ETH zürich

Post-Quantum Cryptography

HAIC Talk
1/11/2019

Professor Kenny Paterson

Applied Cryptography Group

ETH Zürich

www.appliedcrypto.ethz.ch

@kennyog

Agenda

- Quantum Computing and Quantum Supremacy
- Thought Experiment: A World Without Public Key Cryptography
- Post-Quantum Cryptography
- The NIST PQC "Competition"
- Concluding Remarks

Quantum Computing – Basic Concepts

Basic tenet of quantum physics: superposition.

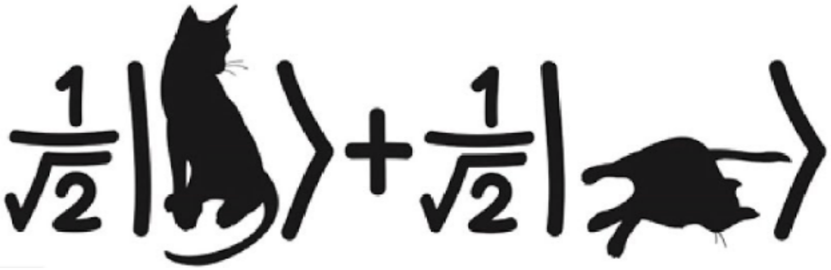

Qubit: basic unit of quantum computation, loosely a superposition of classical “0” and “1” bits.

Quantum gates: analogues of classical computing gates – AND, NOT, etc – acting on qubits.

Quantum computing: execution of a sequence of quantum gates on ”all possible classical states in parallel”.

Quantum Computing – Shor’s Algorithm

For *some* classical computing problems, we can arrange the initial state and the circuit so that something interesting results when the state collapses upon observation…

Shor’s algorithm (1994): finds the *period* of a classical function *f.*

- Leads to *efficient* quantum algorithms for solving integer factorisation and discrete logarithm problem (DLP).
- Number of quantum bits and circuit depth needed are polynomial in number of bits needed to specify the target input.
- cf. subexponential- and exponential-time classical algorithms for factoring and solving DLP.
- These ”hard problems” are the corner-stones of modern cryptography.

Quantum Computing – Shor's Algorithm

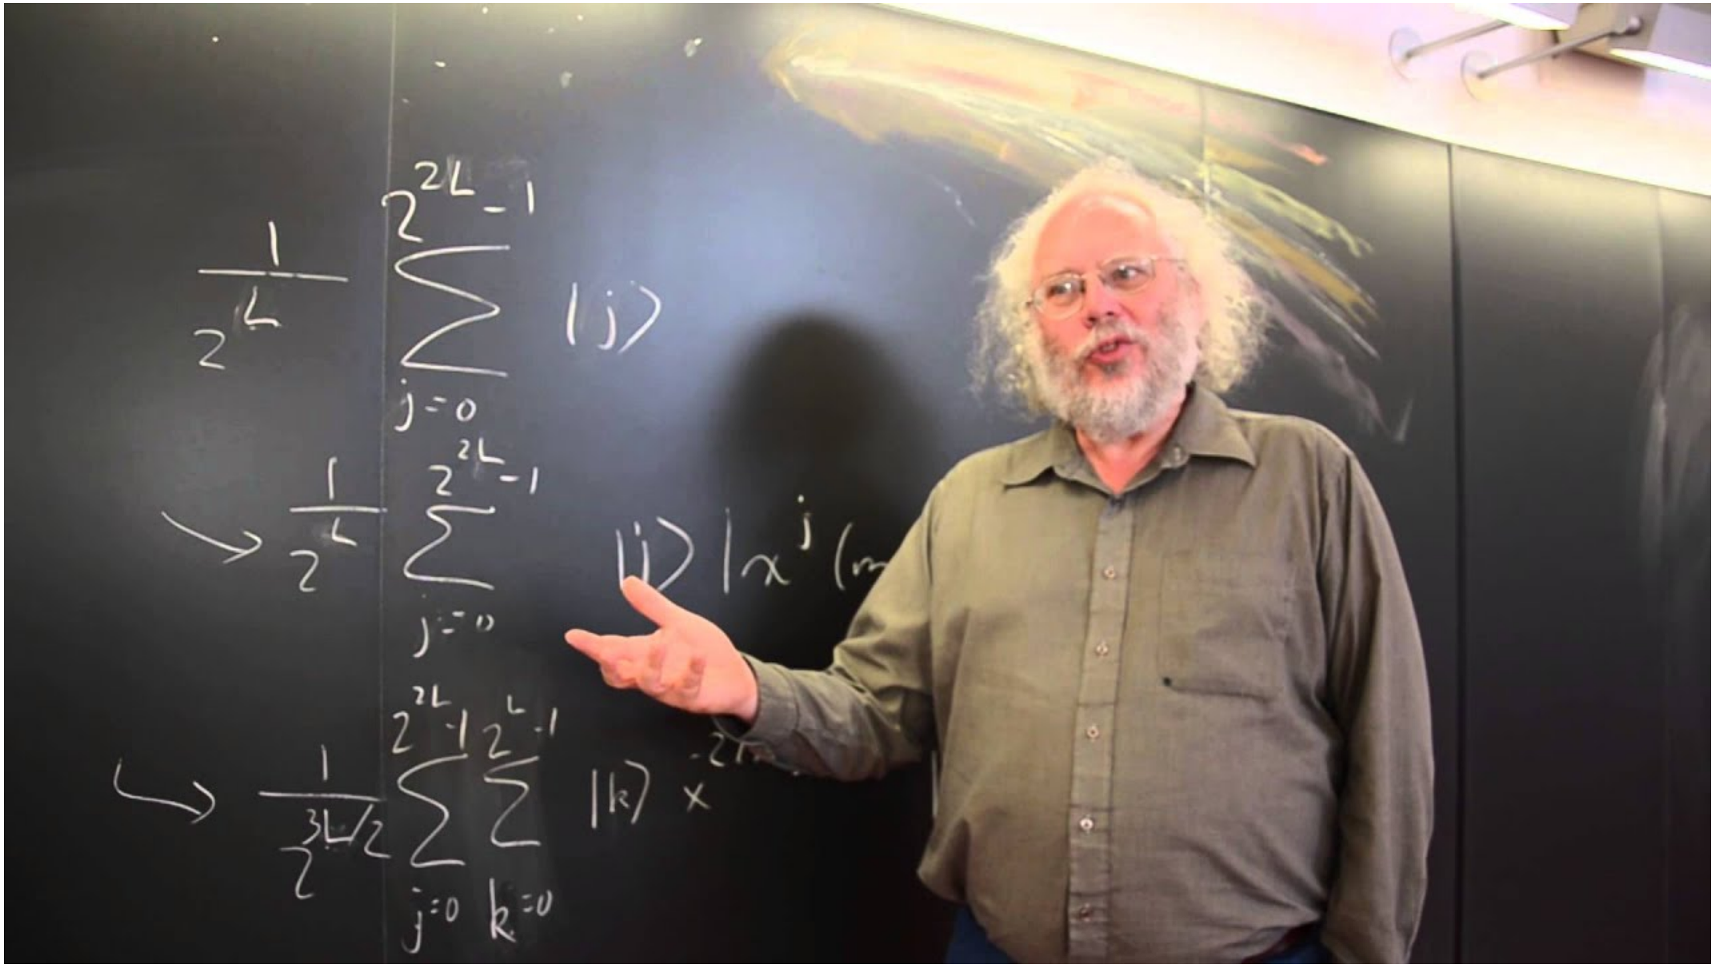
$\frac{1}{2^L}\sum_{j=0}^{2^{2L}-1}|j\rangle$
$\rightarrow \frac{1}{2^L}\sum_{j=0}^{2^{2L}-1}|j\rangle|x^j(m$
$\hookrightarrow \frac{1}{2^{3L/2}}\sum_{j=0}^{2^{2L}-1}\sum_{k=0}^{2^L-1}|k\rangle \times$

https://www.youtube.com/watch?v=hOlOY7NyMfs

Progress Towards Large-Scale Quantum Computing

http://en.wikipedia.org/wiki/Timeline_of_quantum_computing

Pre 1994: isolated contributions by Wiesner, Holevo, Bennett, etc.

1994: Shor's algorithm

1996: Grover's algorithm – quadratic speed up for search problems, applicable to exhaustive key search for symmetric algorithms.

1998: 2-qubit and 3-qubit NMR

2000: 5-qubit and 7-qubit NMR.

2001: The number 15 is factored!

2006: 12 qubits.

2007: 28 qubits.

2008: 128 qubits.

(D-Wave: quantum annealing machine)

D-Wave

Further Progress

http://en.wikipedia.org/wiki/Timeline_of_quantum_computing

2011: 14 qubits.

2012: The number 21 is factored!

2013 - 2017: ???

Late 2016 onwards: physicists switch focus to **quantum supremacy** as their success metric and **quantum systems simulation** as a killer application.

2017: D-Wave 2000Q, with 2000 qubits; IBM unveils 17-qubit machine, then a 50-bit machine.

2018: Google Bristlecone chip: 72 qubits; Intel Tangle Lake: 49 qubits.

2018: Quantum Supremacy? Not quite…

2019: Quantum Supremacy? Maybe…

Quantum Supremacy?

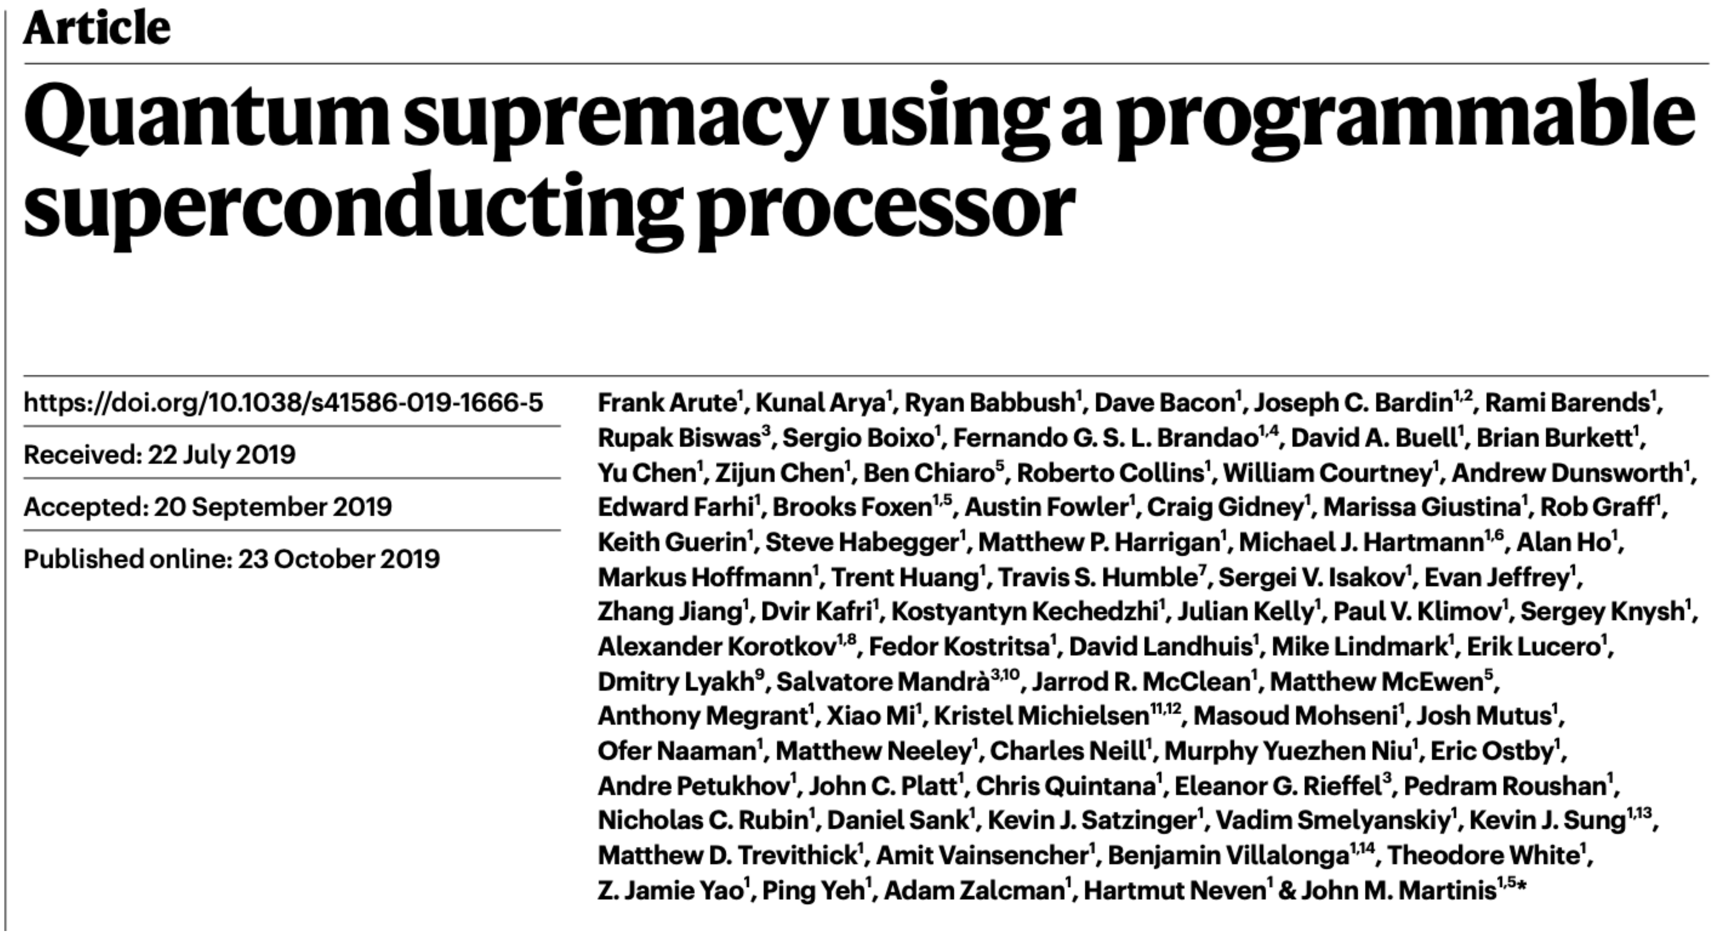

Article

Quantum supremacy using a programmable superconducting processor

https://doi.org/10.1038/s41586-019-1666-5

Received: 22 July 2019

Accepted: 20 September 2019

Published online: 23 October 2019

Frank Arute[1], Kunal Arya[1], Ryan Babbush[1], Dave Bacon[1], Joseph C. Bardin[1,2], Rami Barends[1], Rupak Biswas[3], Sergio Boixo[1], Fernando G. S. L. Brandao[1,4], David A. Buell[1], Brian Burkett[1], Yu Chen[1], Zijun Chen[1], Ben Chiaro[5], Roberto Collins[1], William Courtney[1], Andrew Dunsworth[1], Edward Farhi[1], Brooks Foxen[1,5], Austin Fowler[1], Craig Gidney[1], Marissa Giustina[1], Rob Graff[1], Keith Guerin[1], Steve Habegger[1], Matthew P. Harrigan[1], Michael J. Hartmann[1,6], Alan Ho[1], Markus Hoffmann[1], Trent Huang[1], Travis S. Humble[7], Sergei V. Isakov[1], Evan Jeffrey[1], Zhang Jiang[1], Dvir Kafri[1], Kostyantyn Kechedzhi[1], Julian Kelly[1], Paul V. Klimov[1], Sergey Knysh[1], Alexander Korotkov[1,8], Fedor Kostritsa[1], David Landhuis[1], Mike Lindmark[1], Erik Lucero[1], Dmitry Lyakh[9], Salvatore Mandrà[3,10], Jarrod R. McClean[1], Matthew McEwen[5], Anthony Megrant[1], Xiao Mi[1], Kristel Michielsen[11,12], Masoud Mohseni[1], Josh Mutus[1], Ofer Naaman[1], Matthew Neeley[1], Charles Neill[1], Murphy Yuezhen Niu[1], Eric Ostby[1], Andre Petukhov[1], John C. Platt[1], Chris Quintana[1], Eleanor G. Rieffel[3], Pedram Roushan[1], Nicholas C. Rubin[1], Daniel Sank[1], Kevin J. Satzinger[1], Vadim Smelyanskiy[1], Kevin J. Sung[1,13], Matthew D. Trevithick[1], Amit Vainsencher[1], Benjamin Villalonga[1,14], Theodore White[1], Z. Jamie Yao[1], Ping Yeh[1], Adam Zalcman[1], Hartmut Neven[1] & John M. Martinis[1,5]*

Quantum Supremacy?

- Transmuon qubits – nonlinear superconducting resonators at 5-7GHz.
- 53 qubit machine, operating at 20mK.
- Each qubit coupled to 4 neighbours in rectangular grid.
- Error rates below 1% for 2-qubit gate operations.
- Evaluate random quantum circuits with around 1100 single qubit and 400 two-qubit gates.
- Repeated circuit evaluation to reduce noise, entire computation runs in a few minutes.
- Expect to see "quantum interference" speckle pattern in output.
- Believed to be hard to simulate such a system classically: *Nature* paper estimates 50 trillion core-hours and 1 Petawatt-hour of energy on Google servers.

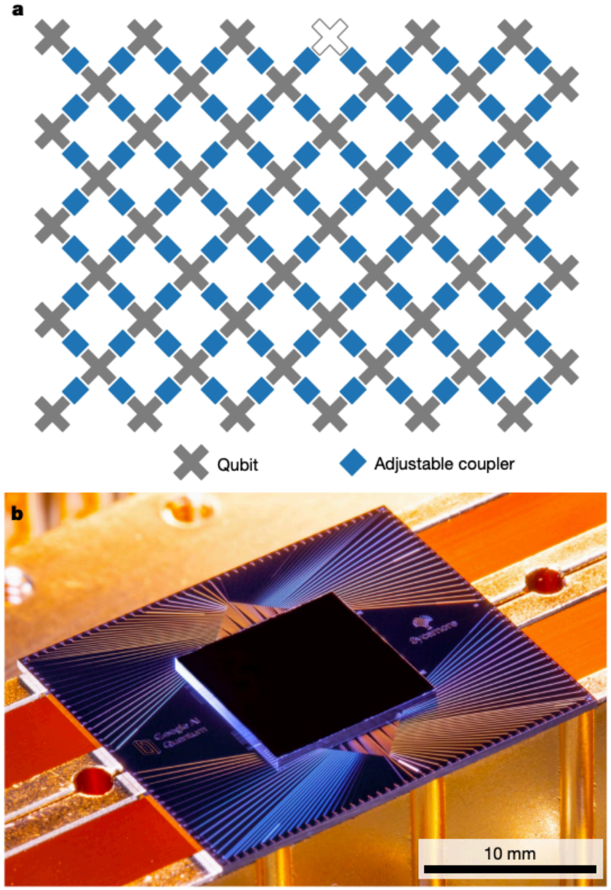

Fig. 1 | The Sycamore processor. a, Layout of processor, showing a rectangular array of 54 qubits (grey), each connected to its four nearest neighbours with couplers (blue). The inoperable qubit is outlined. **b**, Photograph of the Sycamore chip.

https://www.nature.com/articles/s41586-019-1666-5.pdf

Quantum Supremacy – Some Critique

- So what counts as (quantum) computation? If naturally arising quantum systems are hard to simulate classically, then didn’t Nature already achieve quantum supremacy?
- Nature paper claims 10 minutes on their quantum computer and 10,000 years on IBM summit supercomputer; IBM blogpost and paper then claims 2.5 days on IBM summit supercomputer by better use of disk storage. So has supremacy really been achieved?
- The computation is essentially useless, so what does this tell us? When do we get useful computation?
- *Quantum supremacy* is a horrible choice of words. Physicists need to up their game.

Quantum Computing’s Prospects

- The timeline for large-scale quantum computing is difficult to assess.
  - Quantum supremacy is **largely** a red herring in this regard: a race focussed on this goal does not tell us that much about what happens next.
  - Factoring slightly larger numbers is a **complete** red herring.
  - Smart people are working on it and have had a **lot** of research investment.
  - A scaling breakthrough might be imminent, but then again it might not.

Quantum Computing's Prospects

- On the other hand, maybe QC is a bit like fusion research?
  - Do-able in the lab, but very hard and expensive to scale.
  - Key issues are noisy gates and decoherence of quantum state.

…hardware improvements will probably follow a quantum-processor equivalent of Moore's law

– authors of *Nature* paper.

Quantum computing has been ten years away for about 40 years now and qualifies for the ironic title of 'Techno-Ponzi Scheme' – a good way to extract a research budget from a government, an R&D institute or a management.

David Manners, Electronics Weekly, December 2018

The Coming Crypt-Apocalypse?

- We don't know if there will be a QC scaling breakthrough or not.
- If one comes out of the blue, it could bring about the **Crypt-Apocalypse**.
  - Effect on symmetric key cryptography is small (**Grover's algorithm** gives theoretical square root speed up for exhaustive key search, so just double key-size, use AES-256).
  - But **Shor's algorithm** would break all public key crypto deployed on the Internet today – ECC and RSA.
  - Record now, break later attacks.
  - Replacing crypto at scale takes time: years to decades.

Ways Forward?

More usefully:

- Design new and improve existing cryptosystems that we believe resist quantum attack.
  - Lattice-based, code-based, non-linear systems of equations, isogeny-based…

- Consider a world without *any* public key cryptography?

A World Without Public Key Cryptography?

- Digital signature schemes from one-way functions, solving integrity and authentication problems.
    - Lamport signatures (1979) + hash trees.
    - Substantial research effort has gone into optimising constructions.
    - Not as efficient as, e.g. EC-DSA or RSA signatures, but just about usable.
    - IETF RFCs for standards.

- But we wouldn’t know how to do public key encryption, and we don’t know how to do Diffie-Hellman key exchange.

The “Open Systems” Challenge

- Prototypical application: e-commerce, protected by SSL/TLS.
- Requirements:
  - No pre-arranged trust relationships nor pre-established key between communicating parties.
  - Customers and merchants want end-to-end communications security and authentication guarantees.
- Currently achieved using PKC and digital certificates.
  - Either RSA encryption or Elliptic-curve Diffie-Hellman supported by signatures.

The “Open Systems” Challenge

Paterson’s Security Meta-Theorem^{TM} :

> **Any cryptographic problem can be solved by the introduction of sufficiently many trusted third parties.**

Applying the Meta-Theorem to Address the Crypt-Apocalypse

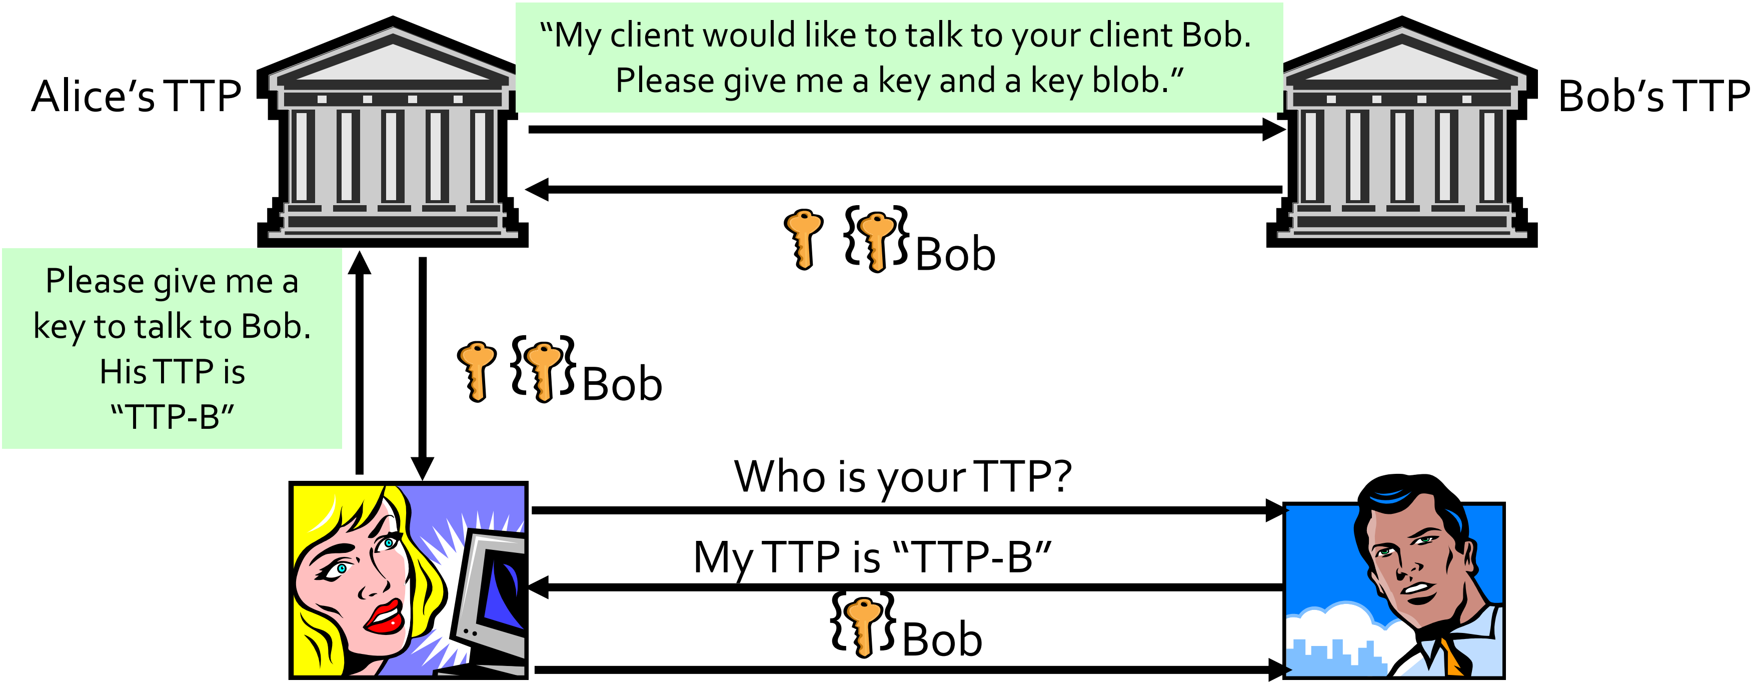

Alice's TTP

"My client would like to talk to your client Bob. Please give me a key and a key blob."

Bob's TTP

Bob

Please give me a key to talk to Bob. His TTP is "TTP-B"

Bob

Who is your TTP?

My TTP is "TTP-B"

Bob

Applying the Meta-Theorem to Address the Crypt-Apocalypse

- A low-tech 4-party protocol to establish keys for authentication and secure communications.
- Alice and Bob contract with individual TTPs for security services.
- Offers significantly reduced privacy for users, and requires very high trust in TTPs compared to current PKC-based solutions.
- Degree of trust needed can be reduced and privacy improved by splitting TTPs and using threshold cryptography.

<End of thought experiment>

Post Quantum Cryptography (PQC)

Conventional public-key cryptosystems that resist known quantum algorithms.

- AKA **quantum-safe** or **quantum-immune** cryptography.
- (Quantum Key Distribution (QKD) is something else entirely.)
- Main candidates are lattice-based, code-based, non-linear systems of equations, elliptic curve isogenies.
- Public key encryption, digital signature schemes, key exchange, and more: replicating classical functionality.
- Until recently, a relatively niche research area within cryptography.
- NIST PQC standardisation process has attracted many new researchers and caused others to redouble their efforts.

PQC – Quantum and Classical Security

- **Some or even all available PQC algorithms may be vulnerable to further advances in quantum algorithms.**
  - cf. Soliloquy paper (GCHQ/NCSC) – certain ways of using ideal lattices are vulnerable to quantum attack.
  - *But also: where are all the new quantum algorithms?*
- **Even conventional security is not yet well understood in all cases.**
  - For lattice-based systems: many ways to tackle underlying hard problems, clear picture still emerging.
  - Many multivariate schemes have been proposed and then broken.
  - Security for vanilla code-based cryptography is by now fairly stable.
  - Hash-based signatures schemes are particularly mature.

PQC Characteristics

- Current PQC schemes offer a different set of trade-offs compared to classical schemes.
  - Public key size vs ciphertext/signature size vs computation time.
  - Can have **faster** cryptographic operations – just matrix multiplication plus noise for lattice based schemes.
  - Can have **short** public keys and compact ciphertexts – isogeny-based crypto.
  - But nothing as “nice” as current situation with ECC.
- Parameter selection is more complex than for RSA/ECC.
  - For example, lattice-based schemes have several tuneable parameters and structural choices – base ring, dimension, noise distribution,…
  - May need to adapt parameters to respond to improved attacks.
  - But RSA/ECC parameter selection was once hard too!

PQC and NIST

NIST: US National Institute for Standards in Technology.

- Has run cryptographic competitions in the past: DES (NBS), AES, SHA-3.
- Publishes many fundamental cryptographic standards based on these and other algorithms.
- These standards become *de facto* global standards for cryptography (notwithstanding algorithms from China, Russia, etc).
- They are often adopted by other standards bodies, e.g. IETF and ISO.

PQC and NIST

NIST process, 2016 – 2023(ish) for standardising post-quantum public key algorithms.

- http://csrc.nist.gov/groups/ST/post-quantum-crypto/
- **Evaluation criteria**: security, cost, flexibility/simplicity/adoptability.
- **Process (5-7 years)**:

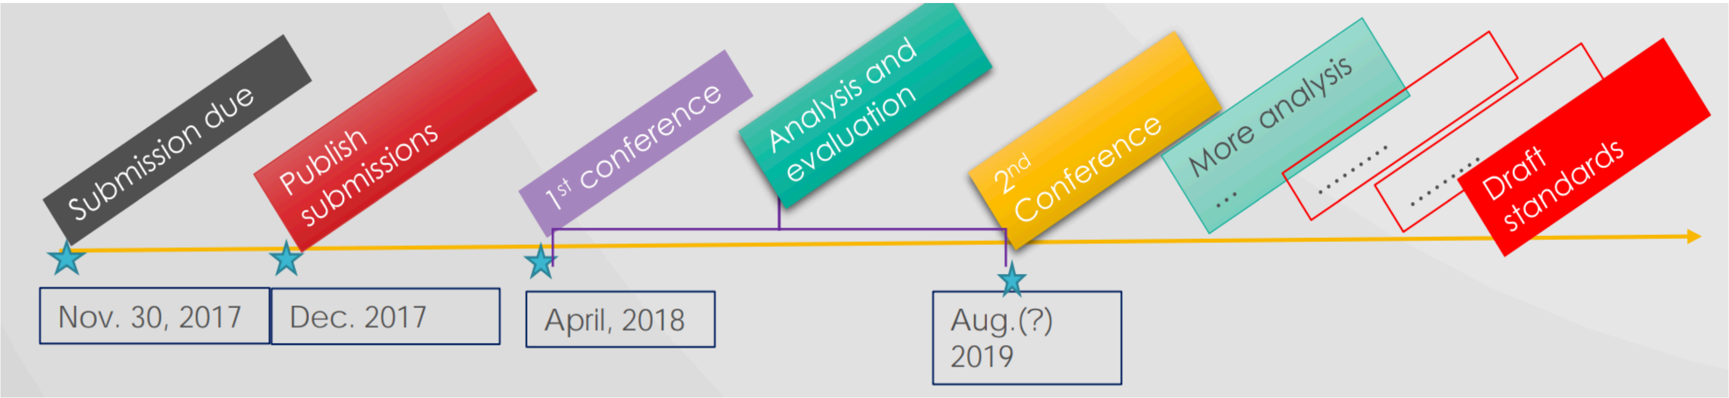

https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/asiacrypt-2017-moody-pqc.pdf

PQC and NIST

NIST asked for 3 types of algorithm:

- Signatures.
- Public key encryption for symmetric key transport.
- “Key-establishment (KEMs)” – “Schemes like Diffie-Hellman key exchange (see SP 800-56A)”.

NIST have said they will not pick a single winner in each category – more likely is a portfolio of choices, to accommodate different application scenarios.

PQC and NIST

NIST identified 5 security levels, and provided guidance on security proofs and classical vs quantum adversaries.

Computational resources should be measured using a variety of metrics:

- Number of classical elementary operations, quantum circuit size, etc…
- Consider realistic limitations on circuit depth (e.g. 2^{40} to 2^{80} logical gates).
- May also consider expected relative cost of quantum and classical gates.
- Estimates from submitters considered to be preliminary.
- Submitters need not provide parameters for all levels.

NIST Submissions – 11/2017

- NIST received 82 submissions from 25 countries, total page count approx 3200.
- 69 submissions deemed to be “complete and proper”
- 5 teams withdrew; several more quickly attacked.

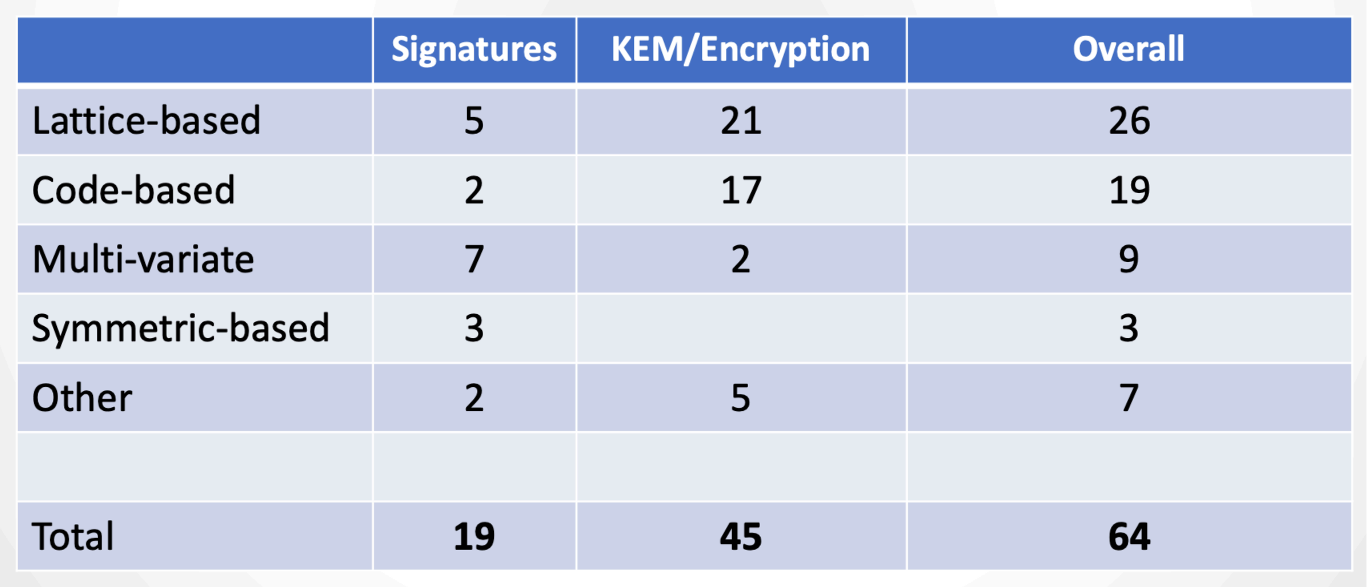

	Signatures	KEM/Encryption	Overall
Lattice-based	5	21	26
Code-based	2	17	19
Multi-variate	7	2	9
Symmetric-based	3		3
Other	2	5	7
Total	**19**	**45**	**64**

https://csrc.nist.gov/CSRC/media/Presentations/Round-2-of-the-NIST-PQC-Competition-What-was-NIST/images-media/pqcrypto-may2019-moody.pdf

Skin in the Game

- Full disclosure: I was a co-submitter on two proposals: LIMA and NTS-KEM.
- LIMA did not make it past the first round, but LIMA team members were invited to join NewHope proposal in second round.
- NTS-KEM survives as a second round candidate:
  - A simple, conservative coding-based KEM based on ideas from McEliece and Niederreiter.
  - Fast algorithms, fairly compact ciphertexts.
  - Tight security proof.
  - Only drawback is large public key size.

First NIST Workshop – 4/2018

- Took place in April 2018 in Florida.
- All submitters (not withdrawn) were invited to present briefly.
- One invited presentation compared attack cost models for lattice-based schemes (based on work of Albrecht *et al*.).
  - This highlighted lack of community agreement on how to measure security and attack costs for lattice-based schemes.
  - Much subsequent discussion on the NIST PQC mailing list.

Second Round Candidates – 1/2019

- January 20th 2019: NIST announced 26 candidates for the second round of analysis.
- 17 public key encryption schemes/KEMs:

  BIKE, Classic McEliece, CRYSTALS-KYBER, FrodoKEM, HQC, LAC, LEDAcrypt (merger of LEDAkem/LEDApkc), NewHope, NTRU (merger of NTRUEncrypt/NTRU-HRSS-KEM), NTRU Prime, NTS-KEM, ROLLO (merger of LAKE/LOCKER/Ouroboros-R), Round5 (merger of Hila5/Round2), RQC, SABER, SIKE, Three Bears
- 9 signature schemes:

  CRYSTALS-DILITHIUM, FALCON, GeMSS, LUOV, MQDSS, Picnic, qTESLA, Rainbow, SPHINCS+
- Still representatives from all major design paradigms.
- Proposers invited to submit tweaked designs, additional analysis, etc, by March 15th 2019.
- Some rationale for selection given by NIST: https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8240.pdf

Second Round Encryption Candidates – KPIs

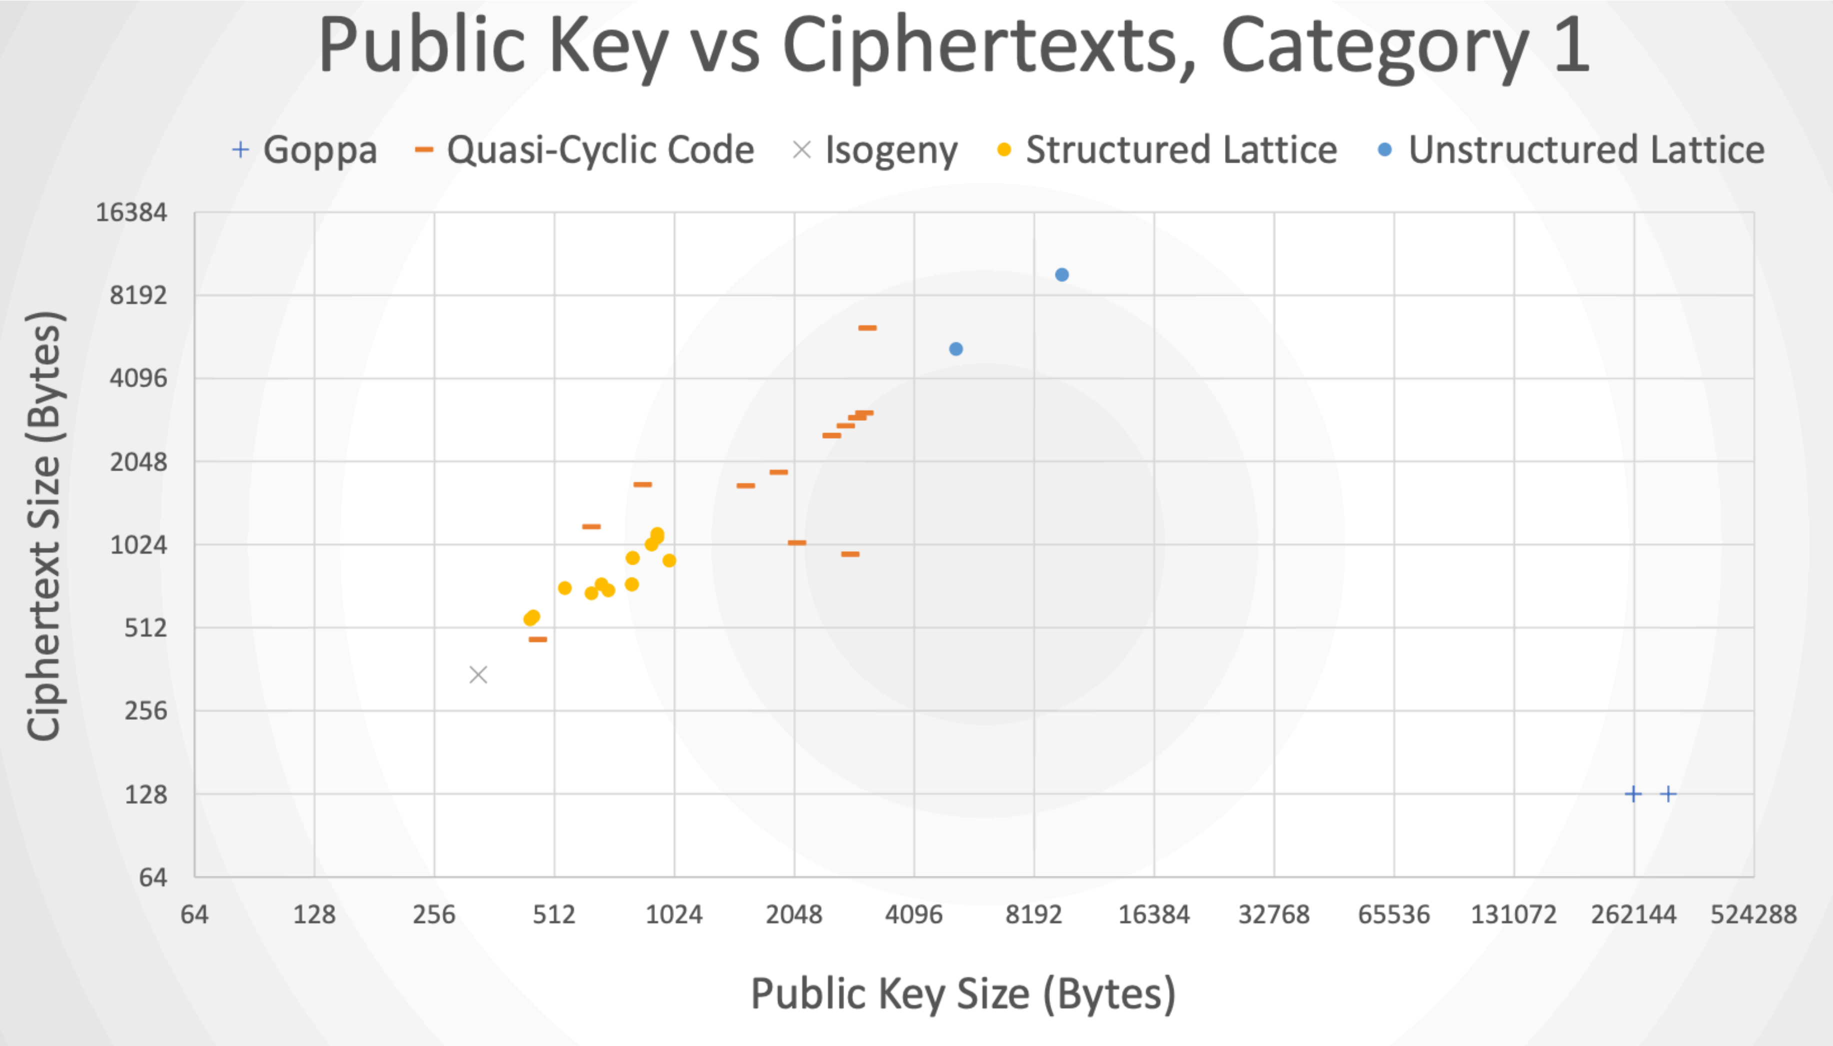

https://csrc.nist.gov/CSRC/media/Presentations/Round-2-of-the-NIST-PQC-Competition-What-was-NIST/images-media/pqcrypto-may2019-moody.pdf

Second Round Signature Candidates – KPIs

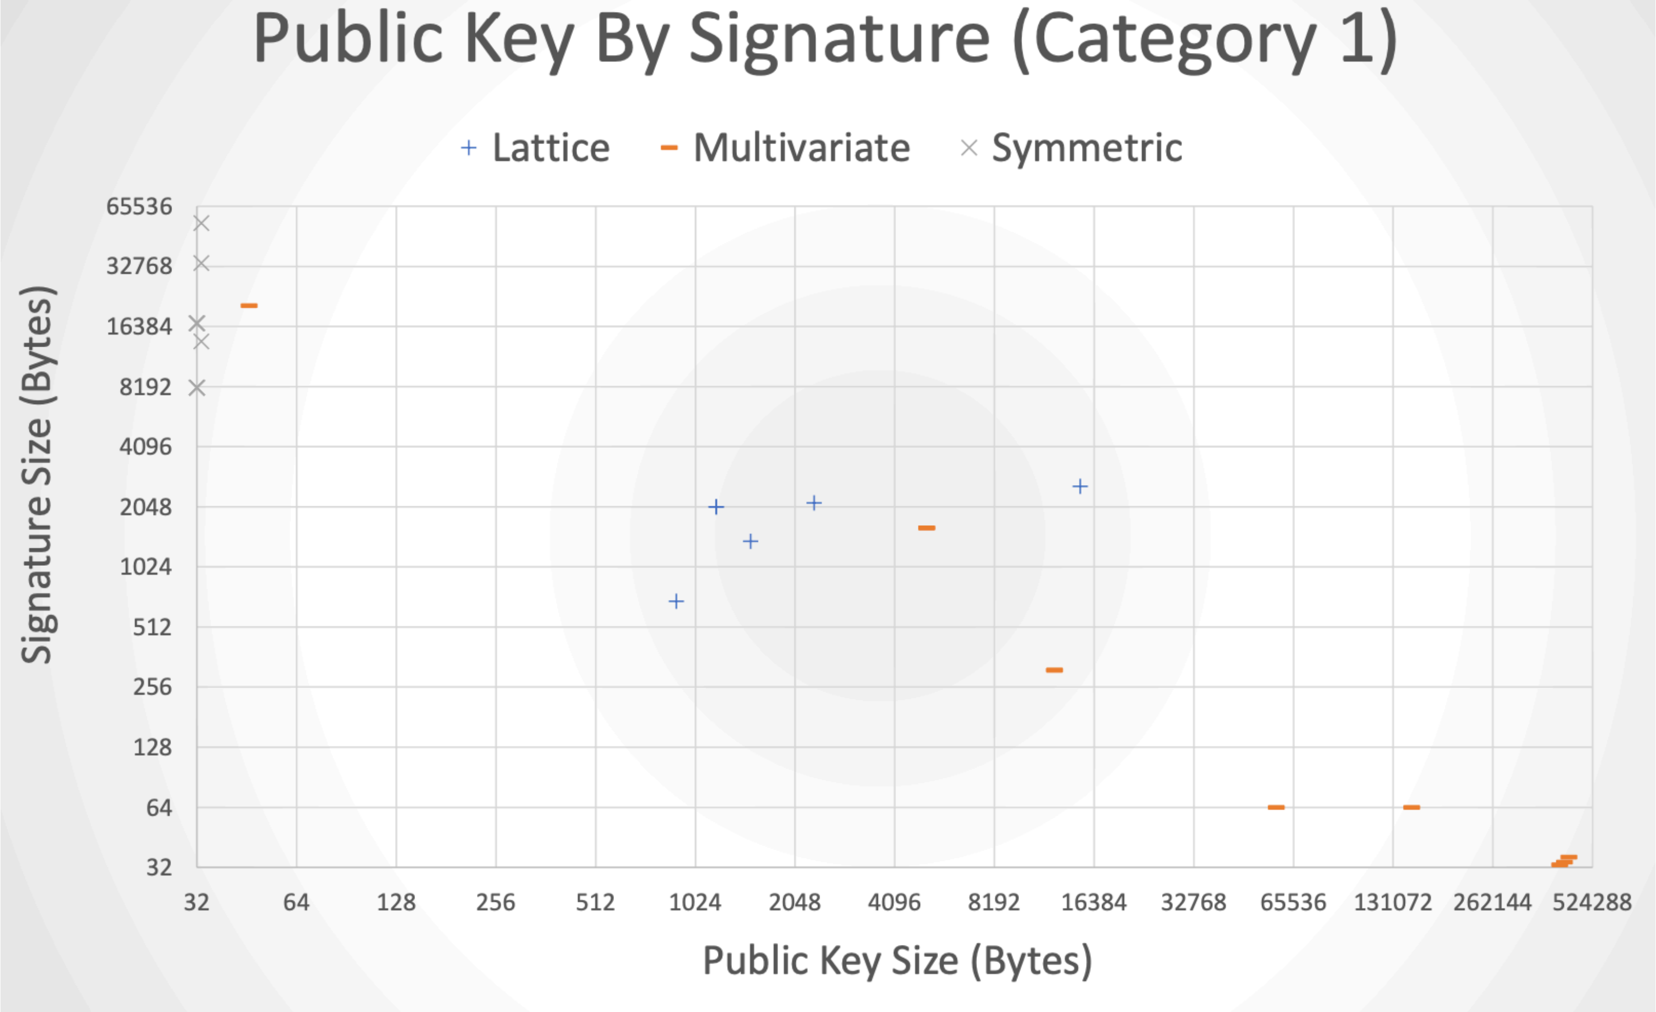

https://csrc.nist.gov/CSRC/media/Presentations/Round-2-of-the-NIST-PQC-Competition-What-was-NIST/images-media/pqcrypto-may2019-moody.pdf

Second NIST Workshop – 8/2019

- August 2019, Santa Barbara, CA.
- All 26 second round candidates were invited to present.
- Plus several submitted talks on benchmarking, library integration, deployment experiments,…
- NIST ran a survey asking for feedback on the process.

"Ensure civility on the mailing list. Certain vocal and at times outright impolite personalities dominate the mailing list, causing others to hesitate to contribution their work or questions…"

The NIST Process From Here

- NIST have indicated that round 2 will last 12-18 months.
- NIST is encouraging further emphasis on hardware+software performance.
- Also still open to mergers between proposals.
- There may be a third round of analysis – reflecting research oriented nature of process.
- Rapid developments in isogeny-based crypto pose a particular challenge.
- NIST still aiming for draft standards in 2022.

NIST Process: Research Community Foci

- Improving quality of implementations, including resistance to side-channel attacks.
- Carefully vetting of security proofs.
- Development of new proof techniques in “quantum ROM” setting.
- Gaining deeper understanding of the underlying hard problems, especially for lattice-based and isogeny-based proposals.
- Understanding impact of different designs on performance of main Internet protocols (TLS, SSH, etc).
- Developing hybrid protocols combining classical and PQ primitives.

PQC: from Research to Deployment

- NIST PQC process is benefitting from deployment experiments.
  - Experimental deployment of NewHope in combination with ECC by Google in TLS in 2016.
    - https://www.imperialviolet.org/2016/11/28/cecpq1.html
  - On-going experimental deployment of HRSS + ECC / SIKE + ECC by Google/Cloudflare in TLS.
    - https://blog.cloudflare.com/the-tls-post-quantum-experiment/

CECPQ2 = HRSS + X25519

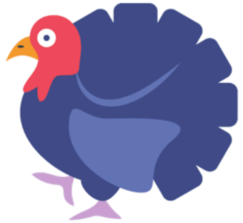

CECPQ2b = SIKE + X25519

Concluding Remarks – 1

- The Crypt-Apocalypse might be coming… or it might not be.
- Symmetric solutions will take us only so far, and would require radical changes to current security infrastructures.
- Post-quantum public key cryptography represents a more attractive route to solving the problem.

Concluding Remarks – 2

- We can hope that the NIST process will proceed in an orderly fashion and produce a sensible and conservative portfolio of options in a reasonable timeframe.
- It will take years beyond NIST’s planned decision point for the selected algorithms to become widely used.
- We can expect lots of deployment pain along the way: computing and bandwidth overheads, integration challenges, unanticipated glitches, immature code-bases,…

Thanks

Questions?

But What About Quantum Cryptography and QKD?

- Quantum Key Distribution (QKD) promises **unconditional** security.
  - "Security based only on the correctness of the laws of quantum physics".

- Often contrasted with security offered by currently deployed public key cryptography (PKC).
  - PKC is vulnerable to quantum computers.
  - PKC is vulnerable to algorithmic advances in conventional algorithms for factoring, discrete logs, etc.

QKD Limitations

Four reasons for slow rate of commercial uptake:

1. QKD has limits on rate and range in the absence of quantum repeaters.
2. Security in theory does not equal security in practice: quantum hacking.
3. QKD requires an initial shared key to be able to produce longer keys.
4. QKD does not offer significant **practical** security advantages over what we can currently do at low-cost with conventional techniques under the same starting assumption of a shared initial key.

QKD or PQC? The NCSC/GCHQ View

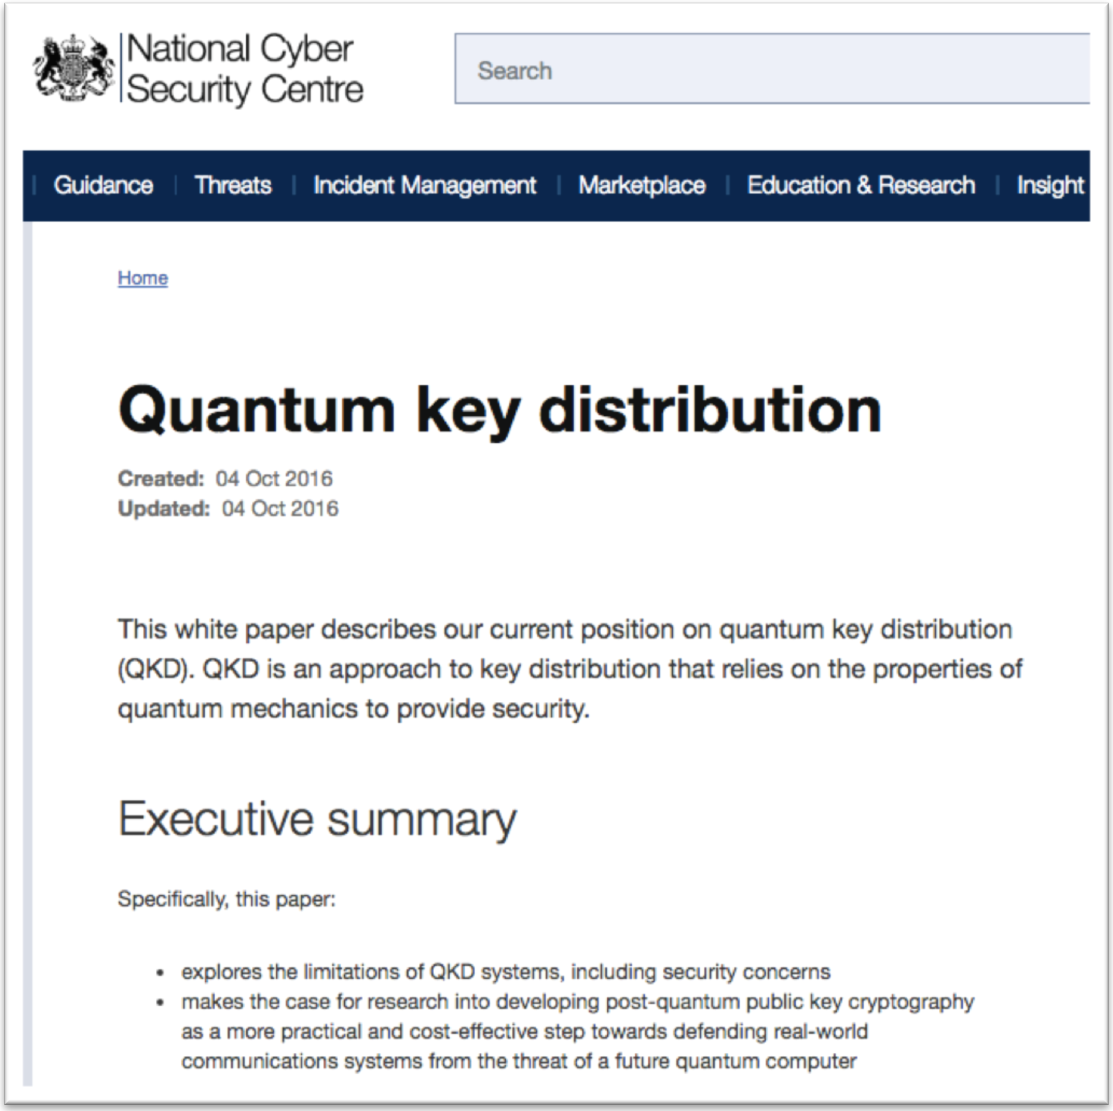

National Cyber Security Centre
Search
Guidance | Threats | Incident Management | Marketplace | Education & Research | Insight
Home

Quantum key distribution

Created: 04 Oct 2016
Updated: 04 Oct 2016

This white paper describes our current position on quantum key distribution (QKD). QKD is an approach to key distribution that relies on the properties of quantum mechanics to provide security.

Executive summary

Specifically, this paper:

- explores the limitations of QKD systems, including security concerns
- makes the case for research into developing post-quantum public key cryptography as a more practical and cost-effective step towards defending real-world communications systems from the threat of a future quantum computer