Quantum Cryptography Deck

Must-Know Concepts, Explained

Coverage	Quantum Foundations, QKD Protocols, Security & Eavesdropping Detection, Practical Systems, Post-Quantum Cryptography, Quantum Threats, Quantum Networks, Real-World Deployment
Format	Question → Concise answer → comparison / intuition where useful
Level	Graduate / Research to Applied Security Engineering rounds

Compiled as a quick-revision reference. Skim the bold question, then use the boxed answer to check your recall before interviews or exams.

Generated September 11, 2026

1 Quantum Foundations for Cryptography 2

2 Quantum Key Distribution (QKD) Protocols 4

3 Security & Eavesdropping Detection 6

4 Practical QKD Systems 8

5 Post-Quantum Cryptography (PQC) 10

6 Quantum Threats to Classical Cryptography 12

7 Quantum Networks & Repeaters 13

8 Real-World Deployment & Challenges 15

Quantum Foundations for Cryptography

Q1.1

What is the no-cloning theorem, and why is it the single most important physical principle underlying quantum cryptography?

The no-cloning theorem states that it is **physically impossible** to create an exact, independent copy of an arbitrary unknown quantum state – this follows directly from the linearity of quantum mechanics (a hypothetical universal cloning operation would have to act linearly on superpositions, and no such linear operation can copy every possible input state correctly). This matters enormously for cryptography because it means an eavesdropper **cannot** silently intercept a quantum-encoded message, make a perfect copy to analyze at leisure, and forward the original undisturbed – any attempt to extract information from an unknown quantum state **necessarily disturbs it** in a statistically detectable way. This is the physical root of quantum key distribution's security guarantee, replacing the computational-hardness assumptions that classical cryptography relies on with a law of physics.

Q1.2

Why does measuring a quantum state in the "wrong" basis disturb it, and how does this enable eavesdropper detection?

A qubit's information can be encoded in different, mutually **incompatible** measurement bases (e.g. the rectilinear basis $\{|0\rangle, |1\rangle\}$ and the diagonal basis $\{|+\rangle, |-\rangle\}$) – a state prepared in one basis, when measured in the other, incompatible basis, yields a genuinely **random** outcome (by the Heisenberg uncertainty principle applied to non-commuting observables), and this measurement **collapses** the state, irreversibly destroying the original information it carried in the basis it was actually prepared in. An eavesdropper who doesn't know which basis a qubit was encoded in must guess, and roughly half the time will guess the wrong basis – introducing errors into the qubit that the legitimate sender and receiver can later detect statistically by comparing a subset of their results, even though they can never directly "see" the eavesdropper.

Q1.3

What is quantum entanglement, and why does it not, by itself, allow faster-than-light communication (avoiding a naive misconception)?

Entanglement is a correlation between two or more quantum systems such that their combined state cannot be described as simply the product of independent individual states – measuring one entangled particle instantaneously determines the correlated outcome of a measurement on its (possibly distant) partner, with correlations stronger than any classical theory allows. This does **not** enable faster-than-light signaling because the outcome of a measurement on either particle is fundamentally **random** from the local observer's perspective – there's no way for the person measuring one particle to **choose** what outcome the other party sees, so no controllable information is transmitted; the correlation is only revealed and made useful once the two parties later compare their results through an ordinary classical channel, which is itself bound by the speed of light.

Q1.4

What is the distinction between a "prepare-and-measure" QKD protocol and an "entanglement-based" QKD protocol?

In a **prepare-and-measure** protocol (e.g. BB84), the sender (Alice) actively prepares individual qubits in specific states/bases and sends them to the receiver (Bob), who measures them – security relies on Alice's active, trusted state preparation. In an **entanglement-based** protocol (e.g. E91), a source (which may even be untrusted, potentially controlled by an adversary) distributes entangled particle pairs to both Alice and Bob, who each independently

measure their own particle in a randomly chosen basis – the raw key emerges from their correlated measurement outcomes rather than from a value Alice actively chose and sent, and security can further be certified by measuring a Bell-inequality violation, verifying the particles were genuinely entangled without needing to trust the source at all.

Q1.5

What is a qubit, and how is quantum information physically encoded for use in a QKD photon (e.g. via polarization)?

A qubit is the quantum analogue of a classical bit – instead of being restricted to a definite 0 or 1, it can exist in a **superposition** $\alpha|0\rangle + \beta|1\rangle$ of both basis states simultaneously, only resolving to a definite classical outcome upon measurement, with probabilities given by $|\alpha|^2$ and $|\beta|^2$. In practical QKD systems using photons, a common encoding is **polarization**: horizontal/vertical polarization might represent $|0\rangle$/$|1\rangle$ in the rectilinear basis, while $+45°$/$-45°$ diagonal polarization represents the complementary diagonal basis – other systems instead encode information in **phase** (via interferometers) or **time-bin** (the photon's arrival time slot), each with different practical trade-offs for transmission over fiber versus free space.

Quantum Key Distribution (QKD) Protocols

Q2.1

Walk through the BB84 protocol step by step – what do Alice and Bob each actually do?

1. **Alice** randomly chooses a bit value (0 or 1) and randomly chooses one of two bases (rectilinear or diagonal) to encode it in, then sends the corresponding photon to Bob.
2. **Bob**, for each received photon, randomly chooses a basis (independently of Alice) to measure it in, and records the outcome.
3. Over a public (but authenticated) classical channel, Alice and Bob announce **which basis** they used for each photon (*not* the bit value) – they keep only the results where their bases happened to match (roughly half, on average), discarding the rest. This surviving set is the "sifted key."
4. They sacrifice a random **sample** of the sifted key, publicly comparing those bit values to estimate the error rate; if it's below a threshold, they proceed to error correction and privacy amplification on the remaining, unrevealed bits to produce the final secret key.

Published by Bennett and Brassard in 1984, this was the first QKD protocol, and remains the most widely referenced.

Q2.2

What is the difference between BB84 and the B92 protocol?

BB84 uses **two** bases, each with two possible states (four states total), and Alice's basis choice is part of what's kept secret until the sifting step. **B92** (proposed by Bennett in 1992) simplifies this to using only **two non-orthogonal states** total (rather than four), one representing bit 0 and the other bit 1 – security still relies on the fact that these two states are non-orthogonal and therefore cannot be perfectly, reliably distinguished by an eavesdropper. B92 is conceptually simpler and requires generating fewer distinct states, but is generally considered somewhat more vulnerable to certain practical attacks and less commonly deployed in real systems compared to BB84 and its more hardened variants.

Q2.3

How does the E91 protocol (Ekert 1991) use Bell's theorem to guarantee security, and how does this differ conceptually from BB84's security argument?

E91 distributes entangled photon pairs to Alice and Bob, who each measure their own photon using **randomly chosen** measurement settings from a predefined set. After many rounds, they use a subset of their correlated results to compute a **Bell inequality** (specifically, the CHSH inequality) – if the measured correlations **violate** the inequality by the amount quantum mechanics predicts for genuinely entangled particles, this itself certifies that the particles were truly entangled and that no eavesdropper could have learned the key without breaking that entanglement (and thus reducing the violation). This is a fundamentally different security argument than BB84's: rather than relying on trusting that qubits were correctly prepared and that basis-mismatch errors reveal tampering, E91's security is grounded directly in the non-local, statistically verifiable signature of entanglement itself – a property BB84 doesn't use or require at all.

Q2.4

What is "sifting," and why does roughly half the raw transmitted data get discarded in BB84?

Sifting is the process of comparing (over the public classical channel) which **basis** Alice used to prepare each qubit against which basis Bob used to measure it, and keeping only the results where the two **happened to match** – since Bob chooses his measurement basis independently and randomly, with two possible bases, he matches Alice's basis only about **50% of the time** purely by chance. When the bases don't match, Bob's measurement outcome is essentially random noise uncorrelated with Alice's original bit, and must be discarded, since keeping it would introduce a high, unpredictable error rate into the key rather than genuine shared secret information. This roughly 50% sifting loss is an inherent, expected overhead of the protocol, not a sign of a problem.

Q2.5

What are error correction and privacy amplification, and why are both necessary *after* the raw sifted key is obtained?

Even in the absence of any eavesdropper, real-world channel noise and detector imperfections mean Alice's and Bob's sifted keys won't be perfectly identical – **error correction** (e.g. via a protocol like Cascade) reconciles these discrepancies over the public channel so both parties end up with an identical key, though this process itself necessarily leaks *some* information about the key to a potential eavesdropper listening on the public channel. **Privacy amplification** then applies a carefully chosen hash function that compresses the reconciled key down to a **shorter** final key, in a way that provably reduces any eavesdropper's residual information about the final key to a negligible amount – deliberately sacrificing some key length in exchange for a firm, quantifiable security guarantee on what remains.

Security & Eavesdropping Detection

Q3.1

What is the Quantum Bit Error Rate (QBER), and how is it used to decide whether a QKD session is secure?

QBER is the fraction of bits in the sifted key that **disagree** between Alice's and Bob's records, estimated by publicly comparing a randomly sampled subset of the sifted key (which is then discarded, since it's now been revealed). A baseline QBER is expected even with **no** eavesdropper present, purely from real-world channel noise, detector dark counts, and optical imperfections. If the measured QBER exceeds a theoretically-derived security **threshold** (for BB84, intuitively related to the maximum error rate an intercept-resend eavesdropping strategy could introduce while still learning a useful amount of the key), Alice and Bob **abort** the session entirely, concluding that an eavesdropper's presence cannot be ruled out – rather than trying to salvage a "probably fine" but possibly compromised key.

Q3.2

What is an "intercept-resend" attack, and how much error does it necessarily introduce, illustrating why it's detectable?

In this basic attack, Eve intercepts each qubit in transit, measures it (having to **guess** a basis, since she doesn't know which one Alice used), and then sends a **newly prepared** qubit onward to Bob based on whatever outcome she measured, attempting to hide her intrusion. Whenever Eve guesses the **wrong** basis (roughly half the time), her measurement collapses the state to a random outcome uncorrelated with Alice's original bit, and the qubit she resends to Bob is consequently also prepared incorrectly relative to what Alice intended – when Bob happens to measure in the *same* basis Alice originally used (which sifting would otherwise guarantee gives a perfectly correlated result), he now gets the wrong answer roughly **25%** of the time overall (a 50% chance Eve guessed wrong, times a 50% chance that wrong guess flips Bob's specific outcome) – a substantial, clearly detectable jump in QBER above the expected noise-only baseline.

Q3.3

What is information-theoretic (unconditional) security, and how does it fundamentally differ from the security guarantees of classical public-key cryptography like RSA?

Information-theoretic security means a cryptographic scheme's security holds **absolutely**, regardless of an adversary's computational power – even an adversary with unlimited computing time and a hypothetical future quantum computer cannot break it, because the security doesn't rest on any assumption about a problem being hard to *compute*; it rests on physical law (no-cloning, measurement disturbance) and information-theoretic arguments about how little information an eavesdropper could possibly have gained. Classical public-key schemes like RSA instead rely on **computational hardness assumptions** (e.g. that factoring large integers is intractable for realistic computers) – secure only as long as no sufficiently fast algorithm or sufficiently powerful computer (classical or quantum) exists to solve that underlying hard problem, which is precisely the assumption a large-scale quantum computer running Shor's algorithm would break for RSA specifically.

Q3.4

What is the difference between "individual attacks," "collective attacks," and "coherent (general) attacks" in QKD security proofs, and why does proving security against the most general case matter?

An **individual attack** has Eve interact with and measure each transmitted qubit **independently**, one at a time, using a fixed strategy decided in advance. A **collective attack** lets Eve interact with each qubit individually (storing the

resulting quantum states rather than measuring immediately), but allows her to perform an optimal **joint measurement** across all her stored states *after* the public sifting/reconciliation discussion, using everything revealed there. A **coherent (general) attack** removes essentially all restrictions – Eve may apply an arbitrary, fully general quantum operation across the *entire* sequence of transmitted qubits jointly, representing the most powerful attack physically allowed. A genuine, rigorous security proof must hold against this most general coherent-attack case, since any weaker proof leaves open the theoretical possibility of a stronger, uncountered attack strategy – modern QKD security proofs (e.g. via the quantum de Finetti theorem or entropic uncertainty relations) are specifically constructed to establish this strongest, most general guarantee.

Practical QKD Systems & Implementation

Q4.1

Why can't real QKD systems use a perfect, ideal single-photon source, and what is the Photon Number Splitting (PNS) attack this limitation opens up?

True, on-demand single-photon sources are technologically difficult and expensive; most practical QKD systems instead use **attenuated laser pulses**, dimmed down so that most pulses contain zero or one photon – but a non-negligible fraction of pulses still occasionally contain **two or more** photons, following Poissonian statistics. In a **Photon Number Splitting attack**, Eve exploits exactly these multi-photon pulses: she can, in principle, split off and keep one photon for herself (measuring it later, after learning the basis during public sifting) while letting the remaining photon(s) continue undisturbed to Bob – gaining information without introducing any detectable extra error at all, since Bob's photon was never touched.

Q4.2

What is the decoy-state method, and how does it defend against the PNS attack without requiring a true single-photon source?

The decoy-state method has Alice randomly vary the **intensity** (average photon number) of her transmitted pulses across several different levels – "signal" pulses at the intensity actually used for the final key, plus one or more "decoy" pulses at deliberately different intensities, with Alice keeping track of which pulses were sent at which intensity (revealed only afterward, over the public channel). Because a PNS attack's effectiveness depends on the multi-photon pulse rate, which differs predictably between signal and decoy intensities, Alice and Bob can statistically compare the observed transmission/error rates **across** the different intensity levels to detect any anomalous, PNS-attack-consistent behavior – effectively closing the security loophole opened by using imperfect, weak coherent-pulse sources, without needing to solve the much harder engineering problem of building true single-photon sources.

Q4.3

What fundamentally limits the maximum transmission distance of fiber-based QKD, and roughly what secret key rate versus distance trade-off results?

Photon loss in optical fiber grows **exponentially** with distance (typically around 0.2 dB/km for standard telecom fiber at the optimal wavelength) – and because of the no-cloning theorem, lost photons **cannot** be amplified and regenerated mid-path the way classical optical signals routinely are with repeater amplifiers (a classical amplifier would need to "read" the quantum state to boost it, which is exactly what no-cloning forbids). This means the secret key rate falls off roughly exponentially with distance, and beyond a few hundred kilometers of direct fiber transmission, the achievable key rate becomes impractically low – current record fiber-based QKD experiments reach several hundred km, but practical deployed systems are generally limited to under $\sim$100–200 km without trusted-node relays or a quantum repeater.

Q4.4

What is a "trusted node," and what security assumption does it require that pure end-to-end QKD otherwise avoids?

To extend QKD beyond fiber's inherent distance limit, a chain of intermediate **trusted nodes** can be used: QKD is run independently over each shorter hop (Alice-to-Node1, Node1-to-Node2, ..., NodeN-to-Bob), and each node in the chain **combines** the keys from its two adjacent links (e.g. via XOR) to relay the overall end-to-end key onward. This extends practical range significantly, but requires **trusting** every intermediate node not to record or leak the key material passing through it – reintroducing exactly the kind of trusted-intermediary assumption that end-to-end

QKD's information-theoretic security was designed to avoid in the first place, since a single compromised trusted node anywhere along the chain compromises the entire end-to-end key.

Q4.5

What is Measurement-Device-Independent QKD (MDI-QKD), and what class of attack does it specifically close off?

Real detectors have imperfections (e.g. limited efficiency, or exploitable timing/blinding vulnerabilities) that an attacker can exploit through **detector side-channel attacks** not accounted for in the idealized theoretical security proof – famously, "detector blinding" attacks have been experimentally demonstrated against real commercial QKD systems, extracting the key without triggering the expected QBER increase, by exploiting real hardware flaws outside the clean mathematical model. **MDI-QKD** restructures the protocol so that both Alice and Bob send quantum states to an **untrusted, possibly adversary-controlled** third-party measurement station in the middle, and the key is extracted purely from the announced pattern of joint measurement outcomes (a Bell-state measurement) – crucially, the security proof no longer depends on trusting the detectors' behavior at all, since even a fully malicious measurement station cannot compromise the resulting key, closing off the entire class of detector side-channel attacks by design.

Post-Quantum Cryptography (PQC)

Q5.1

What is Post-Quantum Cryptography, and how does it fundamentally differ in approach from Quantum Key Distribution?

Post-Quantum Cryptography (also called quantum-resistant cryptography) refers to **classical** cryptographic algorithms, running on ordinary classical computers and classical communication channels, that are specifically designed based on mathematical problems believed to remain hard even for a large-scale **quantum** computer to solve. This is a fundamentally different strategy than QKD: PQC doesn't use any quantum phenomena at all in its operation – it's classical cryptography with different, quantum-resistant *computational hardness assumptions* – whereas QKD uses actual quantum physics (photon transmission and measurement) to achieve information-theoretic security independent of any computational assumption. PQC's major practical advantage is that it can be deployed over ordinary existing internet infrastructure without any new quantum hardware at all.

Q5.2

What are the main families of mathematical problems that current NIST-standardized PQC algorithms are built on?

- **Lattice-based** – security based on the hardness of problems like Learning With Errors (LWE) or finding short vectors in high-dimensional lattices; the basis for **CRYSTALS-Kyber** (selected as the primary NIST standard for key encapsulation) and **CRYSTALS-Dilithium** (digital signatures) – currently the most favored family due to a strong balance of security confidence, performance, and reasonably small key sizes.
- **Hash-based** – security relies only on the hardness of finding collisions in a cryptographic hash function, a very well-understood and conservative assumption; the basis for **SPHINCS+** (a NIST-standardized signature scheme), valued for strong security confidence at the cost of larger signature sizes.
- **Code-based** – security based on the hardness of decoding a random linear error-correcting code; the long-studied **McEliece** cryptosystem is a prominent example, offering high security confidence (some of the longest-standing unbroken proposals) but with notably large public key sizes.
- **Multivariate and isogeny-based** schemes are additional, less widely standardized families, with isogeny-based SIKE notably having been **broken** by a classical (non-quantum!) attack in 2022, a cautionary example of why cryptanalysis confidence matters as much as the underlying quantum-resistance motivation.

Q5.3

Why is "harvest now, decrypt later" considered an urgent threat today, even though large-scale, cryptographically-relevant quantum computers don't yet exist?

An adversary can **record and store** encrypted traffic today (e.g. intercepted internet communications encrypted with RSA or elliptic-curve cryptography) without being able to decrypt it *yet* – but if a sufficiently powerful quantum computer capable of running Shor's algorithm at scale becomes available at some point in the future, that adversary could then retroactively decrypt all of that previously-harvested, stored data. This makes the threat urgent *today*, not just once quantum computers actually arrive, specifically for any data whose confidentiality must be protected for a **long time** (decades) into the future – government/military secrets, long-term medical records, or trade secrets – since data encrypted with vulnerable classical algorithms today is already effectively "at risk" the moment a capable quantum computer eventually exists, regardless of when that happens.

Q5.4

What is a ”hybrid” cryptographic scheme combining a classical algorithm with a post-quantum algorithm, and why do organizations often deploy this during the transition period rather than switching directly to pure PQC?

A hybrid scheme combines the output of **both** a well-established classical algorithm (e.g. ECDH) and a newer post-quantum algorithm (e.g. Kyber) to derive a session's shared secret – an attacker would need to break *both* algorithms to compromise the resulting key. This is favored during the transition period because newer PQC algorithms, despite passing rigorous NIST standardization scrutiny, are comparatively **less battle-tested** in real-world deployment than decades-old classical algorithms (the SIKE break, mentioned above, is a sobering reminder that even carefully-vetted PQC candidates can still have undiscovered classical weaknesses) – a hybrid approach provides defense-in-depth, remaining secure even if either the classical algorithm is eventually broken by a quantum computer *or* an unexpected classical/mathematical weakness is later found in the new PQC algorithm.

Quantum Threats to Classical Cryptography

Q6.1

What does Shor's algorithm do, and why does it specifically threaten RSA and elliptic-curve cryptography?

Shor's algorithm is a quantum algorithm that can factor large integers (and solve the closely related discrete logarithm problem) in **polynomial time** on a sufficiently large, fault-tolerant quantum computer – an exponential speedup over the best known classical factoring algorithms, whose intractability for large numbers is precisely the security foundation RSA relies on. Elliptic-curve cryptography (ECDH, ECDSA) relies instead on the hardness of the elliptic-curve discrete logarithm problem, which a variant of Shor's algorithm can also solve efficiently. This means both of the two dominant families of **public-key** cryptography in use across the internet today would become completely broken (not just weakened) by a large enough, error-corrected quantum computer running Shor's algorithm – the primary motivation driving the urgency of the entire PQC transition effort.

Q6.2

What does Grover's algorithm do, and why does it only require doubling key sizes for symmetric cryptography rather than a complete algorithmic replacement?

Grover's algorithm provides a quantum **quadratic** speedup for unstructured search problems – including brute-forcing a symmetric key – reducing the effective search effort for an n-bit key from $O(2^n)$ classical operations down to $O(2^{n/2})$ quantum operations. This is a meaningfully weaker threat than Shor's exponential break of public-key crypto: a quadratic speedup can be fully compensated for simply by **doubling the key length** (e.g. AES-128's effective quantum-era security drops to roughly the equivalent of AES-64, which is inadequate, but AES-256 still offers roughly AES-128-equivalent quantum security, which remains considered strong) – so symmetric algorithms like AES don't need to be replaced with fundamentally different algorithms the way RSA/ECC do; a straightforward, well-understood key-size increase is sufficient mitigation.

Q6.3

What is meant by the number of "logical qubits" versus "physical qubits" required to actually run Shor's algorithm at a cryptographically relevant scale, and why does this distinction matter for assessing current real-world risk?

Today's quantum computers are built from **physical qubits**, which are highly error-prone and lose their quantum state (decohere) quickly – running a genuinely useful, long computation like factoring a 2048-bit RSA key requires **quantum error correction**, which encodes one reliable, low-error **logical qubit** using many (current estimates run from hundreds to well over a thousand) noisy physical qubits working together. Breaking RSA-2048 is estimated to require on the order of thousands of stable logical qubits, translating to potentially millions of physical qubits with current error-correction overhead – far beyond the few hundred to low-thousands of *noisy, uncorrected* physical qubits that today's most advanced quantum computers actually have. This gap is precisely why cryptographically-relevant quantum computers are still assessed to be years to decades away, even as headline "qubit count" numbers continue to grow.

Quantum Networks & Repeaters

Q7.1

What is a quantum repeater, and how does it extend range without violating the no-cloning theorem?

A quantum repeater extends the effective range of quantum communication by dividing a long distance into several shorter segments, generating **entanglement** independently over each short segment (where loss is manageable), and then using a process called **entanglement swapping** at intermediate nodes to ”stitch” these independent short-range entangled links together into a single, longer-range entangled pair spanning the full distance – without ever needing to directly amplify or clone the original quantum state end-to-end. This sidesteps the no-cloning restriction entirely, since no state is ever copied – new entanglement is generated locally at each segment and then connected, rather than an existing quantum state being read out and duplicated.

Q7.2

What is entanglement swapping, and what is entanglement purification, and why are both typically needed together in a quantum repeater architecture?

Entanglement swapping: if particle A is entangled with particle B (at one node), and particle C is entangled with particle D (at an adjacent node), performing a joint Bell-state measurement on particles B and C (both physically present at the middle node) causes particles A and D – which never directly interacted – to become entangled with each other, effectively extending the entangled link’s range without any single particle physically traveling the full distance. **Entanglement purification**: since real-world entanglement generation and storage is imperfect (noisy), purification takes several lower-quality (weakly) entangled pairs and, through a local operation and classical communication protocol, distills them into a **smaller number of higher-fidelity** entangled pairs. Both are typically needed together because swapping alone would compound and worsen the accumulated noise/errors from each segment as the chain grows longer; purification is what keeps end-to-end entanglement quality usable despite that accumulation.

Q7.3

What is a ”quantum memory,” and why is it considered one of the hardest remaining engineering challenges for building practical quantum repeaters?

A quantum memory is a device capable of storing an arrival quantum state (e.g. the quantum state of a photon) for some duration and later releasing it **coherently** (preserving the delicate quantum information, including phase relationships, rather than just recording classical information about it) – necessary in a quantum repeater because entanglement swapping at a middle node requires **both** adjacent entangled links to be successfully established and **simultaneously available** to perform the joint measurement, and since each segment’s entanglement generation is itself probabilistic, one side often succeeds well before the other. Without a quantum memory to hold the first successfully-generated half of an entangled pair coherently while waiting for its neighboring segment to also succeed, the system would need to wait for pure, exponentially-unlikely simultaneous success across every segment at once – quantum memories with practically useful coherence times, high fidelity, and efficient photon-to-memory interfacing remain an active, unsolved area of experimental physics.

Q7.4

What is envisioned as the eventual ”quantum internet,” and how does its purpose differ from simply making today’s classical internet faster or more secure?

A quantum internet is a proposed future network of interconnected quantum devices/nodes able to **distribute and maintain entanglement** between distant parties on demand, enabling applications fundamentally impossible over

any purely classical network – not just faster classical data transfer. Beyond QKD (secure key distribution) itself, envisioned applications include **distributed/blind quantum computing** (letting a client securely use a remote quantum computer's processing power without revealing their input data to it, verified through quantum properties no classical protocol can replicate), quantum-enhanced distributed sensing (correlating entangled sensors across long distances for higher precision than any independent classical sensor network could achieve), and clock synchronization with fundamentally higher precision – the emphasis is on unlocking entirely new capabilities enabled specifically by distributing genuine quantum correlations, not simply "the internet, but encrypted with quantum keys."

Real-World Deployment & Challenges

Q8.1

Why does free-space (satellite) QKD offer a path to much longer distances than fiber, and what was significant about China's Micius satellite experiments?

Photon loss in free space through the vacuum of space (once above the atmosphere) is dramatically lower than the continuous exponential attenuation photons experience traveling through kilometers of optical fiber – a satellite in low Earth orbit acts as a natural, very high "relay point," with loss dominated mainly by the relatively short atmospheric path near each ground station rather than by the (much longer) total ground distance between two far-apart cities. China's **Micius** satellite (launched 2016) demonstrated exactly this, achieving intercontinental QKD (notably a video call secured between China and Austria) and satellite-to-ground entanglement distribution over roughly 1,200 km – distances utterly impractical for direct, un-relayed fiber-based QKD – establishing satellite-based QKD as a genuinely viable path toward global-scale quantum-secured communication.

Q8.2

What is a side-channel attack in the context of practical QKD hardware, and how does it differ from an attack against the theoretical protocol itself?

A side-channel attack doesn't attempt to break the underlying **mathematical/physical security proof** of the protocol at all – instead, it exploits **implementation-specific physical imperfections** in the real hardware that weren't accounted for in the idealized theoretical model the security proof assumes. The detector-blinding attacks mentioned earlier (Section 4) are a prominent real example: by shining bright, continuous light at a single-photon avalanche detector, an attacker can force it into a classical, linear-response mode where its clicks can be **controlled** by the attacker rather than genuinely reflecting a quantum measurement outcome – letting Eve force particular measurement results at Bob's detector without introducing the QBER increase the theoretical security proof assumes any eavesdropping would cause. This illustrates a recurring theme in applied quantum cryptography: a mathematically airtight proof of security doesn't automatically guarantee a specific physical implementation of that protocol is actually secure.

Q8.3

What are quantum random number generators (QRNGs), and why is genuine, certifiable randomness particularly important for cryptography more broadly (not just QKD)?

QRNGs exploit an inherently, fundamentally random quantum process (e.g. which output port a photon exits from at a 50/50 beam splitter) to generate random bits, in contrast to classical pseudo-random number generators, which are actually **deterministic** algorithms that only *appear* statistically random, and whose output can in principle be predicted entirely if the algorithm and its seed are known. Cryptography broadly depends critically on genuinely unpredictable randomness for generating secret keys, nonces, and initialization vectors – a weakness or predictability in the underlying randomness source has historically been the cause of several real-world cryptographic breaks, even when the cryptographic algorithm itself was mathematically sound; QRNGs offer randomness whose unpredictability is grounded directly in the fundamental, provable randomness of quantum measurement outcomes, rather than relying on the practical difficulty of guessing a deterministic algorithm's internal state.

Q8.4

What are the main practical barriers currently limiting widespread commercial/enterprise adoption of QKD, beyond distance limitations?

- **Cost and specialized hardware** – QKD requires dedicated, specialized optical hardware (single-photon detectors, precisely stabilized interferometers) that's considerably more expensive and less mature than standard tele-

com networking equipment.

- **Point-to-point limitation** – QKD fundamentally secures a link between two specific endpoints; scaling this to a full network of many mutually-communicating parties (as opposed to today's flexible, any-to-any classical internet routing) is architecturally far more complex, generally requiring either trusted-node relay chains or advances in quantum repeaters.
- **No solution for authentication or digital signatures on its own** – QKD solves **key distribution** specifically, but still requires an **authenticated** classical channel to prevent a man-in-the-middle attack during the public discussion phase, meaning some form of pre-shared authentication key or classical (potentially quantum-vulnerable, unless PQC is used) authentication scheme is still needed alongside it.
- **Competing with the "good enough" economics of PQC** – for most organizations, upgrading to post-quantum *classical* cryptography is dramatically cheaper and easier to deploy over existing infrastructure than installing entirely new quantum hardware, limiting QKD's near-term commercial market primarily to the highest-security-assurance use cases (government, defense, critical financial infrastructure).

End of deck.