Post-quantum cryptography: risk assessment

Daniel J. Bernstein

15 June 2023

A bit about me

Designing cryptography to proactively reduce risks. Deployed cryptosystems include X25519, Ed25519, ChaCha20, NTRU Prime in TinySSH and OpenSSH, Classic McEliece in Mullvad and Rosenpass.

Coined the phrase “post-quantum cryptography” in 2003.

A bit about an attacker

2012 "Investigative Report on the U.S. National Security Issues Posed by Chinese Telecommunications Companies Huawei and ZTE" by the U.S. House of Representatives Permanent Select Committee on Intelligence:

Chinese intelligence collection efforts against the U.S. government are growing in "scale, intensity and sophistication."[12] Chinese actors are also the world's most active and persistent perpetrators of economic espionage.[13] U.S. private sector firms and cyber-security specialists report an ongoing onslaught of sophisticated computer network intrusions that originate in China, and are almost certainly the work of, or have the backing of, the Chinese government.[14] Further, Chinese intelligence services, as well as private companies and other entities, often recruit those with direct access to corporate networks to steal trade secrets and other sensitive proprietary data.[15]

R0: The basic quantum-attack risk

R0(Y) definition: attackers in year Y have a large enough quantum computer to break RSA-2048 with Shor's algorithm.

Impact: **security disaster** if RSA-2048 is still in wide use in year Y.

R0: The basic quantum-attack risk

R0(Y) definition: attackers in year Y have
a large enough quantum computer
to break RSA-2048 with Shor's algorithm.

Impact: **security disaster**
if RSA-2048 is still in wide use in year Y.

Mitigation: upgrade to post-quantum
cryptography before year Y.

Probability of the basic risk

Global Risk Institute 2022 survey of
40 people working on quantum computing:

- “Optimistic”: reaches 50% in Y = 2027.
- Median: reaches 50% in Y = 2037.
- “Pessimistic”: under 30% in Y = 2052.

Probability of the basic risk

Global Risk Institute 2022 survey of
40 people working on quantum computing:

- “Optimistic”: reaches 50% in Y = 2027.
- Median: reaches 50% in Y = 2037.
- “Pessimistic”: under 30% in Y = 2052.

My assessment: reaches 50% in Y = 2029;
50% for public demonstration in Y = 2032.

Probability of the basic risk

Global Risk Institute 2022 survey of
40 people working on quantum computing:

- “Optimistic”: reaches 50% in Y = 2027.
- Median: reaches 50% in Y = 2037.
- “Pessimistic”: under 30% in Y = 2052.

My assessment: reaches 50% in Y = 2029;
50% for public demonstration in Y = 2032.

Two common mistakes analyzing this risk:

- Assuming attackers aren’t ahead of us.

Probability of the basic risk

Global Risk Institute 2022 survey of
40 people working on quantum computing:

- “Optimistic”: reaches 50% in Y = 2027.
- Median: reaches 50% in Y = 2037.
- “Pessimistic”: under 30% in Y = 2052.

My assessment: reaches 50% in Y = 2029;
50% for public demonstration in Y = 2032.

Two common mistakes analyzing this risk:

- Assuming attackers aren’t ahead of us.
- Watching advances in #qubits and in qubit error rates but not in algorithms.

2022 Jaques: Quantum Landscape

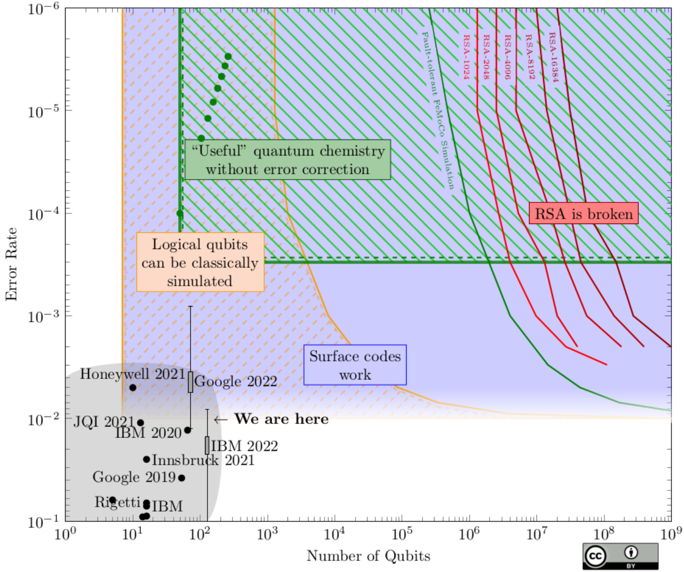

R1: The retroactive-attack risk

R1(Y) definition: attackers are recording RSA-2048-encrypted data today, break RSA-2048 with a quantum computer in year Y, and still find the data useful.

R1: The retroactive-attack risk

R1(Y) definition: attackers are recording RSA-2048-encrypted data today, break RSA-2048 with a quantum computer in year Y, and still find the data useful.

Impact: security disaster **even if we're no longer using RSA-2048 in year *Y***.

R1: The retroactive-attack risk

R1(Y) definition: attackers are recording RSA-2048-encrypted data today, break RSA-2048 with a quantum computer in year Y, and still find the data useful.

Impact: security disaster **even if we're no longer using RSA-2048 in year *Y***.

("Perfect forward secrecy" is broken too.)

R1: The retroactive-attack risk

R1(*Y*) definition: attackers are recording RSA-2048-encrypted data today, break RSA-2048 with a quantum computer in year *Y*, and still find the data useful.

Impact: security disaster **even if we're no longer using RSA-2048 in year *Y***.

("Perfect forward secrecy" is broken too.)

Probability: can be anywhere between 0 and the basic risk, depending on the type of data.

R1: The retroactive-attack risk

R1(Y) definition: attackers are recording RSA-2048-encrypted data today, break RSA-2048 with a quantum computer in year Y, and still find the data useful.

Impact: security disaster **even if we're no longer using RSA-2048 in year *Y***.

("Perfect forward secrecy" is broken too.)

Probability: can be anywhere between 0 and the basic risk, depending on the type of data.

Mitigation: upgrade encryption now!

R2: The upgrade-time risk

R2(U) definition: upgrading takes $>U$ years.

Probability: depends on U and application.

R2: The upgrade-time risk

R2(U) definition: upgrading takes $>U$ years.

Probability: depends on U and application.

Impact: amplifies R0 and R1,
by slowing down mitigations for those.

R2: The upgrade-time risk

R2(U) definition: upgrading takes $>U$ years.

Probability: depends on U and application.

Impact: amplifies R0 and R1,
by slowing down mitigations for those.

Mitigations: upgrade asap; search for paths to faster upgrades; reduce reliance on systems that are hard to upgrade.

R3: The cryptanalysis risk

R3a(C) definition: We’re wrong in thinking that new cryptosystem C is secure against quantum computers.

Impact: Upgrading to that cryptosystem fails to mitigate R0 and R1.

R3: The cryptanalysis risk

R3a(*C*) definition: We're wrong
in thinking that new cryptosystem *C*
is secure against quantum computers.

Impact: Upgrading to that cryptosystem
fails to mitigate R0 and R1.

R3b(*C*) definition: New cryptosystem *C* isn't
secure against *non-quantum* computers.

R3: The cryptanalysis risk

R3a(C) definition: We're wrong
in thinking that new cryptosystem C
is secure against quantum computers.

Impact: Upgrading to that cryptosystem
fails to mitigate R0 and R1.

R3b(C) definition: New cryptosystem C isn't
secure against *non-quantum* computers.

Impact: The upgrade **instantly damages security** against knowledgeable attackers.
Also, the upgrade fails to mitigate R0 and R1.

Submissions to NIST: status in 2017

BIG QUAKE. **BIKE**. **CFPKM**. **Classic McEliece**. **Compact LWE**. **CRYSTALS-DILITHIUM**. **CRYSTALS-KYBER**. **DAGS**. **Ding Key Exchange**. **DME**. **DRS**. **DualModeMS**. **Edon-K**. **EMBLEM and R.EMBLEM**. **FALCON**. **FrodoKEM**. **GeMSS**. **Giophantus**. **Gravity-SPHINCS**. **Guess Again**. **Gui**. **HILA5**. **HiMQ-3**. **HK17**. **HQC**. **KINDI**. **LAC**. **LAKE**. **LEDAkem**. **LEDApkc**. **Lepton**. **LIMA**. **Lizard**. **LOCKER**. **LOTUS**. **LUOV**. **McNie**. **Mersenne-756839**. **MQDSS**. **NewHope**. **NTRU Prime**. **NTRU-HRSS-KEM**. **NTRUEncrypt**. **NTS-KEM**. **Odd Manhattan**. **OKCN/AKCN/CNKE**. **Ouroboros-R**. **Picnic**. **pqNTRUSign**. **pqRSA encryption**. **pqRSA signature**. **pqsigRM**. **QC-MDPC KEM**. **qTESLA**. **RaCoSS**. **Rainbow**. **Ramstake**. **RankSign**. **RLCE-KEM**. **Round2**. **RQC**. **RVB**. **SABER**. **SIKE**. **SPHINCS+**. **SRTPI**. **Three Bears**. **Titanium**. **WalnutDSA**.

Submissions to NIST: status today

BIG QUAKE. BIKE. CFPKM. Classic McEliece. Compact LWE. CRYSTALS-DILITHIUM. CRYSTALS-KYBER. DAGS. Ding Key Exchange. DME. DRS. DualModeMS. Edon-K. EMBLEM and R.EMBLEM. FALCON. FrodoKEM. GeMSS. Giophantus. Gravity-SPHINCS. Guess Again. Gui. HILA5. HiMQ-3. HK17. HQC. KINDI. LAC. LAKE. LEDAkem. LEDApkc. Lepton. LIMA. Lizard. LOCKER. LOTUS. LUOV. McNie. Mersenne-756839. MQDSS. NewHope. NTRU Prime. NTRU-HRSS-KEM. NTRUEncrypt. NTS-KEM. Odd Manhattan. OKCN/AKCN/CNKE. Ouroboros-R. Picnic. pqNTRUSign. pqRSA encryption. pqRSA signature. pqsigRM. QC-MDPC KEM. qTESLA. RaCoSS. Rainbow. Ramstake. RankSign. RLCE-KEM. Round2. RQC. RVB. SABER. SIKE. SPHINCS+. SRTPI. Three Bears. Titanium. WalnutDSA.

Legend: Still in the NIST competition. Less security than claimed. Really broken. Attack scripts.

Probability of cryptanalysis

Quantitative cryptographic risk analysis:

- Define clear mechanisms to quantify the risk of any given cryptosystem.
- Scientifically evaluate these mechanisms.
- Use the best mechanisms to select lowest-risk cryptosystems.

Probability of cryptanalysis

Quantitative cryptographic risk analysis:

- Define clear mechanisms to quantify the risk of any given cryptosystem.
- Scientifically evaluate these mechanisms.
- Use the best mechanisms to select lowest-risk cryptosystems.

Bad news: **This is wishful thinking.**

Probability of cryptanalysis

Quantitative cryptographic risk analysis:

- Define clear mechanisms to quantify the risk of any given cryptosystem.
- Scientifically evaluate these mechanisms.
- Use the best mechanisms to select lowest-risk cryptosystems.

Bad news: **This is wishful thinking.**

Note that success in this direction would eliminate many cryptographic jobs: failures produce cryptographic funding.

Mitigations for cryptanalysis

Minimum mitigation for R3b (not useful for R3a): Don’t throw away the existing encryption layer. Double encrypt using old+new cryptosystems (“hybrid”).

Recommended by, e.g., ANSI, French ANSSI, German BSI. Used in 2019 SIKE experiment; prevented R3b impact there.

Mitigations for cryptanalysis

Minimum mitigation for R3b (not useful for R3a): Don’t throw away the existing encryption layer. Double encrypt using old+new cryptosystems (“hybrid”).

Recommended by, e.g., ANSI, French ANSSI, German BSI. Used in 2019 SIKE experiment; prevented R3b impact there. But NSA has been objecting, and NIST is agnostic.

Mitigations for cryptanalysis

Minimum mitigation for R3b (not useful for R3a): Don't throw away the existing encryption layer. Double encrypt using old+new cryptosystems ("hybrid").

Recommended by, e.g., ANSI, French ANSSI, German BSI. Used in 2019 SIKE experiment; prevented R3b impact there. But NSA has been objecting, and NIST is agnostic.

R3a+R3b mitigation: Take, e.g., Kyber-1024, not Kyber-512. Note that this doesn't *eliminate* risk: largest SIKE version is broken.

R4: The patent risk

R4(C) definition: patent holders ask for money for using new cryptosystem C.

Impact: Most users switch to something else or nothing at all. Upgrades are slower and more expensive. Lack of focus amplifies R3.

R4: The patent risk

R4(C) definition: patent holders ask for money for using new cryptosystem C.

Impact: Most users switch to something else or nothing at all. Upgrades are slower and more expensive. Lack of focus amplifies R3.

Probability: depends on C and on what has been patented. Common mistake: not understanding doctrine of equivalents.

R4: The patent risk

R4(*C*) definition: patent holders ask for money for using new cryptosystem *C*.

Impact: Most users switch to something else or nothing at all. Upgrades are slower and more expensive. Lack of focus amplifies R3.

Probability: depends on *C* and on what has been patented. Common mistake: not understanding doctrine of equivalents.

Mitigations: search patents; analyze patent coverage; use older cryptosystems.

Kyber patent delays, part 1

NIST sets October 2021 deadline
for input regarding the competition.

December 2021: “NIST will be selecting the first post-quantum standards for KEMs (and digital signatures) around the end of December or sometime in early January.”

April 2022: “the delay is not due to technical considerations but is due to some legal and procedural steps that are taking more time than we anticipated”.

Kyber patent delays, part 2

July 2022: NIST announces selection of Kyber for encryption—but says it hasn't signed patent licenses yet.

"NIST expects to execute the various agreements prior to publishing the standard. If the agreements are not executed by the end of 2022, NIST may consider selecting NTRU instead of Kyber."

(NTRU is older; patent expired in 2017. NTRU Prime and Kyber are variants of NTRU.)

Kyber patent delays, part 3

November 2022: NIST says it has signed two patent licenses—but one license won't activate until NIST issues a standard, probably 2024.

License text is only for "a NIST Special Publication or Federal Information Processing Standard" and specifically disallows any "modification, extension, or derivation of the parameters of the PQC ALGORITHM".

Why Kyber rather than NTRU?

Instead of delaying usage until 2024
(and possibly further if more patents apply),
NIST could have announced NTRU in 2021.

Why Kyber rather than NTRU?

Instead of delaying usage until 2024
(and possibly further if more patents apply),
NIST could have announced NTRU in 2021.

July 2022 report says that NIST finds Kyber "marginally more convincing" than NTRU but that "NIST is confident in the security that each provides".

Why Kyber rather than NTRU?

Instead of delaying usage until 2024
(and possibly further if more patents apply),
NIST could have announced NTRU in 2021.

July 2022 report says that NIST finds Kyber "marginally more convincing" than NTRU but that "NIST is confident in the security that each provides".

The only decisive-sounding factor listed is performance: "With regard to performance, Kyber was near the top (if not the top) in most benchmarks."

R5: The computer-overload risk

R5(C) definition: new cryptosystem C is so big or so slow that users cannot afford it.

R5: The computer-overload risk

R5(C) definition: new cryptosystem C is
so big or so slow that users cannot afford it.

Impact: similar to R4.

R5: The computer-overload risk

R5(C) definition: new cryptosystem C is so big or so slow that users cannot afford it.

Impact: similar to R4.

Mitigations: change protocols to use cryptosystem C more efficiently; use smaller, faster cryptosystems.

R5: The computer-overload risk

R5(C) definition: new cryptosystem C is so big or so slow that users cannot afford it.

Impact: similar to R4.

Mitigations: change protocols to use cryptosystem C more efficiently; use smaller, faster cryptosystems.

Probability: low for all C of interest here. But commonly portrayed as high, driving selection of cryptosystems that amplify risks R3 and R4.

Example of how NIST analyzed R5

July 2020 NIST report:

- FrodoKEM in TLS key exchange would cost “around 20,000 bytes” plus “2 million cycles” for the server.
- “NIST’s first priority for standardization is a KEM that would have acceptable performance in widely used applications overall” so NIST is punting on FrodoKEM.

Example of how NIST analyzed R5

July 2020 NIST report:

- FrodoKEM in TLS key exchange would cost “around 20,000 bytes” plus “2 million cycles” for the server.
- “NIST’s first priority for standardization is a KEM that would have acceptable performance in widely used applications overall” so NIST is punting on FrodoKEM.

I request explanation of the basis for the claim that 20000 bytes plus 2 million cycles would not be “acceptable performance” for post-quantum TLS key exchange.

Example, continued

NIST's answer: "While it is not possible to speak for what every user of our standards would or wouldn't find 'acceptable', there is a pretty large difference between the performance of Frodo on the one hand and Kyber, NTRU, and Saber on the other hand. We are therefore more confident that Kyber, NTRU, or Saber will be considered 'acceptable' for most users than that Frodo will."

What's next?

You won't be fired for this strategy: "We'll form a committee to devise an action plan to inventory current usage of cryptography to support future assessment of the steps needed to build a best-practices playbook for meeting the performance challenges of upgrading to post-quantum cryptography, with a target date after I retire."

What’s next?

You won’t be fired for this strategy: “We’ll form a committee to devise an action plan to inventory current usage of cryptography to support future assessment of the steps needed to build a best-practices playbook for meeting the performance challenges of upgrading to post-quantum cryptography, with a target date after I retire.”

But what SSH did is a better strategy.