no margin left. There is no clever modulation, no error-correcting code, no power increase that recovers it. The only engineering responses are

> choosing dry, high sites; spreading stations across regions with uncorrelated weather; and scheduling around forecasts.

This is why the operational model of every deployed system is not real-time encryption but **key accumulation**. On clear nights the satellite fills a buffer of key material at each ground station; that buffer is then drawn down during the day, in fog, and through the monsoon, feeding conventional symmetric encryption. If the buffer holds a month's worth, a month of bad weather is invisible to the people using the network.

The weather problem has a much sharper edge for entanglement-based architectures. There, both ground stations must have clear sky *at the same moment*. If one site is usable on 60 % of nights and another on 70 %, the joint availability is roughly the product — 42 % — and it falls further as you add stations. Prepare-and-measure links need only one site clear per pass, and can collect the other half of the relationship on a different night entirely.

Turbulence, the astronomer's old enemy

Finally, the atmosphere is not a uniform slab. Thermal turbulence in the last few kilometres above the receiver bends and breaks the incoming wavefront — the same effect that makes stars twinkle. For a quantum link it causes the beam to wander off the detector, and it destroys the ability to couple the light into a single-mode fibre, which the best detectors want. Mitigations range from a simple fast tip–tilt mirror to full adaptive optics with a deformable mirror. All of them are cheaper than choosing a bad site and trying to compensate afterwards.

· CHALLENGE FOUR ·

7. Whom Do You Trust?

Suppose every engineering problem above is solved. A subtler one remains, and it is the reason the field has not simply converged on a single global system.

The trusted node problem

Recall the relay: the satellite makes a key with each station and publishes the exclusive-or. It works beautifully and it means the spacecraft holds both secrets in the clear. The physics guarantee — that no eavesdropper on the *channel* can learn the key — says nothing whatsoever about the operator of the relay.

For a network run by and for a single organisation, that may be entirely acceptable: the satellite is inside the trust perimeter, like a safe in the basement. For anyone buying a turnkey system built abroad, it is not a footnote but the central question. The manufacturer of the payload wrote the firmware, supplied the random number generator, and specified the key store. The whole point of an information-theoretic key is to remove trust assumptions; a foreign-built trusted node re-introduces a different one and merely relocates the problem.

Quantum key distribution eliminates the need to trust mathematics. It does not eliminate the need to trust whoever built the box.

The entanglement escape route, and its cost

There is a way out, and it is the reason Ekert's 1991 protocol still matters. If the satellite instead generates *entangled pairs* and sends one photon to each of two ground stations, it never learns the key at all. The correlations are certified by a Bell test performed by the ground stations themselves. Security survives even if the spacecraft was built by your adversary. This is what "untrusted node" means, and as a security property it is close to ideal.

The price is arithmetic. For prepare-and-measure, the key rate falls in proportion to the channel transmittance η. For dual-downlink entanglement, both photons must survive independent journeys:

$$R_{single} \propto \eta \qquad \text{versus} \qquad R_{dual} \propto \eta_A \cdot \eta_B$$

RATE SCALING: SINGLE VS. DUAL DOWNLINK — EQ. (4)

In decibels, the losses add. A downlink of 40 dB becomes a dual downlink of 70 to 90 dB. Recall that 40 dB already meant one part in ten thousand; 80 dB means one part in a hundred million. That is precisely why the best entanglement-based key on record ran at 0.12 bits per second while contemporary prepare-and-measure links delivered megabits per pass. At 0.12 bits per second, a single 256-bit key takes over half an hour of perfect conditions at both stations simultaneously.

There is a further practical obstacle. To serve pairs of stations that are *not* simultaneously visible — which, for a low-orbit satellite, is most pairs on Earth — you would need to store one photon of an entangled pair until the second station comes into view. That requires a space-qualified *quantum memory* with storage times and efficiencies far beyond anything demonstrated. Together with entanglement swapping between satellites, that is the technology that would eventually enable a true global quantum internet with no trusted nodes anywhere. It is a research programme for the 2040s, not a procurement option today.

Attacking the hardware instead of the mathematics

The final trust problem is the oldest lesson in the field, and the one first taught by that audibly clicking bench apparatus in 1989. Security proofs describe idealised devices. Real devices have imperfections, and imperfections are attack surfaces. A partial catalogue of demonstrated attacks on real QKD systems:

- **Detector blinding.** Shine a bright continuous light into the receiver and the single-photon detectors stop behaving quantum-mechanically, becoming ordinary classical light meters an attacker can drive at will — while both parties see a perfectly normal-looking error rate [Lydersen2010].
- **Time-shift and efficiency-mismatch attacks.** Two detectors are never identical; their sensitivity peaks at slightly different instants. Nudge the arrival time and you bias which outcome gets recorded.
- **Trojan-horse probing.** Send light *into* the transmitter and read the reflections to learn what state it is currently preparing. The countermeasure is a sixty-decibel optical isolator and a watchdog photodiode.
- **Source side channels.** If the four polarisation states differ measurably in their spectrum, timing or spatial mode, an eavesdropper may distinguish them without touching the polarisation at all.
- **Weak randomness.** If the basis choices are predictable, everything else is irrelevant. The random number generator is the most security-critical component in the system and the easiest to neglect.

None of these breaks quantum mechanics; all of them break products. The field's response has been to formalise it: international standards for the security requirements and test methods of QKD systems now exist, along with a Common Criteria protection profile for prepare-and-measure modules — meaning independent laboratories can evaluate whether a given box actually deserves the guarantee its brochure claims [Xu2020].

· REALITY CHECK ·

8. The Numbers, Without the Marketing

Quantum communication attracts more than its share of enthusiastic claims. The most useful antidote is a table of what has actually been measured, and what the physics permits.

Architecture	Channel loss	Demonstrated key rate	Status
Trusted node, prepare-and-measure downlink	~35–45 dB	kbit/s during a pass; up to 1.07 Mbit total per pass	Operational; several spacecraft flying
Untrusted node, dual-downlink entanglement	~70–90 dB	0.12 bit/s over 1,120 km	Demonstrated once, as an experiment
Geostationary orbit	Adds ~33 dB over an 800 km orbit	Studied, not flown for QKD	Feasible on paper with metre-class space optics
Terrestrial fibre, no relay	0.2 dB/km — 100 dB at 500 km	Mbit/s below 100 km; effectively zero past ~400 km	Widely deployed at metropolitan scale

Why not geostationary orbit?

It is a fair question, and a popular one, because a geostationary satellite hangs motionless in the sky: no tracking slew, no Doppler sweep, no five-minute scramble. The problem is that geometric loss scales with the square of the distance, and geostationary orbit is 45 times farther away than a typical low orbit. That is an extra 33 decibels — a factor of two thousand — on a budget that was already down to one part in ten thousand. Recovering it means metre-class diffraction-limited optics in space and a spacecraft an order of magnitude larger and more expensive. Feasibility studies conclude it can be done, and it has an interesting advantage for untrusted-node schemes, since a geostationary platform can see two distant ground stations continuously. Nobody has flown one for quantum key distribution.

What a night's worth of key is actually good for

Suppose a station collects a few hundred kilobits to a megabit of secret key on a good night. What does that buy?

- **Plenty:** refreshing the symmetric keys of a few hundred high-value links several times a day. A 256-bit key is 256 bits; a megabit is thousands of them.
- **Plenty:** seeding a key-management system that then derives session keys conventionally, so that the quantum material anchors a much larger classical structure.
- **Just about:** one-time-pad encryption of short, extremely sensitive messages — a megabit is around 125 kilobytes of perfectly unbreakable text per night.
- **Absolutely not:** one-time-pad encryption of ordinary traffic. Video calls, database replication, or a national network's day-to-day communications would exhaust a night's key production in seconds. Any proposal that implies otherwise has an arithmetic problem.

The same arithmetic explains why nobody builds a ground station in every town. A single satellite in low orbit gets one or two usable night passes over a given site, and can realistically service only a handful of stations per night in total. Serving hundreds of sites would require either hundreds of satellites or a wait of many months between visits — and each station is an observatory costing on the order of a million dollars or more to build and a six-figure sum every year to run. Every serious network therefore looks the same: a

small number of regional hubs collecting key from orbit, with conventional cryptography carrying the assurance the last few hundred kilometres to the users.

· THE ARGUMENT ·

9. Is Any of This Worth It?

Given everything above — night-only service, weather dependence, observatory-grade ground stations, trusted relays and a decade of hardware attacks — a reasonable person might ask why anyone bothers, when post-quantum cryptography is a software update.

That is not a rhetorical question. It is the live position of several of the world's most capable signals-intelligence and cyber-security agencies. The United States National Security Agency has published guidance stating that it does not support the use of quantum key distribution for protecting communications in national security systems, and lists five specific limitations: it provides no source authentication on its own; it requires special-purpose hardware; trusted relays add cost and insider risk; validating real implementations is difficult; and it is unusually easy to deny service to [NSA]. The United Kingdom's National Cyber Security Centre reaches a similar conclusion and recommends post-quantum cryptography as the preferred mitigation [NCSC]. A joint position paper from French, German, Dutch and Swedish authorities concluded in 2024 that QKD is presently usable only in niche cases and that clear priority should go to post-quantum cryptography [ANSSI2024].

Each of those five limitations is technically correct, and the first deserves unpacking because it looks fatal at first glance.

THE AUTHENTICATION PARADOX

QKD guarantees that nobody eavesdropped on the key. It does not tell you *who* is at the far end of the beam. Without authentication of the public discussion channel, an attacker could simply impersonate the satellite and establish a perfectly secure key with a victim. So QKD needs classical authentication — either a pre-shared symmetric key or a post-quantum signature. Which prompts the obvious objection: if you need conventional cryptography anyway, why not use it for the key as well?

The answer turns on a difference in *lifetime*. Authentication only has to hold at the instant it is used. If the signature algorithm is broken in 2045, the attacker gains the ability to impersonate people from 2045 onward; he cannot travel back to a night in 2028 and physically insert himself into a ten-metre-wide beam between a spacecraft and a guarded mountain-top telescope. Confidentiality is the opposite: if the key-establishment mathematics falls in 2045, every recorded session from 2028 unlocks retroactively. So the sensible design uses each tool where it is strong — post-quantum signatures for the fleeting act of authentication, quantum physics for the decades-long obligation of secrecy.

The other four limitations are real and are handled, where they are handled at all, by architecture rather than denial: hybrid systems keep a conventional key pool alongside the quantum one and fall back automatically when cloud or jamming stops the photons; independent evaluation against published standards addresses implementation flaws; and owning the spacecraft addresses the trusted node.

Meanwhile, some of the same countries whose agencies signed those cautions are funding satellite QKD programmes. That is less contradictory than it appears. The agencies' guidance addresses general-purpose government and commercial networks, where QKD is genuinely a poor fit. The funding addresses a narrower question: whether a state wants the option of a key whose secrecy does not depend on a mathematical assumption, and whose root of trust it controls end to end. Framed that way, the technology is not a competitor to post-quantum cryptography. It is an insurance policy on it, purchased for a small number of links where the cost of eventual failure is effectively unbounded.

The honest summary

For 99 % of the world's traffic — banking, messaging, commerce, ordinary government business — post-quantum cryptography is the correct and sufficient answer, and satellite QKD would be an absurd extravagance. For a small class of secrets that must survive fifty years, and for organisations unwilling to bet those secrets on a mathematical conjecture holding for half a century, a physics-based key has a property no algorithm can offer. Both statements are true at once. Most of the public argument consists of people asserting one of them at people asserting the other.

· HORIZON ·

10. What Comes Next

Four lines of work will determine whether satellite quantum communication becomes ordinary infrastructure or remains a specialist instrument.

01 **Smaller, cheaper, more numerous**

The trajectory from a 635-kg science satellite to a 100-kg microsatellite to CubeSat demonstrators, and from multi-tonne observatories to portable ground terminals under 100 kg, is the single most consequential trend in the field. A technology that requires a national observatory serves a handful of users; one that fits in a van serves many.

02 **Daylight operation**

Moving to longer wavelengths where the sky is dimmer, filtering in wavelength, angle and time simultaneously, and using single-mode fibre as an ultra-tight spatial filter have all produced working daylight links on the ground. Carrying that into routine orbital service would roughly double the available contact time and remove the field's most awkward operational constraint.

03 **Getting rid of trusted nodes**

Measurement-device-independent and twin-field protocols close the detector attack surface and change how the key rate scales with loss. Entanglement payloads keep improving. But the decisive step is a space-qualified quantum memory that can hold half of an entangled pair until a second ground station comes into view — the enabling technology for entanglement swapping and, eventually, a genuine quantum internet.

04 **Integration, which is where value actually appears**

Key that nothing consumes is worthless. The unglamorous work — standard interfaces between quantum systems and key-management infrastructure, policy engines that decide which key source serves which application, automatic fallback when the sky closes — is what turns a physics demonstration into a service. Every mature programme has discovered this the same way: by building the optics first and then finding nothing plugged into them.

It is worth stepping back to appreciate what has already happened. In 1989 the first quantum key crossed 32 centimetres of laboratory air. Within thirty years, single photons prepared aboard a spacecraft were being caught by telescopes on the ground and turned into keys that linked cities on opposite sides of the planet. The physics never changed — Wiesner's insight about conjugate variables is the same one exploited today. What changed was that engineers learned to point a beam of individual photons at a moving target from eight hundred kilometres away, to catch them with superconducting wire, and to know, to within a nanosecond, exactly when to look.

The security of quantum key distribution is guaranteed by the laws of physics. Everything else about it — the pointing, the clocks, the cryogenics, the weather, the trust — is engineering. And engineering is where it will be won or lost.

· REFERENCE ·

Quick Reference Glossary

Term	Meaning
QKD	Quantum key distribution: deriving shared secret bits from quantum measurements, with an eavesdropper's knowledge bounded by physics rather than by computational difficulty.
BB84	The 1984 prepare-and-measure protocol in which random bits are encoded in randomly chosen conjugate bases. Still the basis of every operational satellite system.
Decoy states	Randomly varying the brightness of transmitted pulses so that photon-number-splitting attacks reveal themselves in the statistics.
No-cloning theorem	An unknown quantum state cannot be copied. The source of QKD's security — and the reason quantum signals cannot be amplified.

Term	Meaning
QBER	Quantum bit error rate: the fraction of mismatched bits in the sifted key. Above roughly 11 % for BB84, no secret key can be extracted at all.
Sifting	Publicly comparing measurement bases and discarding events where they disagreed, without revealing any bit values.
Privacy amplification	Hashing a partially compromised string down to a shorter one about which an eavesdropper's information is provably negligible.
Finite-key analysis	Statistical corrections accounting for the limited number of photons collected in a short pass. Ignoring it inflates quoted key rates.
Trusted node	A relay that holds key material in the clear — such as a satellite combining two keys by exclusive-or. Secure against channel eavesdroppers, not against its own operator.
Dual downlink	Entanglement architecture sending one photon of a pair to each of two ground stations. Removes the trusted node; multiplies the loss.
Optical ground station	The receiving end: a diffraction-limited telescope, precision mount, dome, filters, single-photon detectors and timing electronics. An observatory, not a dish.
SNSPD	Superconducting nanowire single-photon detector: over 90 % efficient with picosecond timing, at the cost of cooling to a couple of degrees above absolute zero.
Post-quantum cryptography	Classical algorithms believed hard for quantum computers, standardised in 2024. Software, scalable, and the primary answer for almost all traffic.
Harvest now, decrypt later	Recording encrypted traffic today to decrypt it once a capable quantum computer exists. The reason long-lived secrets are already exposed.

References

[Wiesner1983] Wiesner, S.: Conjugate coding. *ACM SIGACT News* 15(1), 78–88 (1983). Written c. 1970; rejected by journals for over a decade.

[BB84] Bennett, C.H., Brassard, G.: Quantum cryptography: public key distribution and coin tossing. *Proc. IEEE Int. Conf. on Computers, Systems and Signal Processing*, Bangalore, 175–179 (1984). Reprinted in *Theor. Comput. Sci.* 560, 7–11 (2014).

[Ekert1991] Ekert, A.K.: Quantum cryptography based on Bell's theorem. *Phys. Rev. Lett.* 67(6), 661–663 (1991).

[Bennett1992] Bennett, C.H., Bessette, F., Brassard, G., Salvail, L., Smolin, J.: Experimental quantum cryptography. *J. Cryptology* 5, 3–28 (1992).

[Kurtsiefer2002] Kurtsiefer, C., et al.: A step towards global key distribution. *Nature* 419, 450 (2002).

[Schmitt-Manderbach2007] Schmitt-Manderbach, T., et al.: Experimental demonstration of free-space decoy-state quantum key distribution over 144 km. *Phys. Rev. Lett.* 98, 010504 (2007).

[Ursin2007] Ursin, R., et al.: Entanglement-based quantum communication over 144 km. *Nature Physics* 3, 481–486 (2007).

[Nauerth2013] Nauerth, S., et al.: Air-to-ground quantum communication. *Nature Photonics* 7, 382–386 (2013).

[WangJY2013] Wang, J.-Y., et al.: Direct and full-scale experimental verifications towards ground–satellite quantum key distribution. *Nature Photonics* 7, 387–393 (2013).

[Liao2017] Liao, S.-K., et al.: Satellite-to-ground quantum key distribution. *Nature* 549, 43–47 (2017). doi:10.1038/nature23655

[Yin2017] Yin, J., et al.: Satellite-based entanglement distribution over 1200 kilometers. *Science* 356, 1140–1144 (2017). doi:10.1126/science.aan3211

[Liao2018] Liao, S.-K., et al.: Satellite-relayed intercontinental quantum network. *Phys. Rev. Lett.* 120, 030501 (2018).

[Yin2020] Yin, J., et al.: Entanglement-based secure quantum cryptography over 1,120 kilometres. *Nature* 582, 501–505 (2020). doi:10.1038/s41586-020-2401-y

[Villar2020] Villar, A., et al.: Entanglement demonstration on board a nano-satellite. *Optica* 7(7), 734–737 (2020).

[Lu2022] Lu, C.-Y., Cao, Y., Peng, C.-Z., Pan, J.-W.: Micius quantum experiments in space. *Rev. Mod. Phys.* 94, 035001 (2022).

[Li2025] Li, Y., et al.: Microsatellite-based real-time quantum key distribution. *Nature* 640, 47–54 (2025). doi:10.1038/s41586-025-08739-z

[Chen2025] Chen, H.-Z., et al.: China quantum communication network. *npj Quantum Information* 11, 137 (2025).

[Lydersen2010] Lydersen, L., et al.: Hacking commercial quantum cryptography systems by tailored bright illumination. *Nature Photonics* 4, 686–689 (2010).

[Xu2020] Xu, F., Ma, X., Zhang, Q., Lo, H.-K., Pan, J.-W.: Secure quantum key distribution with realistic devices. *Rev. Mod. Phys.* 92, 025002 (2020).

[Pirandola2020] Pirandola, S., et al.: Advances in quantum cryptography. *Adv. Opt. Photon.* 12(4), 1012–1236 (2020).

[FIPS2024] National Institute of Standards and Technology: FIPS 203 (ML-KEM), FIPS 204 (ML-DSA) and FIPS 205 (SLH-DSA), August 2024.

[NSA] National Security Agency / Central Security Service: *Quantum Key Distribution (QKD) and Quantum Cryptography (QC)*, cybersecurity guidance.

[NCSC] UK National Cyber Security Centre: *Quantum security technologies*, white paper.

[ANSSI2024] ANSSI (France), BSI (Germany), NLNCSA (Netherlands) and the Swedish National Communications Security Authority: *Position paper on quantum key distribution* (2024).