Crypto News

Compiled by Dhananjoy Dey, Indian Institute of Information Technology,
Lucknow, U. P. - 226 002, India, ddey@iiitl.ac.in

October 01, 2024

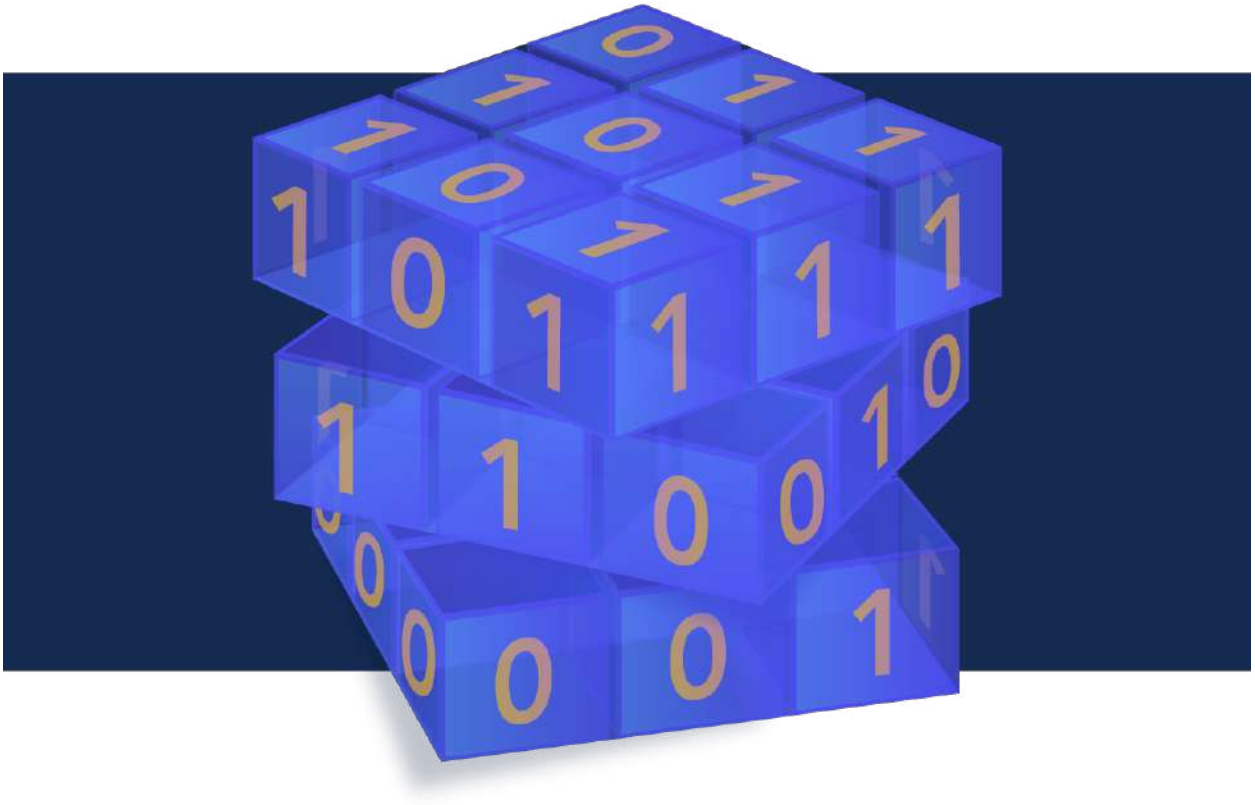

TABLE OF CONTENTS

Editorial

Dear Readers,

Here is your monthly newsletter, bringing you the best and newest in quantum and cryptography, courtesy of Dhananjoy.

To me, it is rather amazing that there is always so much new and interesting stuff every month. Selecting a few topics only for my Editorial is the hard part… But since this is my task, here we go…

If you have time for one read only, I would definitely recommend article 14, from Scott Aaronson, for a very lucid analysis on the status of quantum computing. Scott is providing us with a first-hand account of the current developments in quantum computing, and why he believes that it is finally becoming real. And this is a good introduction to the next topic.

With the final release of the PQC standards by NIST, we have entered into a new phase. Now that the new algorithms are chosen, we have to start implementing them. Many articles this month are devoted to the need to start the transition to post-quantum. See for example: 1, 6, 7, 9, 12, and 15. No more waiting, action!

Finally, as presented in 13 and 16, it seems that fully homomorphic encryption is also becoming real. This will have major consequences for the Cloud. The idea that you can keep your data encrypted in the Cloud and use it as required is truly appealing.

Of course, my choices this month are restricted and subjective. Many other articles are also worth a peak.

Have a great read!

The Crypto News editorial is authored by the Chair of the Quantum-Safe Security Working Group (QSS WG) of the Cloud Security Alliance (CSA), , Bruno Huttner, Director of Strategic Quantum Initiatives at ID Quantique SA and it is compiled by Dhananjoy Dey.

Disclaimer. The QSS WG does not express an opinion on the validity of the ideas and the claims presented in the articles in this newsletter.

1.Begin Transitioning to Post-Quantum Cryptography Now

by Mark Horvath
https://www.gartner.com/en/articles/post-quantum-cryptography

Post-quantum cryptography provides the next generation in data protection

Asymmetric encryption is in almost all software, billions of devices worldwide and most of the communications over the internet. Yet by 2029, advances in quantum computing will make asymmetric cryptography unsafe and by 2034 fully breakable. “Harvest-now, decrypt-later” attacks may already exist.

To resist attacks from both classical and quantum computers, organizations must transition to post-quantum cryptography (PQC). But that’s hardly a simple switch. It will require more work than preparing for Y2K, and failure could have dangerous consequences. Further, many organizations haven’t yet planned or budgeted for this shift.

Make the post-quantum cryptography transition an urgent priority

Governments are already issuing mandates and legal frameworks for organizations to put in place a post-quantum cryptography strategy. Start developing yours now.

Address the challenges of adopting post-quantum cryptography

Hurdles to overcome include:

- **Lack of easy replacement options.** No drop-in alternatives exist for current cryptographic algorithms. This creates the need for discovery, categorization and reimplementation.
- **Varied performance requirements.** New algorithms have different performance characteristics than asymmetric ones. Key and ciphertext sizes are larger, for example, and encryption and decryption times are longer. This may impact performance and require restating or rewriting of current applications.
- **Lack of organizational knowledge.** Few organizations know how their cryptography works, where keys and algorithms are used, or how secrets are stored and managed.
- **Lack of vendor preparedness.** Don’t assume that your vendors are equipped to handle post-quantum cryptography. Most are unprepared to upgrade and may not recognize that they need to unless you push them.

Enact clear policies to enable post-quantum cryptography adoption

To address these challenges and smooth the transition to new algorithms, start by developing policies on algorithm substitution, data retention and the mechanics of swapping or modifying your existing use of cryptography. A policy-based program will reduce confusion and arbitrary choices, and increase manageability.

From there, build a metadata database of all in-use cryptographic algorithms. This enables your organization to scope the impact of new cryptography, determine the risk to specific applications, reduce existing cryptographic security debt and reprioritize incident response plans. Create a life cycle policy to reflect the risks to asymmetric keys.

Ask the vendors you identify in your inventory about their plans for moving to post-quantum cryptography, the overall roadmap for implementation and the potential impact of upgrades to existing systems.

Finally, implement crypto-agile application development. Vet and test new post-quantum cryptographic algorithms to understand their characteristics, uses and performance. Upgrade or replace hardware where necessary.

2.New Quantum Error Correction Method

by Danielle Ellis
https://www.azoquantum.com/News.aspx?newsID=10513

Hayato Goto of the RIKEN Center for Quantum Computing in Japan has presented a novel quantum error correction method utilizing "many-hypercube codes" in a study that was published in Science Advances.

This method’s elegant geometry could facilitate incredibly effective error corrections and pave the way for highly parallel techniques that will enable fault-tolerant quantum computing—the next phase of quantum computing evolution.

> *Thanks to recent experimental progress, there is now great hope that we will be able to build fault-tolerant quantum computers, meaning quantum computers that can correct errors and surpass the power of conventional computers on certain tasks. To achieve this, however, it is important to develop efficient quantum error correction*.
>
> Hayato Goto, RIKEN Center for Quantum Computing

Over the past few decades, scientists have proposed a wide range of error correction techniques. The standard method for quantum error correction involves encoding a single logical qubit, which is similar to a bit in a classical computer, onto numerous entangled physical qubits.

A decoder is then used to extract the logical qubits from the physical ones. This method's scalability is hampered by the enormous increase in the number of physical qubits needed, which results in significant resource overheads.

High-rate quantum codes, like quantum low-density parity-check codes, have been explored as a solution to this issue. However, this method is less efficient in terms of time because it requires the logical gates—which enable calculations—to be set up quite sequentially rather than fully parallel.

To address this, Goto suggested employing a method he refers to as "many-hypercube codes." This method is specifically known by the complicated term "high-rate concatenated quantum codes." What makes it novel is that the logical qubits can be mathematically visualized as forming a "hypercube," a type of shape that includes squares, cubes, and higher-order shapes like tesseracts. Given that most high-rate quantum codes have complex structures, the code's exquisite mathematical and geometric structure is astounding.

Goto highlights that for the new codes to produce better performance, he had to create a brand-new, specialized decoder that could comprehend the output from the actual qubits. High performance is possible with this novel method because it is based on level-by-level minimum distance decoding.

Goto named the system “high-performance fault tolerant computing” about "high-performance computing,” which is used for massively parallel computing. This is because, in contrast to other similar methods, it also permits the placement of logical gates in parallel rather than in series, making the system analogous to parallel processing in classical computers.

The labor was fruitful. Goto states that the codes appear to be the highest in the world when it comes to fault-tolerant quantum computing, as they achieve an encoding rate of up to 30%, which is a measure of the ratio between logical and physical qubits. Furthermore, despite the high rate, the performance is on par with traditional low-rate codes.

Goto concluded, “In practice, this code could be implementing with physical qubit systems such as laser-trapped neutral-atom qubits.”

3.IBM opens its quantum-computing stack to third parties

by John Timmer
https://arstechnica.com/science/2024/09/ibm-opens-its-quantum-computing-stack-to-third-parties/

As we described earlier this year, operating a quantum computer will require a significant investment in classical computing resources, given the amount of measurements and control operations that need to be executed and interpreted. That means that operating a quantum computer will also require a software stack to control and interpret the flow of information from the quantum side.

But software also gets involved well before anything gets executed. While it's possible to execute algorithms on quantum hardware by defining the full set of commands sent to the hardware, most users are going to want to focus on algorithm development, rather than the details of controlling any single piece of quantum hardware. "If everyone's got to get down and know what the noise is, [use] performance management tools, they've got to know how to compile a quantum circuit through hardware, you've got to become an expert in too much to be able to do the algorithm discovery," said IBM's Jay Gambetta. So, part of the software stack that companies are developing to control their quantum hardware includes software that converts abstract representations of quantum algorithms into the series of commands needed to execute them.

IBM's version of this software is called Qiskit (although it was made open source and has since been adopted by other companies). Recently, IBM made a couple of announcements regarding Qiskit, both benchmarking it in comparison to other software stacks and opening it up to third-party modules. We'll take a look at what software stacks do before getting into the details of what's new.

What’s the software stack do?

It's tempting to view IBM's Qiskit as the equivalent of a compiler. And at the most basic level, that's a reasonable analogy, in that it takes algorithms defined by humans and converts them to things that can be executed by hardware. But there are significant differences in the details. A compiler for a classical computer produces code that the computer's processor converts to internal instructions that are used to configure the processor hardware and execute operations.

Even when using what's termed "machine language," programmers don't directly control the hardware; programmers have no control over where on the hardware things are executed (ie, which processor or execution unit within that processor), or even the order instructions are executed in.

Things are very different for quantum computers, at least at present. For starters, everything that happens on the processor is controlled by external hardware, which typically act by generating a series of laser or microwave pulses. So, software like IBM's Qiskit or Microsoft's Q# act by converting the code they're given into commands that are sent to hardware that's external to the processor.

These "compilers" must also keep track of exactly which part of the processor things are happening on. Quantum computers act by performing specific operations (called gates) on individual or pairs of qubits; to do that, you have to know exactly which qubit you're addressing. And, for things like superconducting qubits, where there can be device-to-device variations, which hardware qubits you end up using can have a significant effect on the outcome of the calculations.

As a result, most things like Qiskit provide the option of directly addressing the hardware. If a programmer chooses not to, however, the software can transform generic instructions into a precise series of actions that will execute whatever algorithm has been encoded. That involves the software stack making choices about which physical qubits to use, what gates and measurements to execute, and what order to execute them in.

The role of the software stack, however, is likely to expand considerably over the next few years. A number of companies are experimenting with hardware qubit designs that can flag when one type of common error occurs, and there has been progress with developing logical qubits that enable error correction. Ultimately, any company providing access to quantum computers will want to modify its software stack so that these features are enabled without requiring effort on the part of the people designing the algorithms.

Benchmarking the stack

Gambetta said that one of the things that IBM did recently is rewrite the Qiskit stack in Rust, rather than its original Python. The expectation was that compiled code would lead to better performance overall, which likely motivated the decision to develop a cross-platform benchmarking suite. The suite involves generating quantum circuits from abstract representations of algorithms, and manipulating and optimizing the ensuing circuit. The resulting circuits were then examined for how many gate operations they required (fewer is generally better) and how quickly they could be transformed.

Unsurprisingly, IBM feels that Qiskit did extremely well on its tests, but the details provide a window into where things stand in the world of quantum-computing software stacks, as seven different ones were included in the tests. One detail is that building quantum circuits isn't a major time sink; most tasks completed in around a second or less, and even the worst result for one software package took only a bit more than a minute. That said, a number of the software packages failed to complete a subset of the tests.

There was also clearly a tradeoff. For example, a toolkit called Staq tended to complete circuit building faster than Qiskit, but the resulting algorithm would require more operations to complete on quantum hardware. Meanwhile, software called Tket took longer than Qiskit, but tended to produce quantum algorithms that required fewer operations to execute, especially when targeting non-IBM hardware. Fewer operations can be critical, as each operation is an opportunity for an error to ruin a calculation.

This is all valuable because many of these toolkits are open source, so it should be possible to use the code from one to inform the performance of some of the others.

Beyond Qiskit

In a somewhat related announcement, IBM is also opening up its software stack so that users don't have to be entirely dependent on its tools. Right now, the company is supporting six third-party Qiskit functions that break down into two categories.

The first can be used as stand-alone applications and are focused on providing solutions to problems for users who have no expertise programming quantum computers. One calculates the ground-state energy of molecules, and the second performs optimizations.

But the remainder are focused on letting users get more out of existing quantum hardware, which tends to be error prone. But some errors occur more often than others. These errors can be due to specific quirks of individual hardware qubits or simply because some specific operations are more error prone than others. These can be handled in two different ways. One is to design the circuit being executed to avoid the situations that are most likely to produce an error. The second is to examine the final state of the algorithm to assess whether errors likely occurred and adjust to compensate for any. And third parties are providing software that can handle both of these.

One of those third parties is Q-CTRL, and we talked to its CEO, Michael Biercuk. "We build software that is really focused on everything from the lowest level of hardware manipulation, something that we call quantum firmware, up through compilation and strategies that help users map their problem onto what has to be executed on hardware," he told Ars. (Q-CTRL is also providing the optimization tool that's part of this Qiskit update.)

"We're focused on suppressing errors everywhere that they can occur inside the processor," he continued. "That means the individual gate or logic operations, but it also means the execution of the circuit. There are some errors that only occur in the whole execution of a circuit as opposed to manipulating an individual quantum device." Biercuk said Q-CTRL's techniques are hardware agnostic and have been demonstrated on machines that use very different types of qubits, like trapped ions. While the sources of error on the different hardware may be distinct, the manifestations of those problems are often quite similar, making it easier for Q-CTRL's approach to work around the problems.

Those work-arounds include things like altering the properties of the microwave pulses that perform operations on IBM's hardware, and replacing the portion of Qiskit that converts an algorithm to a series of gate operations. The software will also perform operations that suppress errors that can occur when qubits are left idle during the circuit execution.

As a result of all these differences, he claimed that using Q-CTRL's software allows the execution of more complex algorithms than are possible via Qiskit's default compilation and execution. "We've shown, for instance, optimization with all 156 qubits on [an IBM] system, and importantly—I want to emphasize this word—successful optimization," Biercuk told Ars. "What it means is you run it and you get the right answer, as opposed to I ran it and I kind of got close."

Compensating for errors

Asif Sinay, the CEO of QEDMA, says one of the things his software does is quite similar. "When someone sends us an algorithm, we are just looking at the specific gates that he's going to run, and then we analyze the device as a function of what he sent to us," Sinay said. "Sometimes there are some tricks that you can compress the algorithm, you can do some changes inside your algorithm."

But QEDMA's software also handles errors after they occur. If a user is willing to run the algorithm multiple times, it's possible to profile the errors that occur during its execution, which will often differ between runs. "If you don't know what the noise looks like, you cannot fix the errors," Sinay said. With a sufficient profile of the noise, you can recognize when it's likely to have occurred and can compensate for that. (IBM has also demonstrated a form of this, although it's computationally expensive.) Like Biercuk, Sinay was confident that this approach would start producing useful calculations on quantum computers before we are able to perform those calculations with error-corrected quantum computers. "The middleware is the bridge to close the gap between the hardware and the application, because without the middleware, without reducing these errors, these devices will not be useful," Sinay told Ars.

As for IBM, Gambetta said that he's excited about what might be possible when algorithm developers no longer have to worry as much about reducing noise enough to get a calculation to run to completion. So far, quantum algorithms have been defined through rigorous mathematical proofs, rather than by testing different approaches on actual hardware. As the hardware and software stacks continue to mature, that may start to change. "Most of what people have been doing is working out how to get rid of the effects of noise on our quantum circuits," Gambetta said. "And so the abstraction of the circuit function is the first time algorithm developers are not going to need to worry."

4.Australian Organisations Targeted by Phishing Attacks Disguised as Atlassian

by Ben Abbott
https://www.techrepublic.com/article/phishing-attacks-australia-atlassian/

Enterprises across Australia and the APAC region have been warned that cyber criminals are exploiting popular platforms like Atlassian to launch more convincing phishing attacks on law firms and other corporations. These attacks aim to steal employee credentials and breach company cyber security defences.

Ryan Economos, APAC field chief technology officer at email security firm Mimecast, told TechRepublic that such phishing attacks are rare in their use of Atlassian as a cover. But he noted that phishing attacks are becoming increasingly sophisticated, thanks to phishing kits and AI, which make it easier for cyber criminals to execute their activities.

Atlassian workspaces, Japanese ISPs, and a compliance cover story

Mimecast's Global Threat Intelligence Report 2024 H1 reported on the emergence of a new phishing tactic that used a compliance update cover story to target law firm employees. The phishing attacks:

- Leveraged popular local brand Atlassian's workspaces, as well as other unified workspace platforms, including Archbee and Nuclino, to send employees harmful emails that looked familiar and legitimate.

- Used device compliance updates as a cover, instructing employees via email that they needed to update their devices to remain compliant with company policy.

- Were designed to redirect those who clicked the link to a fake company portal, where attackers could harvest credentials and other sensitive information.

- Embedded the phishing link in an email sent from addresses associated with Japanese ISPs.

"There's quite a lot of personalisation in the emails such as details of a 'device' and several references to the company domain they are sending these campaigns to increase validity," Mimecast's report said.

"The sender address name always refers to the target organisation's domain name with the aim of fooling end users into thinking it is from their internal department."

The growing sophistication of phishing attacks

Economos noted that while the campaign initially targeted Australian law firms, it has since expanded to other industries and is no longer confined to the legal sector. He highlighted several aspects of the campaign that indicate increasing sophistication among threat actors.

Use of Atlassian and other workspaces

Economos said the growing use of Atlassian workspaces was a newer development for the market.

"Mimecast continues to see threat actors making use of services such as OneDrive and Google Docs to host files or links in their campaigns, but the use of workspaces such as Atlassian has not been heavily abused previously," he said.

Part of the campaign was an email that appeared to be from Atlassian's Confluence product. Mimecast referred to a "noticeable increase in the use of Atlassian" to evade detection in recent times.

"Abuse of legitimate services is an ongoing and evolving challenge," Economos said. "Attackers will continue to leverage reputable sources to launch and host their campaigns, in an attempt to evade detection."

Harvesting of tracker data intelligence

The campaign used postmark URLs to redirect users to the unified workspace solutions. Postmark URLs allow attackers to gather data such as location, browser details, and which part of the email was clicked, enabling them to leverage this intelligence to make the phishing lure more convincing.

Multiple URL obfuscation techniques

Making it more difficult for users to identify the true destination of the URL, the phishing campaign used "multiple obfuscation techniques," Mimecast said. This includes multiple redirections within the URL, encoded characters, and the insertion of tracking parameters.

Enlisting unsuspecting Japanese ISPs

Although the use of Japanese ISPs is not unique to this phishing campaign, Economos noted that they were exploited once again, as they had in several previous attacks.

"It continues to expose the lengths that threat actors will go to in order to successfully generate attacks on organisations," he commented.

Phishing attacks will get easier to mount — and more convincing

Phishing is still among the most common cyber threats among organisations, Economos said.

Generative AI and machine learning, while also helping defenders stop attacks, is expected to increase the sophistication and improve the targeting and content of phishing campaigns. This will drive defenders' need to detect and quickly respond to new and novel attack techniques.

"The biggest evolution has been the velocity and accuracy of phishing threats, through the use of phishing kits, automation, and AI-based technologies," Economos said. "These platforms allow even low-skill-level attackers to launch large-scale campaigns and an ability to quickly craft more convincing phishing emails to evade detection by traditional security tools."

Economos also noted the rise of pretexting — where a cyber criminal will research and pose as a character to provide a convincing story or "pretext" to trick the phishing victim — as well as Business Email Compromise, as significant factors in the evolution in the phishing threat landscape.

"As our work surfaces continue to diversify, threat actors are diversifying the vectors they exploit beyond email, targeting social media platforms, collaboration tools like Microsoft Teams, Slack, and OneDrive right through to vishing and smishing attacks using phone calls or text messages to deceive victims," he said.

5.NetSfere Intros Crypto-Agile Architecture with Post-Quantum Cryptography

by Ray Sharma
https://www.thefastmode.com/technology-solutions/37467-netsfere-intros-crypto-agile-architecture-with-post-quantum-cryptography

At the NetSfere Connections 2024 event in Munich, Germany, NetSfere, a global leader in next-generation secure and compliant messaging and mobility solutions, announced a significant leap forward in secure communications with the unveiling of its crypto-agile architecture, featuring seamless integration of Post-Quantum Cryptography (PQC).

As quantum computing technology advances, one of the most critical cybersecurity challenges on the horizon is the rise of "Harvest Now, Decrypt Later" attacks. In these scenarios, sensitive data harvested today could be decrypted by quantum capabilities in the future, posing a substantial risk to enterprises worldwide. NetSfere's PQC solutions directly address this threat by equipping its platform with cutting-edge post-quantum encryption, ensuring that data remains secure now and in the quantum era.

NetSfere's crypto-agile architecture is designed to meet future cryptographic challenges by integrating the NIST-recommended Module Lattice-based Key Encryption Module (ML-KEM) algorithm. This advanced post-quantum encryption ensures that NetSfere's security remains resilient and robust, even in the era of quantum computing, safeguarding enterprise data against the complex threats of tomorrow. Featuring NIST's FIPS 203 ML-KEM with Kyber-1024 strength security, the PQC-enabled platform is optimized for seamless deployment across iOS, Android, and web/desktop applications. NetSfere plans to extend the quantum-safe encryption to all its users by January 2025.

Prof. Dr. Pierre-Michael, CEO, CHCDO, CHCIO of AHIME Academy of Health Information Management Executives and ENTSCHEIDERFABRIK eHealth Incubator, emphasized the importance of continuous innovation in protecting health data.

These advanced technologies give enterprise IT full control over secure communications, ensuring seamless integration into existing workflows while safeguarding sensitive data across all platforms.

6.CISA guidance focuses on post-quantum cryptography tools

by Alexandra Kelley
https://www.nextgov.com/cybersecurity/2024/09/cisa-guidance-focuses-post-quantum-cryptography-tools/399904/?oref=ng-homepage-river

The Cybersecurity and Infrastructure Security Agency publicly released its post-quantum cryptography migration guidance on Friday, focusing on prepping the most vulnerable federal digital systems for the potential advent of a cryptographically-relevant quantum computer.

Accompanying the earlier release of the first standardized algorithms suited to protect information stored on classical computers from a hypothetical quantum computer attack, CISA's guidance, dated Aug. 15, lays out recommendations for federal civilian executive branch agencies conducting initial system inventories using automated cryptography discovery and inventory softwares.

CISA's wants agencies to launch their migration processes sooner rather than later. The guidance notes that the inventory process requires both manual data collection and the use of automated support.

"The primary goal of this strategy is to enable the assessment of agency [post-quantum cryptography] transition progress," the guidance reads. "Included is the use of [automated cryptography discovery and inventory] tools to support a [federal civilian executive branch] agency in its creation of an inventory of its information systems and assets that contain CRQC-vulnerable cryptography."

CISA is asking civilian agencies to first identify potential vulnerabilities and migrate high-impact information systems, or assets storing sensitive information on a given network. The guidance also prioritizes assets that "contain data expected to remain mission-sensitive in 2035."

CISA also identified three on-going research arenas that are expected to inform ongoing PQC migration efforts. They hinge upon understanding the industry offerings of automated cryptographic discovery tools and how well they function to accurately detect embedded algorithms in use within certain software packages.

"CISA has not been able to confirm the full scope of cryptographic algorithm detection capabilities that will be available via automated cryptographic discovery tools," the guidance reads.

Future work in this arena will be helmed by CISA along with other partners in the National Institute of Standards and Technology's National Cybersecurity Center of Excellence's "Migration to PQC" project.

Utilizing automated inventory scanning softwares has been a minor point of contention between government and industry, as government partners work to establish baseline standards of trust in automated cryptography discovery and inventory tools while industry says manual network inventories are too cumbersome to undertake efficiently.

The new guidance does state that such tools — operating individually or in combination with other network analysis efforts — may help with gathering certain cryptographic inventory data from networks, file systems, database systems, and software packages. It adds that further steps are required for entities to integrate automated discovery tools in their network scans. CISA notes that some automated capabilities are available to agencies via the Continuous Diagnostics and Mitigation program, but these offerings aren't yet adequate to CISA's goals. CDM's current dashboard and analytics “will need expanding to support data elements provided by ACDI tools,” the report states.

“This pilot would determine the optimal level of integration including data elements and interfaces,” the guidance states. “As part of this pilot program, a comparative analysis should be conducted to determine the extent that ACDI tools can discover cryptographic assets vice those assets known via manual means.”

Following the guidance’s release, CISA and other federal partners will take more steps in the coming months to support PQC overhauls within the federal government networks, including updating reporting requirements and further evaluating tools suitable for government network analyses.

7.How to prepare for post-quantum computing security

by Kyle Johnson
https://www.techtarget.com/searchsecurity/feature/How-to-prepare-for-post-quantum-computing-security

Quantum computers are projected to break many of the cryptographic standards that have adequately protected data for decades -- a scary thought for security professionals and organizations alike.

While companies don't need to hit the panic button over quantum quite yet -- it will likely be five to 10 years before the technology is ready -- that doesn't mean they can ignore it.

President Joe Biden signed two quantum computing presidential directives in 2022, signaling it was time to figure out how to handle the emerging technology. The directives called for the creation of quantum-resistant cryptographic standards -- a task NIST revealed results for in 2024 after more than half a decade of effort -- and the preparation for federal agencies to adopt these future standards.

"The culmination of the work NIST has been doing is a starting gun for upgrading to post-quantum cryptography," said Colin Soutar, managing director at Deloitte.

With the gun sounded, companies need to figure out how quantum computing will affect them once it arrives, which could call for better data protection now and preparation for post-quantum cryptography (PQC).

The quantum security worry

The major concern with quantum computing is how easily it could crack data transmission cryptography algorithms. The asymmetric RSA algorithm, for example, which is based on integer factoring and provides sufficient security on classical computers, will be breakable by quantum computers.

Attackers are aware of this issue and have begun to do what is known as *data scraping* -- collecting encrypted data in hopes it will be useful later. Because storage is cheap, attackers are harvesting encrypted data now to crack once quantum computing matures.

Post-quantum computing also shines a spotlight on the ongoing issue of legacy systems and devices, said Jon France, CISO at ISC2. "History shows us that we're really bad at dealing with legacy.”

The classic solution to protecting legacy systems generally involves wrapping security around these systems -- a Band-Aid approach that won't work in a post-quantum world. "Quantum is going to be that point of inflection that will rapidly undo the notion we can protect classic systems and devices," France said.

How to prepare for PQC security

Organizations should expect a complete PQC migration to be a multiyear effort, Soutar said, due to the number of services that need updating for PQC and the difficulty for each, as well as dependence on third parties implementing PQC on their systems to secure the entire supply chain.

To prepare for migration now that PQC is standardized, companies should consider the following steps.

1. **Inventory and classify data**

   Review data, and decide what is deemed sensitive. Conduct a data inventory to understand what data the company has and its data classification to understand what data needs which cryptography protections.

   Consider what data needs stronger protection now in terms of the data scraping threat. Not all data a company currently stores will matter beyond the next five to 10 years.

   "What data is OK four years from now that I am not worried about someone scraping?" said Christopher Savoie, CEO of AI vendor Zapata. "On the other hand, what would I be worried about for years?" Data to pay attention to includes corporate or trade secrets and other business-critical information. Take the appropriate actions to ensure data is safe now and into the future.

2. **Understand future exposure**

   With data inventoried and classified, it's time to conduct risk assessments to understand how data is protected against future risks.

   "Organizations should start looking at their potential exposure to understand what their reliance on cryptography is," Soutar said. "It might be deeply embedded in third-party tools; it might be proprietary, transactional capabilities. You need a sense of where cryptography is embedded into your systems and how data is being protected."

   Understanding current and future exposure enables organizations to determine urgency around PQC adoption and start building their roadmap.

   Consider PQC from a business impact perspective -- not just the technical aspects of implementing new cryptographic algorithms. Select someone to lead the PQC migration effort who can explain to executives the importance of PQC and how it can mitigate security incidents and breaches.

   Also, consider the encryption needs of IoT and other embedded devices -- many of which are incapable of handling the increased memory and compute required for PQC algorithms, said Chris Hickman, CSO of identity and access management vendor Keyfactor. Organizations should vet PQC algorithms, such as Falcon and Kyber, that could meet PQC requirements on smaller devices with limited RAM.

3. **Create a mitigation strategy**

   With data inventoried and potential exposure understood, the next step is to create mitigation strategies and a team of employees to lead those efforts.

   "Using a mitigation group, start looking at what policies and procedures need to be in place for when the inevitable happens," Savoie said.

   This should include a data security policy, incident response plan and business recovery plan, at minimum. Also, assess what company data might already be exposed and stored by attackers, and determine how to handle those situations. Next, look at the critical data stored now, and decide whether it needs additional layers of encryption to protect it.

   Symmetric encryption, commonly used by organizations to keep stored data secure, won't be largely affected by quantum computing. Grover's algorithm, which demonstrates how quantum computing quadratically speeds up database searches, has shown it halves the time needed to break symmetric encryption. NIST, therefore, recommended organizations use at least AES-192 or AES-256 to encrypt stored data.

   Data in transit, however, is at risk of being broken by quantum computing. To counter this, replace asymmetric algorithms with PQC encryption standards. This task plays into the last aspect of mitigation, Savoie added -- that organizations need to start thinking about how to become and remain crypto-agile.

   "As standards change going forward, we need to ensure infrastructure is in a place where we can adapt to new threats and new technologies to mitigate those threats," Savoie said. "Getting your systems crypto-agile and forward-compatible to new standards takes time and is something you need to start working on now."

PQC implementation options

In August 2024, NIST announced it had selected the following three PQC algorithms designed to withstand classical and quantum computing cracking efforts:

- **Kyber**, public key encapsulation.
- **Dilithium**, a lattice-based digital signature scheme.
- **SPHINCS+**, a stateless hash-based signature scheme.

NIST continues to evaluate additional algorithms, including Falcon, which is expected to be standardized later in 2024. Further evaluation of other algorithms helps NIST ensure that, if a current algorithm doesn't work as expected, then organizations have other options to use.

France recommended organizations select more than one algorithm -- and ones that don't rely on the same math. "This provides some protection against future failure," he said.

Beyond PQC algorithms, organizations can also consider quantum key distribution (QKD), which uses quantum mechanics to securely exchange encryption keys. Data encrypted via QKD creates a random quantum state that is difficult to copy. Many QKD protocols can also detect eavesdroppers. The National Security Agency, however, has stated this option is not viable on its own as it now stands.

Organizations could, therefore, combine PQC encryption standards and QKD, suggested Rik Turner, analyst at Omdia. This would make it more difficult for attackers, he noted, because they would need to break through both encryption and QKD to access data in transit.

Organizations aren't on their own in preparing for a post-quantum security world. Turner advised reaching out to vendors to learn if and how they're adding PQC into their tools and services. This could reduce the costs of a migration, especially as QKD can be expensive to implement.

8.HSBC tests post-quantum VPN tunnel for digital ledgers

by Cliff Saran
https://www.computerweekly.com/news/366611375/HSBC-tests-post-quantum-VPN-tunnel-for-digital-ledgers

HSBC has worked with Quantinuum to trial “the first application” of quantum-secure technology for distributing tokenised physical gold. The HSBC Gold Token for retail investors in Hong Kong allows the bank’s customers to acquire fractional ownership of physical gold. Moving these tokens over financial networks requires encryption that cannot be broken by high-performance quantum computers and does not impact performance.

HSBC said that the trial demonstrates the commitment of its foreign exchange and commodities businesses to safeguard critical applications from potential future quantum computing attacks.

According to HSBC, the work with Quantinuum also shows a cost- effective approach to protecting existing production distributed ledger technology (DLT) in the short and medium term without the need for re-architecting the DLT.

HSBC said it has also demonstrated the interoperability of its gold tokens by using post-quantum cryptography (PQC) to move digital assets safely across distributed ledgers via secure networks. This includes the capability to convert HSBC’s gold tokens into ERC-20 fungible tokens, thereby enhancing distribution and interoperability with other DLTs and digital wallets. HSBC said the approach it has taken addresses clients’ evolving needs and regulations.

In a whitepaper looking at the trial, HSBC said that although DLT and asset tokenisation offer significant benefits, it is crucial to enhance and future-proof the security measures surrounding these technologies, which involves migrating to quantum-safe cryptography to ensure the continued resilience of its financial systems against both current and emerging cryptographic threats.

“Despite the robust security offered by distributed ledgers and blockchain through encryption and decentralised consensus mechanisms, the rapid advancement of the quantum computing cyber threats necessitates a proactive approach to future-proofing these systems,” the company said.

“It is crucial to not only maintain but to continue to enhance the security measures surrounding distributed ledgers, ensuring that they remain resilient against both current and emerging cyber threats.”

PQC is the approach the industry is taking to protect against cryptographic keys being broken by powerful quantum computers. HSBC said that a major concern of implementing PQC into distributed ledgers is the potential impact on performance.

"These new algorithms have larger key sizes for the cryptography and may have significant effect on the operational use of distributed ledgers," the company explained in the whitepaper.

"It has been argued that the signature and larger key sizes used in PQC systems would cause an increase in block size and signature time. This has the unfortunate consequence of affecting the performance, efficiency and execution speed of the whole distributed network."

The trial, which is aligned with Project Leap – an initiative to quantum-proof financial systems – uses a PQC-VPN tunnel in a gold tokenisation environment. HSBC said it observed minimal impact on performance levels when sending data through the tunnel, irrespective of the size of the data.

Philip Intallura, global head of quantum technologies at HSBC, said: "HSBC was the first international bank to offer tokenised physical gold and is now building on that innovation with cutting-edge cybersecurity protection for the future. This pilot successfully demonstrated the viability of deploying these advanced technologies for a real-world business environment."

9.Google Confirms New Quantum Encryption For Chrome Is Coming Nov. 6

by Davey Winder
https://www.forbes.com/sites/daveywinder/2024/09/18/google-confirms-new-quantum-encryption-for-chrome-is-coming-nov-6/

It may well come as a surprise to you that Google has been experimenting with bringing post-quantum encryption to the Chrome desktop browser, but it's a fact. Google announced that it was helping protect Chrome traffic at the transport layer security level back in August 2023. That was using a yet-to-be standardized algorithm called Kyber. Now members of the Chrome team have confirmed that, as from Chrome 131 coming November 6, the experiment is over and that post-quantum protection will be available to all browser users in the form of the fully standardized Module Lattice Key Encapsulation Mechanism. Here's what you need to know.

A New Quantum Security Path For The Google Chrome Web Browser

In a Google security blog posting dated September 13, David Adrian, David Benjamin, Bob Beck & Devon O'Brien confirmed that the Chrome browser is heading down a new quantum security path. However, it's more accurately a path that has had the builders in to lay a new tarmac top over the existing foundations. The non-standard, as defined by the U.S. National Institute of Standards and Technology, Kyber algorithm has now been tweaked and standardized at last. NIST said that it has "released a final set of encryption tools designed to withstand the attack of a quantum computer," and is encouraging server admins to begin the transition as soon as possible.

Thanks to the minor technical changes made to the post-quantum algorithm Kyber, part of a hybrid key exchange also employing the pre-quantum X25519 algorithm, Google has implemented the new Kyber, now called the Module Lattice Key Encapsulation Mechanism, into its cryptography library, Boring SSL. This means that the standardized version can now be used across services depending upon this library for securing transport layer security. It also means that the new version of ML-KEM is no longer compatible with previously deployed versions of Kyber, so Google is having to make changes, coming in Chrome 131, to allow for this.

What Is Post-Quantum Cryptography And Why Does Chrome Need It?

The point of Transport Layer Security protocols in networking is to protect your data while it is in transit and also for website authentication to validate identity. The cryptographic part of this equation makes it harder for an attacker to intercept, access, or alter either of these properties. However, as Google points out, the evolution of quantum computers threatens to make such protections impotent. “Many types of asymmetric cryptography used today are considered strong against attacks using existing technology,” Google said, “but do not protect against attackers with a sufficiently-capable quantum computer.” Post-quantum cryptography, also referred to as quantum-resistant cryptography, is designed to protect against both quantum and classic attack methodologies.

As you might imagine, there are plenty of obstacles to overcome when thinking about implementing a post-quantum cryptography TLS solution. Google lists these primarily as:

- Post-quantum cryptography is too big to be able to offer two post-quantum key share predictions at the same time.
- The danger of regressing users’ post-quantum security, requiring the delay until Chrome 131 is available to make this change so give server admins a chance to update their implementations.
- Kyber was always experimental, continuing to support it risks ossification on non-standard algorithms.

“We’re excited to continue to improve security for Chrome users against both current and future computers,” Google said.

Talking about the dangers of someone, likely a nation-state, developing quantum technology capable of breaking encryption and keeping it secret, Tim Callan, chief experience officer at Sectigo, said, “It is imperative that businesses take their own proactive measures to prepare for this eventuality by transitioning to quantum-safe algorithms before it is too late. Transitioning to quantum-resistant cryptography will become a mainstream boardroom discussion.”

10.NCSC exposes Chinese company running malicious Mirai botnet

by Alex Scroxton

https://www.computerweekly.com/news/366611295/NCSC-exposes-Chinese-company-running-malicious-Mirai-botnet?utm_campaign=20240919_NCSC+exposes+Chinese+company+running+malicious+Mirai+botnet&utm_medium=email&utm_source=MDN&source_ad_id=366611295&asrc=EM_MDN_302349086&bt_ee=O6ubZqH6IOb5GrvdYo6zEzM1zYTgGkP4KQU1Rfa3FfsJMqJ1yBCjamA5MnW1ym9a&bt_ts=1726746732176

The UK’s National Cyber Security Centre (NCSC) and its counterpart Five Eyes agencies have accused a China-based company acting as a front for the state of running a massive botnet comprising over 250,000 internet-connected devices, about 8,500 of them located in the UK.

The compromised devices include enterprise network and security tools such as routers and firewalls, and internet of things (IoT) products such as CCTV cameras and webcams. Unbeknownst to their owners, they are being used to conduct coordinated cyber attacks, including distributed denial of service (DDoS) attacks and malware delivery.

"Botnet operations represent a significant threat to the UK by exploiting vulnerabilities in everyday internet-connected devices with the potential to carry out large-scale cyber attacks," said NCSC operations director Paul Chichester.

"Whilst the majority of botnets are used to conduct coordinated DDoS attacks, we know that some also have the ability to steal sensitive information.

"That's why the NCSC, along with our partners in Five Eyes countries, is strongly encouraging organisations and individuals to act on the guidance set out in this advisory – which includes applying updates to internet-connected devices – to help prevent their devices from joining a botnet."

The company in question, Integrity Technology Group, is based out of Beijing and on the surface appears to operate as a legitimate provider of network security services.

However, according to the joint advisory, which can be read in full here, it has also put its expertise to use in the service of the Chinese government – Integrity's China Unicom Beijing Province IP addresses are known to have been used to access other operational infrastructure used in cyber attacks on victims in the US.

According to the authorities, the FBI has engaged with a number of these victims and has uncovered activity consistent with the tactics, techniques and procedures (TTPs) favoured by a state-backed advanced persistent threat (APT) actor tracked as Flax Typhoon, but also known as RedJuliett and Ethereal Panda, among other things.

Its botnet uses the infamous Mirai malware family to hijack devices running Linux-based operating systems. Integrity targets these devices via a number of disclosed common vulnerabilities and exposures (CVEs).

Once the Mirai payload has been downloaded and executed, it begins processes on the device to establish a connection with Integrity's command-and-control (C2) infrastructure using Transport Layer Security (TLS) over port 443. It also gathers and exfiltrates system information, including operating system versions, memory and bandwidth details for enumeration purposes. It also sends requests to "c.speedtest.net" to gather additional internet connection details. Additionally, the investigation found, some of the Mirai payloads are self-deleting to avoid detection.

Meanwhile, upstream, Integrity operates a tier of management servers via TCP port 34125 to run the botnet's C2 infrastructure. The servers host a MySQL database holding information on the compromised devices that, as of June this year, is thought to contain over 1.2 million records. The servers also host an application known as Sparrow to interact with the botnet – the code for this application is stored in a Git repository and defines various functions that enable users to send tasking and exploitation commands to the compromised devices, among other things. Sparrow can also provide device vulnerability information to users, and a subcomponent called "vulnerability arsenal" lets them exploit traditional networks via the victim devices.

11.Cops infiltrate ‘Ghost’ encryption app used by drug lords, mafia

by Antoaneta Roussi

https://www.politico.eu/article/police-ghost-encryption-app-drug-lords-mafia-europol/

In the war between law enforcement and encrypted messaging apps, cops won one battle on Wednesday.

Police agencies across the world announced they had infiltrated an encrypted chat platform called **Ghost**, tapping into reams of private communications of criminal networks and leading to the arrest of 51 suspects so far.

Criminals involved in drug trafficking and money laundering coordinated activities on the platform daily, said officials at a press conference at the headquarters of the European Union's law enforcement agency, Europol, in The Hague. The platform is known for its advanced security features, which include layers of encryption and the fact that users can purchase it without disclosing personal details.

Authorities in Australia, Canada, France, Ireland, Italy, the Netherlands, Sweden and the United States worked with Europol and Eurojust, the European Union Agency for Criminal Justice Cooperation, to map the platform’s global infrastructure.

They found servers in France and Iceland and located the owners of Ghost in Australia, where authorities arrested a 32-year-old administrator, The Associated Press reported.

“No matter how advanced the technology, no matter how secure they think their communications are, we will find them and we will shut down their criminal activities,” Europol Executive Director Catherine De Bolle said.

Authorities also busted members of the Italian mafia, motorcycle gangs and organized crime groups in arrests spanning from Canada to Sweden, Ireland and Italy, officials said.

This is the latest high-profile encrypted communication bust to be carried out by a team of international law enforcement in recent years.

In 2020, France and the Netherlands led the investigation into EncroChat, a "cryptophone" company selling encrypted communication services and devices that were used by criminal networks, many of which were involved in drug trafficking and organized crime.

Authorities penetrated EncroChat to get access to "more than 100 million encrypted messages by criminals" and monitored the chats of "thousands of criminals" in real time, leading to more than 120 people charged with drug and arms trafficking, extortion, acts of torture and attempted murder in Belgium.

In March 2021, police in Belgium and the Netherlands raided more than 200 homes and arrested dozens of suspected criminals after cracking into Sky ECC, an encrypted chat app based in Canada and the U.S. that claims to offer “the world’s most secure messaging platform.”

And in June of that year, the U.S. Federal Bureau of Investigation revealed it created and ran an app called ANOM that got picked up by criminal gangs seeking new ways to communicate securely, leading to a whopping 500 arrests globally in a two-day takedown of criminal networks.

12.It's time for CIOs to think about post-quantum cryptography

by Paul Barker
https://www.cio.com/article/3526658/its-time-for-cios-to-think-about-post-quantum-cryptography.html

Quantum is key among six critical trends that Info-Tech's new Tech Trends report predicts will shape the future of IT in 2025.

With the threat that quantum computers will break current encryption methods "looming on the horizon," the adoption of post-quantum cryptography is now a critical priority, particularly for those industries that handle sensitive data, a report released Tuesday revealed.

The current and future state of quantum are among six trends contained in Info-Tech Research Group's Tech Trends 2025 report, based on responses from an estimated 1,000 IT decision-makers in the US, Canada, Asia Pacific (APAC), and Europe, who were surveyed between March and July of this year.

According to the report, "quantum computing has moved beyond theoretical exploration and is now accessible through cloud platforms, enabling real-world business experiments. As organizations begin leveraging quantum hardware to solve complex problems, industries such as media, government, and financial services are leading the charge in quantum investments."

When it comes to plans to invest in quantum computing and post-quantum cryptography, there were clear similarities among respondents in various sectors. Results showed 33% of organizations in the media, telecom, and technology sectors investing in quantum computing, followed by 27% in the public sector, and 20% in financial services.

On the post-quantum cryptography front, 31% of what Info-Tech described as "advanced IT departments" are planning to invest in the technology before the end of next year. That compares to 16% of "average IT departments."

Brian Jackson, principal research director and lead author of the report, said, "preparedness is critical as organizations face increasing risks from 'Harvest Now, Decrypt Later' cyberattacks."

In the report, he wrote, "the progress of quantum computing harkens the risks organizations will face as a result of the new computing capabilities it promises. The most urgent among them is that quantum computing will inevitably be able to crack a majority of the encryption methods that we use today."

He pointed out that the US National Institute of Standards and Technology (NIST) "has been working on new quantum-resistant methods of encryption known as post-quantum cryptography since 2016. It published the first set of those standards in the summer of 2024, marking a one-year countdown for US federal agencies to create a plan to migrate their encryption standards as mandated by the US Quantum Computing Cybersecurity Preparedness Act. If other organizations are wise, they won't be too far behind."

Four more predictions for 2025

The other four trends Info-Tech sees as being pervasive next year are:

Deepfake Defense: Deepfakes are, the research firm said, becoming a "powerful tool for fraud and misinformation. IT leaders are prioritizing AI-powered detection tools and content authentication methods, such as blockchain, to combat the rising threat of AI-powered cyberattacks and ensure the integrity of their data."

Expert Models: According to the report, "as AI matures, organizations are increasingly developing custom AI models tailored to their specific industries. These expert models improve the accuracy and relevance of AI outputs, enabling businesses to derive exponential value from IT investments."

AI Sovereignty: The reports stated that organizations "are focusing on balancing AI adoption with governance and control to protect sensitive data, reduce costs, and ensure AI performance. By 2026, more companies will run localized AI models to improve cost-effectiveness and maintain control over their AI initiatives." Upwards of 65% of respondents cited privacy and security concerns being among the top factors influencing AI investment decisions.

Lessons for CIOs

According to Jackson, there are two lessons to draw from the findings for IT leaders looking ahead to 2025. Historically, he wrote, "chief information officers have been accountable for the recordkeeping at their organizations. Like a resident historian, CIOs maintained the integrity of an organization's past, making it verifiable and auditable. With digital transformation, CIOs were asked to do more to report on the current state of the organization – in as real-time as possible."

"[The business intelligence and analytics] required to drive decision-making couldn't be based on old information. Now, as firms push their investment into artificial intelligence (AI), and more specifically generative AI, the focus shifts to simulating the future. In a fast-changing and uncertain world, AI predicts different scenarios based on the probability of the outcome. Generative AI provides an output that aims to simulate a human response prompted by an input."

13.Mathematical Certainty in Security: The Rise of Fully Homomorphic Encryption

by David Archer
https://embeddedcomputing.com/technology/security/software-security/mathematical-certainty-in-security-the-rise-of-fully-homomorphic-encryption

Interest in Fully Homomorphic Encryption (FHE) is growing as companies seek more robust data privacy solutions in an increasingly regulated world. FHE allows data to be processed without ever being decrypted, a breakthrough that could revolutionize industries where data security is paramount. To delve deeper into this technology, let's look at the development, challenges, and future of FHE, shedding light on its potential to redefine secure data processing.

The development of Fully Homomorphic Encryption (FHE)

Homomorphic encryption developed gradually over the past several decades, first by the accidental discovery of partially homomorphic systems and then more purposefully until its full capability emerged in 2009 and the following decade. The implications of these developments were staggering: we could send data to the cloud, an AI engine, or another third party for processing without ever worrying about a resulting privacy breach. However, the computational power required to implement FHE remained many orders of magnitude greater than computing "in the clear", making broad adoption a difficult sell, and relegating FHE to an academic endeavor.

Now, however, interest in and advancement of FHE is driven by new forces. Companies must navigate a complex legal framework encompassing GDPR in Europe, CCPA in California, and diverse regulations in at least 14 other U.S. states. And yet, the commercial appetite for third-party data continues to grow: Companies are counting on the ability to ingest new data to solve hard problems ranging from detecting financial fraud to researching medical treatments.

At the same time, privacy-assuring alternatives to FHE face significant headwinds. Confidential computing methods such as Trusted Execution Environments (TEEs) have been shown time and again to be vulnerable to both side-channel attacks and direct breaches, placing the companies that rely on them at risk. Other privacy-assuring approaches such as secure multi-party computation typically require networks of computers to be and stay online together throughout computations, requiring complex network configurations and vulnerable to failure if any one of the participating machines or network links fails.

FHE, on the other hand, has cryptographically sound proofs of privacy, requires no complex network configurations, and relies only on a single compute server's reliability. This pairing of cryptographically strong privacy guarantees with simplicity of deployment makes FHE a strong contender for practical, secure privacy assurance in fields such as finance and healthcare, where privacy is paramount. With FHE, companies can perform computations on encrypted data, ensuring that data remains protected throughout storage, transit, and processing. Now, we're at the forefront of a new wave of hardware accelerators that will take FHE the last mile to commercial performance viability. We're on the brink of a whole new era in data privacy. Within a generation, there will be no such thing as sharing or outsourcing computation on unencrypted data.

Computing on encrypted data

In the past, we've encrypted data at rest - in storage media such as disk drives - and in transit on networks. However, to process data, we needed to decrypt it, because no practical encryption mechanisms also allowed computation. Decrypting the data also made it visible to anyone performing that computation, requiring the data's owner to trust those performing the computation. Novel encryption schemes such as those used in FHE not only keep the data from being revealed, but also allow computation on the data in its encrypted state. As a result, data owners need not trust those performing computations to keep the data private. This "zero trust, full computation" breakthrough is a sea change in the relationship between data owner and data processor, enabling outsourcing of computation without risk of data compromise.

Challenges in implementing FHE

There are three main challenges:

1. The computational complexity of FHE is a performance challenge. FHE computations are dramatically slower than unencrypted computations, often by several orders of magnitude, making it difficult to achieve practical performance levels. This slowdown is due to the additional work required by CPUs and GPUs to manage the complicated data representations used in FHE.

2. The data expansion typically seen in FHE encryptions is a storage and network bandwidth challenge. Homomorphically encrypted data is also substantially larger than unencrypted data, requiring times more storage space. Current research ideas such as hybrid FHE are insufficiently developed to answer this challenge so far.

3. The complex algorithms required to compute in FHE are a usability challenge. Programming in FHE - even with the advent of some fantastic FHE libraries - is a major challenge because of the many parameters that must be correctly chosen for FHE, and because of the many auxiliary operations needed to “manage” FHE computations, which (due to lack of tooling) cannot be automatically handled by the programmer’s tools.

How to encrypt data for FHE and the role of homomorphism

To answer that question in full, we’d need to talk about Gaussian noise sampling, polynomial representations of data, residue number systems, the Learning With Errors math problem, public key encryptions, prime modular arithmetic, and high-dimensional vector spaces – not really great topics for polite company! Instead, let’s do a quick summary. In FHE, we move data from the normal number line into an alternative space. What’s important is that the movement of the data is an encryption - something much harder to undo if you don’t possess a specific key. The other thing that’s important is that the alternative space be homomorphic (homo- for “same”, -morphic for “shape”) to the regular number line with respect to multiplication and addition so that once the data is moved, you can add and multiply at will knowing that when you move the data back (via decryption), those multiplies and adds did the expected thing to the data.

Programs, Computations, and Limitations

The security provided by FHE is based (in part) on adding a little “noise” to data during the encryption process. One of the problems with FHE is that when you add or multiply the data, that noise grows - just as you’d expect. After a certain limited number of operations on a data item, the noise can grow large enough that decryption is no longer possible. To deal with that problem, FHE uses a special but very expensive process to remove noise without revealing data, so that computation can continue. This special process must be done every few operations to keep the data fresh, but it is by far the most expensive operation used in FHE - consuming as much as 95% of computation time.

That’s a long introduction to say that the best-suited computations for FHE are those that don’t require very many operations sequentially on data, so the noise removal process is unused or seldom-used. What kinds of computations fall into that regime? Linear algebra and private information query are two examples. Extending those ideas, statistical computation such as regressions, certain kinds of image processing, and even relatively simple neural networks can be good targets for FHE-assured privacy.

Optimization and careful selection of use cases are essential to maximize the benefits of FHE.

Practical applications for FHE

FHE unlocks entirely new applications across industries that would be impossible without mathematically guaranteed privacy. While some of the following applications are still challenging at scale for FHE, they are all good targets for hardware-accelerated FHE in the near future.

Healthcare statistics: FHE facilitates large-scale analysis of health records while maintaining patient privacy. Clinicians and insurance providers can analyze data on patient satisfaction, hospital readmission, and other factors across their patient populations. This comprehensive analysis helps uncover more effective treatments and personalized care plans, improving overall quality of life.

Finance: FHE enables the secure sharing of financial transaction data across institutions and borders, allowing banks to identify fraudulent accounts and transactions regardless of their origin. This enhanced ability to detect and prevent fraud strengthens the financial system's integrity.

Machine Learning: FHE allows machine learning models to analyze sensitive data without exposing the data itself. For example, image recognition can identify security threats or legal violations without inappropriate surveillance, and medical scans can be analyzed without risking patient data exposure.

Market Intelligence: FHE enables manufacturers to share inventory, sales, distribution data, and more with analysts, data brokers, and even competitors. This collaboration enhances the ability to predict and respond to market changes and manage supply chains effectively. Data brokers can also perform computations on private data such as GPS locations, uncovering valuable population-level insights without compromising individual privacy.

Cross-Organizational Coordination: FHE enables secure and private data sharing between different jurisdictions.

Advancements in FHE, and comparison to unencrypted computations

Advancements in FHE have focused on optimizing algorithms, developing specialized hardware to accelerate processing, and to a lesser degree, leveraging parallel processing. Niobium's FHE Hardware Acceleration chip, for example, includes proprietary optimizations and a hardware-software co-design approach to enhance performance. Despite these improvements, FHE is still slower than traditional unencrypted computations. Even with the first generation of Niobium's hardware acceleration, well-chosen FHE applications can be hundreds of times slower than "in the clear" computations, though efforts are ongoing to reduce this gap significantly.

Already, we have made meaningful progress. By accelerating FHE computation by a factor of up to 10,000, we have met the practical requirements for several applications, particularly in sectors like finance, insurance, and healthcare. According to feedback from potential users, the current speed of our chip is sufficient to unlock many new and previously infeasible use cases for FHE, allowing them to guarantee data security and privacy during processing. These companies will be able to use and share data for processing and analysis while keeping the actual contents of the data utterly confidential. This is the key to a safe, profitable data economy. Using FHE hardware acceleration, companies can conduct analysis of rich data sets without compromising security or privacy.

Proof-of-concept work in various industries has shown us that our current computation performance is more than adequate for many practical applications, such as neural network-based machine learning or financial fraud detection across international borders. We are also leveraging feedback from this work to refine our hardware architecture and guide future developments to achieve even better performance.

The Future of FHE

Many companies and investors have recognized the potential of FHE to revolutionize the data economy – there has been at least $200M worth of venture investment in FHE hardware acceleration alone, as well as substantial investment by the US Government in the form of a dedicated DARPA program, DPRIVE. These investments are driving toward the future goal of FHE achieving performance parity with traditional unencrypted computations, making FHE practical for a wider range of applications. Developments in the next few years are expected to focus further on optimizing algorithms and hardware to reduce computational overhead, developing user-friendly programming interfaces and tools to simplify FHE implementation, and expanding the range of practical applications to demonstrate the value of FHE in various industries through proof-of-concept projects and real-world deployments.

14.Quantum Computing: Between Hope and Hype

by Scott Aaronson
https://scottaaronson.blog/?p=8329

When Rafi (from UCLA) invited me to open this event (a binational US/India workshop, for national security officials from both countries to learn about the current status of quantum computing and post-quantum cryptography), it sounded like he wanted big-picture pontification more than technical results, which is just as well, since I'm getting old for the latter. Also, I'm just now getting back into quantum computing after a two-year leave at OpenAI to think about the theoretical foundations of AI safety. Luckily for me, that was a relaxing experience, since not much happened in AI these past two years. [Pause for laughs] So then, did anything happen in quantum computing while I was away?

This, of course, has been an extraordinary time for both quantum computing and AI, and not only because the two fields were mentioned for the first time in an American presidential debate (along with, I think, the problem of immigrants eating pets). But it's extraordinary for quantum computing and for AI in very different ways. In AI, practice is wildly ahead of theory, and there's a race for scientific understanding to catch up to where we've gotten via the pure scaling of neural nets and the compute and data used to train them. In quantum computing, it's just the opposite: there's right now a race for practice to catch up to where theory has been since the mid-1990s.

I started in quantum computing around 1998, which is not quite as long as some people here, but which does cover most of the time since Shor's algorithm and the rest were discovered. So I can say: this past year or two is the first time I've felt like the race to build a scalable fault-tolerant quantum computer is actually underway. Like people are no longer merely giving talks about the race or warming up for the race, but running the race.

Within just the last few weeks, we saw the group at Google announce that they'd used the Kitaev surface code, with distance 7, to encode one logical qubit using 100 or so physical qubits, in superconducting architecture. They got a net gain: their logical qubit stays alive for maybe twice as long as the underlying physical qubits do. And crucially, they find that their logical coherence time increases as they pass to larger codes, with higher distance, on more physical qubits. With superconducting, there are still limits to how many physical qubits you can stuff onto a chip, and eventually you'll need communication of qubits between chips, which has yet to be demonstrated. But if you could scale Google's current experiment even to 1500 physical qubits, you'd probably be below the threshold where you could use that as a building block for a future scalable fault-tolerant device.

Then, just last week, a collaboration between Microsoft and Quantinuum announced that, in the trapped-ion architecture, they applied pretty substantial circuits to logically-encoded qubits—-again in a way that gets a net gain in fidelity over not doing error-correction, modulo a debate about whether they're relying too much on postselection. So, they made a GHZ state, which is basically like a Schrödinger cat, out of 12 logically encoded qubits. They also did a "quantum chemistry simulation," which had only two logical qubits, but which required three logical non-Clifford gates—which is the hard kind of gate when you're doing error-correction.

Because of these advances, as well as others—what QuEra is doing with neutral atoms, what PsiQuantum and Xanadu are doing with photonics, etc.—I'm now more optimistic than I've ever been that, if things continue at the current rate, either there are useful fault-tolerant QCs in the next decade, or else something surprising happens to stop that. Plausibly we'll get there not just with one hardware architecture, but with multiple ones, much like the Manhattan Project got a uranium bomb and a plutonium bomb around the same time, so the question will become which one is most economic.

If someone asks me why I'm now so optimistic, the core of the argument is 2-qubit gate fidelities. We've known for years that, at least on paper, quantum fault-tolerance becomes a net win (that is, you sustainably correct errors faster than you introduce new ones) once you have physical 2-qubit gates that are ~99.99% reliable. The problem has "merely" been how far we were from that. When I entered the field, in the late 1990s, it would've been like a Science or Nature paper to do a 2-qubit gate with 50% fidelity. But then at some point the 50% became 90%, became 95%, became 99%, and within the past year, multiple groups have reported 99.9%. So, if you just plot the log of the infidelity as a function of year and stare at it—yeah, you'd feel pretty optimistic about the next decade too!

Or pessimistic, as the case may be! To any of you who are worried about post-quantum cryptography—by now I'm so used to delivering a message of, maybe, eventually, someone will need to start thinking about migrating from RSA and Diffie-Hellman and elliptic curve crypto to lattice-based crypto, or other systems that could plausibly withstand quantum attack. I think today that message needs to change. I think today the message needs to be: yes, unequivocally, worry about this now. Have a plan.

So, I think this moment is a good one for reflection. We're used to quantum computing having this air of unreality about it. Like sure, we go to conferences, we prove theorems about these complexity classes like BQP and QMA, the experimenters do little toy demos that don't scale. But if this will ever be practical at all, then for all we know, not for another 200 years. It feels really different to think of this as something plausibly imminent. So what I want to do for the rest of this talk is to step back and ask, what are the main reasons why people regarded this as not entirely real? And what can we say about those reasons in light of where we are today?

Reason #1

For the general public, maybe the overriding reason not to take QC seriously has just been that it sounded too good to be true. Like, great, you'll have this magic machine that's gonna exponentially speed up every problem in optimization and machine learning and finance by trying out every possible solution simultaneously, in different parallel universes. Does it also dice peppers?

For this objection, I'd say that our response hasn't changed at all in 30 years, and it's simply, "No, that's not what it will do and not how it will work." We should acknowledge that laypeople and journalists and unfortunately even some investors and government officials have been misled by the people whose job it was to explain this stuff to them.

I think it's important to tell people that the only hope of getting a speedup from a QC is to exploit the way that QM works differently from classical probability theory — in particular, that it involves these numbers called amplitudes, which can be positive, negative, or even complex. With every quantum algorithm, what you're trying to do is choreograph a pattern of interference where for each wrong answer, the contributions to its amplitude cancel each other out, whereas the contributions to the amplitude of the right answer reinforce each other. The trouble is, it's only for a few practical problems that we know how to do that in a way that vastly outperforms the best known classical algorithms.

What are those problems? Here, for all the theoretical progress that's been made in these past decades, I'm going to give the same answer in 2024 that I would've given in 1998. Namely, there's the simulation of chemistry, materials, nuclear physics, or anything else where many-body quantum effects matter. This was Feynman's original application from 1981, but probably still the most important one commercially. It could plausibly help with batteries, drugs, solar cells, high-temperature superconductors, all kinds of other things, maybe even in the next few years.

And then there's breaking public-key cryptography, which is not commercially important, but is important for other reasons well-known to everyone here.

And then there's everything else. For problems in optimization, machine learning, finance, and so on, there's typically a Grover's speedup, but that of course is "only" a square root and not an exponential, which means that it will take much longer before it's relevant in practice. And one of the earliest things we learned in quantum computing theory is that there's no "black-box" way to beat the Grover speedup. By the way, that's also relevant to breaking cryptography — other than the subset of cryptography that's based on abelian groups and can be broken by Shor's algorithm or the like. The centerpiece of my PhD thesis, twenty years ago, was the theorem that you can't get more than a Grover-type polynomial speedup for the black-box problem of finding collisions in cryptographic hash functions.

So then what remains? Well, there are all sorts heuristic quantum algorithms for classical optimization and machine learning problems — QAOA (Quantum Approximate Optimization Algorithm), quantum annealing, and so on — and we can hope that sometimes they'll beat the best classical heuristics for the same problems, but it will be trench warfare, not just magically speeding up everything. There are lots of quantum algorithms somehow inspired by the HHL (Harrow-Hassidim-Lloyd) algorithm for solving linear systems, and we can hope that some of those algorithms will get exponential speedups for end-to-end problems that matter, as opposed to problems of transforming one quantum state to another quantum state. We can of course hope that new quantum algorithms will be discovered. And most of all, we can look for entirely new problem domains, where people hadn't even considered using quantum computers before—new orchards in which to pick low-hanging fruit. Recently, Shih-Han Hung and I, along with others, have proposed using current QCs to generate cryptographically certified random numbers, which could be used in post-state cryptocurrencies like Ethereum. I'm hopeful that people will find other protocol applications of QC like that one — "proof of quantum work." [Another major potential protocol application, which Dan Boneh brought up after my talk, is quantum one-shot signatures.]

Anyway, taken together, I don't think any of this is too good to be true. I think it's genuinely good and probably true!

Reason #2

A second reason people didn't take seriously that QC was actually going to happen was the general thesis of technological stagnation, at least in the physical world. You know, maybe in the 40s and 50s, humans built entirely new types of machines, but nowadays what do we do? We issue press releases. We make promises. We argue on social media.

Nowadays, of course, pessimism about technological progress seems hard to square with the revolution that's happening in AI, another field that spent decades being ridiculed for unfulfilled promises and that's now fulfilling the promises. I'd also speculate that, to the extent there is technological stagnation, most of it is simply that it's become really hard to build new infrastructure—high-speed rail, nuclear power plants, futuristic cities—for legal reasons and NIMBY reasons and environmental review reasons and Baumol's cost disease reasons. But none of that really applies to QC, just like it hasn't applied so far to AI.

Reason #3

A third reason people didn't take this seriously was the sense of "It's been 20 years already, where's my quantum computer?" QC is often compared to fusion power, another technology that's "eternally just over the horizon." (Except, I'm no expert, but there seems to be dramatic progress these days in fusion power too!)

My response to the people who make that complaint was always, like, how much do you know about the history of technology? It took more than a century for heavier-than-air flight to go from correct statements of the basic principle to reality. Universal programmable classical computers surely seemed more fantastical from the standpoint of 1920 than quantum computers seem today, but then a few decades later they were built. Today, AI provides a particularly dramatic example where ideas were proposed a long time ago—neural nets, backpropagation—those ideas were then written off as failures, but no, we now know that the ideas were perfectly sound; it just took a few decades for the scaling of hardware to catch up to the ideas. That's why this objection never had much purchase by me, even before the dramatic advances in experimental quantum error-correction of the last year or two.

Reason #4

A fourth reason why people didn't take QC seriously is that, a century after the discovery of QM, some people still harbor doubts about quantum mechanics itself. Either they explicitly doubt it, like Leonid Levin, Roger Penrose, or Gerard 't Hooft. Or they say things like, "complex Hilbert space in 2^n dimensions is a nice mathematical formalism, but mathematical formalism is not reality"—the kind of thing you say when you want to doubt, but not take full intellectual responsibility for your doubts.

I think the only thing for us to say in response, as quantum computing researchers—and the thing I consistently have said—is man, we welcome that confrontation! Let's test quantum mechanics in this new regime. And if, instead of building a QC, we have to settle for "merely" overthrowing quantum mechanics and opening up a new era in physics—well then, I guess we'll have to find some way to live with that.

Reason #5

My final reason why people didn't take QC seriously is the only technical one I'll discuss here. Namely, maybe quantum mechanics is fine but fault-tolerant quantum computing is fundamentally "screened off" or "censored" by decoherence or noise—and maybe the theory of quantum fault-tolerance, which seemed to indicate the opposite, makes unjustified assumptions. This has been the position of Gil Kalai, for example.

The challenge for that position has always been to articulate, what is true about the world instead? Can every realistic quantum system be simulated efficiently by a classical computer? If so, how? What is a model of correlated noise that kills QC without also killing scalable classical computing?—which turns out to be a hard problem.

In any case, I think this position has been dealt a severe blow by the Random Circuit Sampling quantum supremacy experiments of the past five years. Scientifically, the most important thing we've learned from these experiments is that the fidelity seems to decay exponentially with the number of qubits, but "only" exponentially — as it would if the errors were independent from one gate to the next, precisely as the theory of quantum fault-tolerance assumes. So for anyone who believes this objection, I'd say that the ball is now firmly in their court.

So, if we accept that QC is on the threshold of becoming real, what are the next steps? There are the obvious ones: push forward with building better hardware and using it to demonstrate logical qubits and fault-tolerant operations on them. Continue developing better error-correction methods. Continue looking for new quantum algorithms and new problems for those algorithms to solve.

But there's also a less obvious decision right now. Namely, do we put everything into fault-tolerant qubits, or do we continue trying to demonstrate quantum advantage in the NISQ (pre-fault-tolerant) era? There's a case to be made that fault-tolerance will ultimately be needed for scaling, and anything you do without fault-tolerance is some variety of non-scalable circus trick, so we might as well get over the hump now.

But I'd like to advocate putting at least some thought into how to demonstrate a quantum advantage in the near-term. Thay could be via cryptographic protocols, like those that Kahanamoku-Meyer et al. have proposed. It could be via pseudorandom peaked quantum circuits, a recent proposal by me and Yuxuan Zhang—if we can figure out an efficient way to generate the circuits. Or we could try to demonstrate what William Kretschmer, Harry Buhrman, and I have called “quantum information supremacy,” where, instead of computational advantage, you try to do an experiment that directly shows the vastness of Hilbert space, via exponential advantages for quantum communication complexity, for example. I’m optimistic that that might be doable in the very near future, and have been working with Quantinuum to try to do it.

On the one hand, when I started in quantum computing 25 years ago, I reconciled myself to the prospect that I’m going to study what fundamental physics implies about the limits of computation, and maybe I’ll never live to see any of it experimentally tested, and that’s fine. On the other hand, once you tell me that there is a serious prospect of testing it soon, then I become kind of impatient. Some part of me says, let’s do this! Let’s try to achieve forthwith what I’ve always regarded as the #1 application of quantum computers, more important than codebreaking or even quantum simulation: namely, disproving the people who said that scalable quantum computing was impossible.

15.Microsoft's quantum-resistant cryptography is here

by Aabha Thipsay
https://techcommunity.microsoft.com/t5/security-compliance-and-identity/microsoft-s-quantum-resistant-cryptography-is-here/ba-p/4238780

How we are preparing for the future of cryptography

Cryptography is the science of securing information from unauthorized access or modification. It is essential for protecting the privacy and integrity of data in the digital world. However, cryptography is not static. It evolves with advances in mathematics, computer science, and technology. One of the biggest challenges that cryptography faces today is the future threat from substantially more powerful quantum computers

Quantum computing leverages the properties of quantum physics, such as superposition and entanglement, operations that are impossible or impractical for classical computers. While quantum computers have the potential to help us solve some of the most complex problems in science, engineering, and medicine, they also have the potential to upend public-key algorithms, which form the foundation of today’s encryption and security for most existing information and communication technology products.

In an earlier blog post we explored how quantum computing could disrupt the most commonly used asymmetric algorithms, such as Rivest–Shamir–Adleman (RSA) and Elliptic Curve Cryptography (ECC), and why symmetric cryptography largely withstands quantum threats. While a capable enough quantum computer to break public-key cryptography is still in the future, threat actors are preparing today. There are increasing concerns related to attackers recording data now with a view to being able to decrypt it later when quantum computers are sufficiently mature – in so-called “Harvest-now, Decrypt-later” attacks.

To address this challenge, researchers have been developing post-quantum cryptography (PQC) algorithms that are resistant to quantum attacks. PQC is based on mathematical problems that are hard for both classical and quantum computers. PQC algorithms offer a promising solution for the future of cryptography, but they also come with some trade-offs. For example, these typically require larger key sizes, longer computation times, and more bandwidth than classical algorithms. Therefore, implementing PQC in real-world applications requires careful optimization and integration with existing systems and standards.

Microsoft is a key participant in and contributor to the National Institute of Standards and Technology (NIST) Post-Quantum Cryptography Standardization Process, which announced the first set of PQC algorithms which can be used by public and private sectors. Microsoft Research's work on PQC includes several proposals for PQC algorithms in collaboration with academics and industry partners, and we have provided feedback and analysis on other submissions. Microsoft is also a core member and contributor in Open Quantum Safe (OQS) and National Cybersecurity Center of Excellence (NCCoE). We are also actively engaged in the Internet Engineering Task Force (IETF) to define standard interoperable ways to use PQC algorithms for safeguarding communication. This step is crucial before we see mainstream PQC adoption in software products and services across the industry.

Microsoft has developed a comprehensive strategy to support quantum resistance, acknowledging the significant impact quantum computing may have on existing public-key encryption methods. To address this, we established the Microsoft Quantum Safe Program (QSP), which unifies and accelerates all quantum-safe initiatives across the company from both technical and business perspectives. The goal of QSP is to achieve quantum readiness by integrating PQC algorithms and other security measures into Microsoft products, services, and infrastructures. Additionally, QSP is dedicated to supporting and empowering our customers, partners, and ecosystems as they work toward their own quantum-safe transitions.

Introducing PQC Algorithms in SymCrypt

At Microsoft, we strive to provide our customers with top security solutions for their data and communications. That is why we are proud to announce that we have begun releasing support for post-quantum algorithms in SymCrypt, Microsoft's open-source core cryptographic library. Last week we published a SymCrypt update that includes the ML-KEM and XMSS algorithms, to be followed in the coming months with additional algorithms described below. This is a major milestone in our journey to prepare for the quantum era and to help protect our customers from future quantum threats.

SymCrypt is Microsoft's main cryptographic library used in products and services such as Azure, Microsoft 365, Windows 11, Windows 10, Windows Server 2025, Windows Server 2022, Azure Stack HCI, and Azure Linux. These products and services use SymCrypt to provide cryptographic security for scenarios such as email security, cloud storage, web browsing, remote access, and device management. SymCrypt offers a consistent interface for encryption, decryption, signing, verification, hashing, and key exchange using both symmetric and asymmetric algorithms. It is built to be fast, secure, and portable across multiple platforms and architectures. In Windows operating systems, the SymCrypt cryptographic library is embedded in the Cryptographic Primitives Libraries (bcryptprimitives.dll and cng.sys) which have gone through multiple FIPS 140 validations; SymCrypt is also going through a FIPS 140 validation as a cryptographic module for Linux-based operating systems. Microsoft maintains an active commitment to meeting the requirements of the FIPS 140 standard. We will continue to update and pursue evaluations for our products and services as standards evolve to support PQC algorithms.

With NIST releasing an initial group of finalized post-quantum encryption standards, we are excited to bring these into SymCrypt, starting with ML-KEM (FIPS 203, formerly Kyber), a lattice-based key encapsulation mechanism (KEM). In the coming months, we will incorporate ML-DSA (FIPS 204, formerly Dilithium), a lattice-based digital signature scheme and SLH-DSA (FIPS 205, formerly SPHINCS+), a stateless hash-based signature scheme.

In addition to the above PQC FIPS standards, in 2020 NIST published the SP 800-208 recommendation for stateful hash-based signature schemes which are also resistant to quantum computers. As NIST themselves called out, these algorithms are not suitable for general use because their security depends on careful state management, however, they can be useful in specific contexts like firmware signing. In accordance with the above NIST recommendation we have added eXtended Merkle Signature Scheme (XMSS) to SymCrypt, and the Leighton-Micali Signature Scheme (LMS) will be added soon along with the other algorithms mentioned above.

PQC algorithms have been meticulously chosen by NIST to offer high security, performance, and compatibility. They have been fine-tuned for efficiency in speed and size and have gone through rigorous tests for security and robustness. Efforts are ongoing within multiple industry standards organizations to ensure these algorithms are adopted into and compatible with existing standards and protocols such as Transport Layer Security (TLS), Secure Socket Shell (SSH), and Internet Protocol Security (IPSec), and that they can operate in hybrid mode alongside classical algorithms like RSA, Elliptic Curve Diffie–Hellman (ECDH), and Elliptic Curve Digital Signature Algorithm(ECDSA). As PQC standards develop, we will seek to incorporate additional algorithms into SymCrypt to maintain compliance, security, and compatibility.

The recommended path for leveraging SymCrypt is via Cryptography API: Next Generation (CNG) on Windows, while on Linux there are several options: direct use of SymCrypt APIs, the SymCrypt engine for OpenSSL (SCOSSL), or the SymCrypt Rust Wrapper. Over the coming months, these layers will add support for PQC algorithms, giving our customers the ability to experiment in their own environments and applications.

The use of PQC algorithms to secure TLS communications is an area experiencing rapid development. Although the finalization of NIST algorithms represents a key milestone in this advancement, two critical standards are required for widespread adoption: quantum safe key exchange and quantum safe signature authentication. We are working closely with the IETF to develop and standardize quantum-safe key exchange and authentication for TLS and other IETF protocols. As these standards get finalized, we will make these available through the Windows TLS stack (Schannel) and SymCrypt engine for OpenSSL on Linux.

PQC algorithms are relatively new, and it is prudent not to consider the initial generation of PQC algorithms as the definitive solution but rather view this as an evolving field. This underscores the importance of "Crypto Agility" which involves designing solutions to be resilient to the use of different algorithms and/or upgradable to use future algorithms as the PQ standards evolve. Recognizing this, Microsoft is a strong advocate of building solutions which are crypto agile, as well as deploying PQC solutions which make use of a hybrid PQ mode of operation. In time, we expect a shift towards pure PQ deployments, as PQ algorithms and standards mature.

Adding post-quantum algorithm support to the underlying crypto engine is the first step towards a quantum safe world. As we enable support for PQC in additional system components and applications, we will see services light up end-to-end scenarios protected by PQC while also giving our customers the option to experiment with and adopt it in their own environments and applications.

Start your PQC transition journey

The transition to PQC is a complex, multi-year and iterative process, which requires attention and careful planning. One of the first steps that we recommend organizations to take is creating an inventory of cryptographic assets in use. By that, organizations can better understand the scope of the effort and establish a risk-based plan for their PQC transition.

Also, we recommend familiarizing the organization with the PQC algorithms and approaches for implementations.

Microsoft is here to assist its customers, partners and ecosystems in navigating their transition to quantum safety and optimizing safety in the quantum era. Fill out this questionnaire to get started with Microsoft.

Conclusion

PQC algorithm support in SymCrypt is a significant step forward in our efforts to prepare for the quantum era, and to help protect our customers from future quantum threats. We are excited to share this update with you and to hear your feedback and suggestions. We also look forward to collaborating with the research community, industry, and standards bodies to advance the state of the art in post-quantum cryptography and to make it more widely available and adopted. By working together, we can maintain cryptography as a strong method for protecting information in the digital age.

16.Quantum Communication – Different Visions of the Quantum Internet

by Amara Graps
https://quantumcomputingreport.com/quantum-communication-different-visions-of-the-quantum-internet/

The BB84 protocol

In 1984, Brassard and Bennett presented their 4-page, conference paper, which described the BB84 algorithm, thus fixing in history the first quantum cryptography protocol for the Quantum Key Distribution (QKD) concept. The BB84 algorithm is still in-use today. See BB84's roots to physics principles in the first words of their Abstract:

> *When elementary quantum systems, such as polarized photons, are used to transmit digital information, the uncertainty principle gives rise to novel cryptographic phenomena unachievable with traditional transmission media […] –Brassard and Bennett*

Given that QKD distributes cryptographic keys produced by quantum techniques across terrestrial and satellite-based global communications networks, QKD will undoubtedly play a role in the future Quantum Internet. Indeed, Stephanie Wehner, David Elkouss, and Ronald Hanson, published in 2018, a *science* review: Quantum internet: A vision for the road ahead with QKD at its trusted nodes.

Trusted nodes, secured by QKD protocols, are also inside GQI's perspective of the evolution of the Quantum Internet. You can find Global Quantum Intelligence's far-ranging vision of the deployment of the Quantum Internet in GQI's Quantum Safe Outlook Report. See the next Figure.

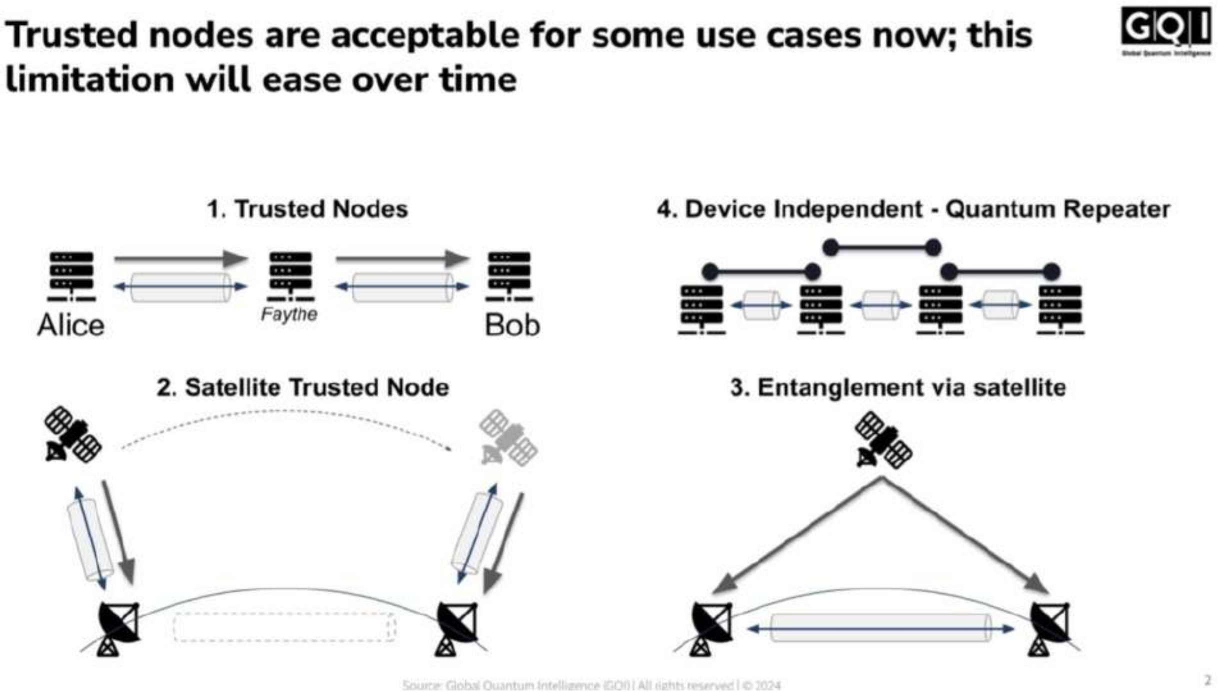

Trusted nodes are acceptable for some use cases now; this limitation will ease over time

GQI
1. Trusted Nodes
Alice
Faythe
Bob
4. Device Independent - Quantum Repeater
2. Satellite Trusted Node
3. Entanglement via satellite
2

Source: Global Quantum Intelligence (GQI) | All rights reserved | © 2024

Figure. An illustration of Trusted nodes in GQI's Quantum Safe Outlook Report[1].

QKD versus PQC

Recent news about NIST's selection of post-quantum cryptography (PQC) algorithms, highlights divergent views of "quantum secure" technology: one ruled by physics (QKD) and one by mathematical approaches (PQC). IEEE Spectrum recently interviewed GQI's Doug Finke, to explain these quantum technology differences. In brief:

> *Theoretically, people cannot predict that these PQC algorithms won't be broken at some point. On the other hand, QKD- there are theoretical arguments based on quantum physics that you can't break a QKD network.*

These different views, feed into a variety of approaches for building the Quantum Internet.

Competing visions for the Quantum Internet

Our Quantum Safe Outlook Report and corresponding *Quantum Safe State of Play Presentation* available for GQI customers, explains the journey to build the Quantum Internet, with the Network-led approaches, the Compute-led approaches, and our vision which is agnostic to both, because these are intermediate steps.

> *[…] investors, companies and governments are faced with overlapping opportunities across four of the great deep tech sectors: cryptography, quantum, AI and space. It's a challenge to avoid being locked into any one technology or narrow field of expertise. It's a challenge to evaluate the trade-off of short-, medium- and long-term revenue opportunities; doubly so when we consider the interplay of*

[1] Quantum Safe Outlook Report. Despite the challenges, the transition to quantum-safe cryptography is essential to ensure the security of our digital infrastructure in the future. Organizations that start planning now will be well-positioned to meet this challenge.

> *economic vs geopolitical factors that will likely influence the development of the wider sector. […] Some are focused on how to deploy quantum safe cryptography to build a quantum resistant Internet. Some are focused on building 'prepare & measure' quantum networks using trusted nodes to create quantum-enabled networks. Others envision the true entanglement-based Quantum Internet.*

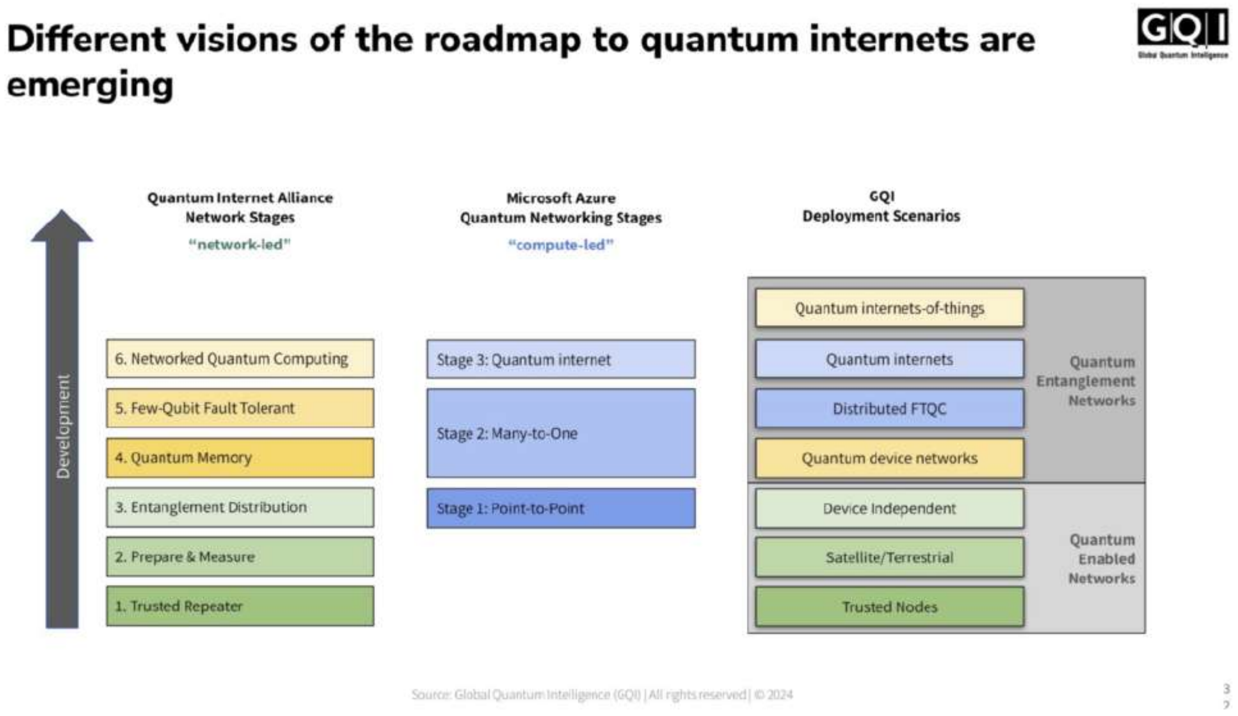

Different visions of the roadmap to quantum internets are emerging

Quantum Internet Alliance Network Stages
"network-led"

6. Networked Quantum Computing
5. Few-Qubit Fault Tolerant
4. Quantum Memory
3. Entanglement Distribution
2. Prepare & Measure
1. Trusted Repeater

Microsoft Azure Quantum Networking Stages
"compute-led"

Stage 3: Quantum internet
Stage 2: Many-to-One
Stage 1: Point-to-Point

GQI Deployment Scenarios

Quantum internets-of-things
Quantum internets
Distributed FTQC
Quantum device networks
Quantum Entanglement Networks

Device Independent
Satellite/Terrestrial
Trusted Nodes
Quantum Enabled Networks

Source: Global Quantum Intelligence (GQI) | All rights reserved | © 2024

32

Figure. An illustration of Trusted nodes in GQI's Quantum Safe Outlook Report.

QCR's article: Ingredients for a Killer App illustrates some of the needed quantum devices, quantum networks and quantum protocols in the Quantum Internet journey. Our key question is:

> *How important are the intermediate steps are in winning the race to build an entanglement-based end goal?*

Back to Physics

If you explore some of the Quantum Internet articles with Connected Papers, you may discover that QKD has direct linkages to Einstein and his 1935 public quantum entanglement discussions with Schroedinger and Bohr**.** Moreover, you can use the (often-misused) academic citation metrics of today to judge the winne*r* of their 1935 debates: Einstein, Podolsky, and Rosen (24695 Citations), Schroedinger (5385 Citations), and Bohr (3265 Citations). Who won? Einstein, of course.

17.Quantum Technology Algorithms – The Ansatz Zoo

by Amara Graps
https://quantumcomputingreport.com/quantum-technology-algorithms-the-ansatz-zoo/

The most astonishing feature of the quantum computing field to newcomers is the variety of methods ('modalities') to make a qubit. With the state-of-the-art today, what evidence, if any, is there for particular applications that are better suited to particular quantum computing modalities? GQI's Doug Finke says in a recent Laser Focus article :

> *We aren't at the point yet where we can definitely say which quantum applications will be able to provide commercially useful results on which machines. But the one thing that makes us optimistic is the diversity of innovative approaches and rapid advances organizations are making in both hardware and software to get us to the point where the systems can be used for quantum production for useful applications.*

We know the answer in a general sense: from quantum hardware that is gate-based, versus annealing, versus bosonic (photonic). Gates refers to logic gates using the computational method we are familiar from classical computing. Quantum annealers use adiabatic models as a paradigm in the device's quantum fluctuations to find the global minimum of a given objective function. Gaussian boson sampling (GBS) is how photonic quantum devices sample from the output distribution of a specific linear optical circuit.

In today's NISQ era, we can also learn which hybrid quantum technology modality is suited for which problem in an algorithm sense, by the choice of ansatz to solve the problem. Ansatz (from German Ansätze): the first guess or initial condition, in quantum computing, refers to a trial wavefunction or trial state used as a starting point for approximations or optimizations. We introduced the use of ansatzes in Part 3 of The Many Faces of Hybrid Classical-Quantum Computing. In GQI's Quantum Algorithms Framework, the ansatzes fit inside the *Notable Early NISQ Heuristics* and *Variable Quantum Algorithm* labels in the next figure, inside the red rectangle.

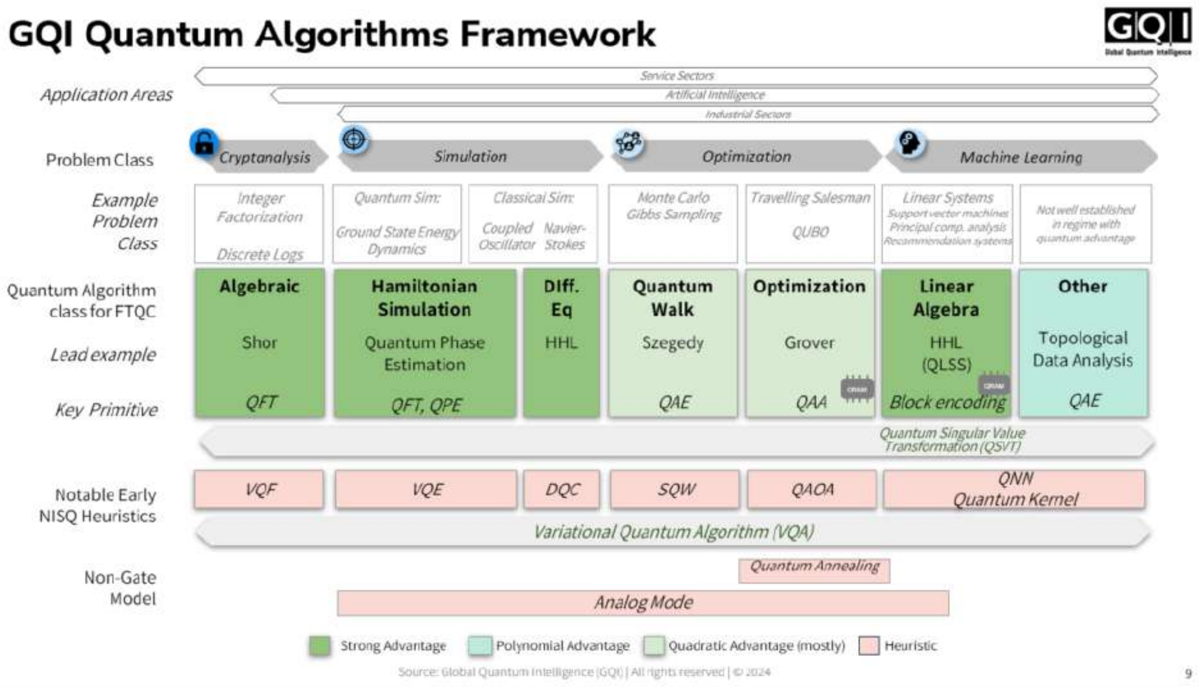

GQI Quantum Algorithms Framework

Application Areas: *Service Sectors*; *Artificial Intelligence*; *Industrial Sectors*

Problem Class	*Cryptanalysis*	*Simulation*		*Optimization*		*Machine Learning*	
Example Problem Class	*Integer Factorization* *Discrete Logs*	*Quantum Sim:* *Ground State Energy Dynamics*	*Classical Sim:* *Coupled Oscillator* *Navier-Stokes*	*Monte Carlo Gibbs Sampling*	*Travelling Salesman* *QUBO*	*Linear Systems* *Support vector machines* *Principal comp. analysis* *Recommendation systems*	*Not well established in regime with quantum advantage*
Quantum Algorithm class for FTQC	**Algebraic**	**Hamiltonian Simulation**	**Diff. Eq**	**Quantum Walk**	**Optimization**	**Linear Algebra**	**Other**
Lead example	Shor	Quantum Phase Estimation	HHL	Szegedy	Grover	HHL (QLSS)	Topological Data Analysis
Key Primitive	*QFT*	*QFT, QPE*		*QAE*	*QAA*	*Block encoding*	*QAE*
						Quantum Singular Value Transformation (QSVT)	
Notable Early NISQ Heuristics	*VQF*	*VQE*	*DQC*	*SQW*	*QAOA*	*QNN* *Quantum Kernel*	
	Variational Quantum Algorithm (VQA)						
Non-Gate Model					*Quantum Annealing*		
		Analog Mode					

🟩 Strong Advantage 🟦 Polynomial Advantage 🟩 Quadratic Advantage (mostly) 🟥 Heuristic

Source: Global Quantum Intelligence (GQI) | All rights reserved | © 2024

Figure. Slide from Presentation *GQI Quantum Software State of Play* showing GQI's own Algorithm Framework that incorporates algorithmic themes for different hardware eras.

Ansatz Research

Ansatzes are a key to answer this application | modality question. The answer will become clear in the near future as the hybrid quantum-classical Use Cases continue to be developed. Today I will show you the zoo of ansatz research, where a quick search of ArXiv for all years returns ~400 ansatz research papers, with ~360 of those research articles in the last five years.

The best education I've found so far for clarity to understand the ansatz zoo is Sophia Economou's May 7, 2021, 30 min. presentation for the Quantum Research Seminars Toronto : How to create a good ansatz for variational quantum algorithms She presents three different classes of ansatzes, with a theme of problem focus. see the next figure. In the left side of the arrow in red are the Hardware-Efficient Ansatzes (HEA), which are agnostic for the problem, and towards the right, is the highest degree of problem-tailoring. In her seminar, she explains the forms of the ansatz function and how further to develop them while customizing it for the problem at hand.

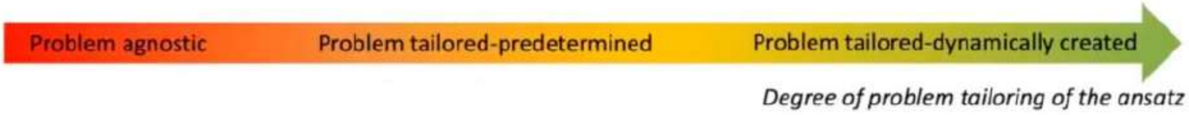
Problem agnostic
Problem tailored-predetermined
Problem tailored-dynamically created
Degree of problem tailoring of the ansatz

Figure. Schematic for classes of ansatz development from Sophia Economou's May 7, 2021, 30 min. presentation for the Quantum Research Seminars Toronto: How to create a good ansatz for variational quantum algorithms

The Ansatz Library, circa 2019

Cao et al, 2019's 194-pg tome Quantum Chemistry in the Age of Quantum Computing includes a table (Table 5) that identified particular chemistry problems with their ansatzes and quantum computing architecture (i.e. modality). This was the paper which inspired my thinking that we have some evidence for preferences for quantum hardware modalities to particular hybrid problems. The readable research article maps chemistry Hamiltonians to qubit Hamiltonians and describes in 25p a handful of ansatzes for state preparation in chemistry problems.

The Ansatz Library, circa 2022

Several years later, Tilley et al, 2022 in their 156-pg tome: The Variational Quantum Eigensolver: a review of methods and best practices provided a VQE Survey with practical advice and descriptions of the ansätzes, indicating a good growth of the field. The 'Fixed Structure' classification in Tilley's work would lie on the left of Economou's figure Arrow as 'Problem Agnostic' and the 'Adaptative structure' would lie in the middle of her figure Arrow.

Fixed structure ansatz

- Hardware-efficient ansatz
- The Unitary Coupled Cluster (UCC)
- Symmetry-preserving methods
- Hamiltonian Variational Ansatz

Adaptative structure ansatz

- Iterative ansatz growth methods (ADAPT-VQE and extensions)
- Iterative Hamiltonian dressing (iterative Qubit Coupled Cluster (iQCC) and extensions)

The Ansatz Library, circa 2024

In 2024, the ansatz zoo expanded significantly. Blekos, et a.l, 2024 in *A review on Quantum Approximate Optimization Algorithm and its variants* provide at least 14 different classes, where the Hardware-efficient (now called Hardware-specific) is one of the 14. It's a detailed 67-pg review with 10-pgs of ansatz descriptions plus practical advice, including this Section 6.1: *Which QAOA Ansatz Variant Should I Use for My Problem?* For quantum computer scientists today, who are implementing real Use Cases on hybrid quantum-classical hardware, this might be one of your most valuable papers to have in your library.

As we've just scratched the surface of this deep topic, we'll be returning to the topic again in the future.

18.Homomorphic encryption pioneer on Apple, DARPA and cracking the code for mainstream adoption

by Jasper Hamill
https://www.thestack.technology/homomorphic-encryption-pioneer-on-apple-darpa-and-cracking-the-code-for-mainstream-adoption/

Roughly 15 years ago, a quiet revolution began. Following the publication of a paper about homomorphic encryption (HE), DARPA began funding research into this potentially game-checking tech, which allows data within encrypted files to be "accessed" and compute on it to happen without requiring full decryption — a potential game-changer for cross-jurisdictional data collaboration, among other use cases.

Now Apple has written and open-sourced an HE application in its programming language Swift and introduced it to its mobile operating system iOS, where it powers a Live Caller ID Lookup for caller ID and spam blocking services. This lets Apple send an encrypted query to a server that can provide information about a phone number without the server needing to know/store the number. It's another big step towards the mainstream – and top-tier standards bodies are also getting serious about the tech.

At DARPA, it's often argued that new encryption techniques take roughly 20 years to achieve adoption. Does this mean HE is about to have its moment in the sun?

Kurt Rohloff, Duality: "This starts in the 70s."

To find out, we spoke to Kurt Rohloff, who is now CTO and co-founder of Duality, but has been involved in the development of HE since 2010.

He spent nine years working at the defence giant Raytheon, where he was Principal Investigator on a multi-million dollar DARPA-funded R&D drive intended to reduce run-time of fully homomorphic encryption (FHE) and somewhat homomorphic encryption (SHE). Rohloff then went on to play a leading part in building one of the main open-source libraries in his space: OpenFHE.

"Like all good stories, this starts in the 70s," he tells The Stack. "One of the challenges was that when data was encrypted, you couldn't do anything with it. Think of the word 'crypt'. It's the grave. Yet there was always a vision of enabling collaboration on data or extracting knowledge from it without actually revealing the data itself."

"You can see the very obvious implications of the ability to compute on data and collaborate on sensitive data on the cloud without fear of it leaking," he adds. "I ran a team that was funded to provide the first prototypes of this technology and make it real."

What are the capabilities of homomorphic encryption?

At its current stage of adoption, HE is a "CISO or Chief Data Officer domain" rather than a consumer-grade technology, Rohloff says. Apple might just help to change this.

His firm has deployed HE to enable secure collaboration and anonymised data-sharing across clinical institutions without the need for data-sharing agreements. Mastercard has also used its tech to exchange data across borders.

A major reason for the growing maturity of this form of encryption is the sheer increase in compute horsepower that has taken place over the past decade and a half, enabling Rohloff's team to achieve "Moore's Law-style performance improvements".

"So, for example, the team that I've been running has been improving performance by an order of magnitude every six months," he reveals.

"The underlying compute model of homomorphic encryption looks a lot like operations on very, very long vectors," he adds. "These are supported very well by highly parallel processors like GPUs or FPGAs.

"There's been a bunch of work in organisations including DARPA to design custom ASICs and the growth of modern compute from commercial, off-the-shelf technologies like GPUs, FPGAs and ASICs and things like that. This stuff is just getting a lot faster."

At the same time, the open-source community has reliably generated software that steadily improves upon previous generations.

"It's not all about the hardware," Rohloff points out. "There have been tremendous improvements in the software. This is the value of open source in particular."

The road has also been smoothed by standards and regulatory acceptance of the technology. NIST has considered it as part of a project focused on privacy-enhancing cryptography. The UK Information Commissioner's Office also included homomorphic encryption in its guidance around deploying privacy-enhancing tech (PET).

Encryption in the enterprise

For enterprises, the advantage of opening up encrypted data to safe, secure and resource-light collaboration extends beyond security.

"Like Apple, what we do as a company is deploy on legacy environments, which shows the real value-add of homomorphic encryption," the Duality co-founder adds. "It becomes efficient to deploy on what have traditionally been less secure environments, like commercial clouds, with a high degree of security - including post-quantum security and protection against nation-state level attacks. And this will run on the kind of Dell servers you could buy with your credit card and have shipped the next day."

As a strong supporter of the open source ethos, Rohloff says that Apple is "pulling from the same well" as his firm but insists he "does not want to take any credit."

"When Apple is doing this work, it's their thing. Several of the staff members over there have been collaborators with our open-source team before they went to Apple. The underlying protocol that Apple uses is one of the core protocols in our open source library, although it's not the same implementation and Apple has its own library."

We conclude our conversation by asking about some of the lessons learned from working at DARPA - which is also famed for its open source work.

Rohloff offers an insight that "sounds a little tongue in cheek and pithy, but is also very sincere at the same time".

"If you want folks to use advanced tech, make it as boring as possible," he says. "Make it usable. Then, once the kind of technology is showing value, you can start adding features and capabilities.

"If you want to get over that first hump – the valley of death – to get your tech from research into operational use, make it boring, get it into people's hands and go from there."

19.Department of Commerce Implements Controls on Quantum Computing and Other Advanced Technologies Alongside International Partners

https://www.bis.gov/press-release/department-commerce-implements-controls-quantum-computing-and-other-advanced#agency

The U.S. Commerce Department's Bureau of Industry and Security (BIS) published an interim final rule (IFR) today implementing controls on critical and emerging technologies that have reached broad technical agreement among our international partners. This IFR includes controls related to quantum computing, semiconductor manufacturing, and other advanced technologies. Today's action strengthens our international relationships with like-minded countries and ensures that U.S. export controls keep pace with rapidly advancing technologies that pose serious threats to our national security when in the wrong hands.

"Today's action ensures our national export controls keep step with rapidly evolving technologies and are more effective when we work in concert with international partners," **said Alan Estevez, Under Secretary for the Bureau of Industry and Security**. "Aligning our controls on quantum and other advanced technologies makes it significantly more difficult for our adversaries to develop and deploy these technologies in ways that threaten our collective security."

"The most effective way to protect our national security is to develop and coordinate our controls alongside like-minded partners, and today's actions demonstrate our flexibility in how we craft such controls to achieve our national security objective," **said Assistant Secretary of Commerce for Export Administration Thea D. Rozman Kendler**, adding "Not only are we adopting new emerging technology controls with this rule, we are also building increased agility into our system with a new license exception for trusted partners."

As critical technologies with military applications continue to emerge and evolve, there is an increased need to regulate their movement to ensure these items are not used for purposes contrary to U.S. national security or foreign policy.

In today's IFR, BIS is implementing worldwide export controls on specific types of items, including:

- **Quantum Computing Items**: quantum computers, related equipment, components, materials, software, and technology that can be used in the development and maintenance of quantum computers.
- **Advanced Semiconductor Manufacturing Equipment**: tools and machines that are essential for the production of advanced semiconductor devices.
- **Gate All-Around Field-Effect Transistor (GAAFET) Technology**: technology that produces or develops high-performance computing chips that can be used in supercomputers.
- **Additive Manufacturing Items:** Equipment, components and related technology and software designed to produce metal or metal alloy components.

Several like-minded countries have already announced or implemented new national controls for export of items under their jurisdiction related to quantum computing and advanced semiconductor manufacturing. We anticipate additional countries will implement similar controls soon. This unified approach, built on a foundation of shared values and security interests, enhances the effectiveness of our controls. Continued international collaboration in advanced technologies is paramount to national security. For this reason, this IFR establishes a new **License Exception Implemented Export Controls (IEC)** so that countries may meet the terms of IEC by implementing equivalent national controls which would eliminate the need to submit license applications for those items, thereby fostering innovation opportunities among implementing countries. BIS has posted on its website HERE a table that lists IEC eligibility status for countries and items.

Furthermore, this IFR implements certain exclusion clauses for deemed exports (i.e., the sharing or release of controlled technology or source code to a foreign person within the U.S.) and reexports to avoid disrupting the ongoing research and development of these critical and emerging technologies. BIS has also added a general license for deemed exports/reexports of certain technology/software, subject to annual reporting requirements, to provide the U.S. Government with necessary visibility and oversight for national security reasons. Additionally, there is a 60-day delayed compliance date for quantum items to certain destinations to allow for the submission of license applications and revisions to internal compliance procedures.

20.Post-Quantum Cryptography Coalition Publishes Comparison of International PQC Standards

by MITRE
https://www.mitre.org/news-insights/news-release/post-quantum-cryptography-comparison-international-pqc-standards

Growing to more than 125 participating cyber researchers from industry and academia, the global Post-Quantum Cryptography (PQC) Coalition published a comparison of PQC standards being defined by international government regulatory bodies. The coalition is helping ready the cyber community for the post-quantum transition, working in parallel with National Institute of Standards and Technology (NIST) and the National Cybersecurity Center of Excellence's PQC Migration Project to provide comprehensive assessment and guidance.

"NIST has approved its first three federal standards for PQC algorithms addressing cryptographic and digital signature schemes, allowing for broad deployment of PQC," said Wen Masters, vice president, cyber technologies, MITRE. "Through our collective efforts, we want the nation to take advantage of quantum opportunities and be prepared to defend itself from quantum-enabled threats to national and economic security."

Data encrypted online with today's methods can be harvested now for decryption later by an adversary with access to a cryptographically relevant quantum computer.

"The coalition comes at an important time," said Matt Mickelson, lead coordinator of the coalition and senior cyber principal for science and technology at MITRE. "These threats target not just personal, commercial, military, and intelligence data but also the digital signatures that identify trusted entities and contracts. National and international standards create a common framework for addressing such threats."

Celebrating its one-year anniversary this month, the PQC Coalition aims to accelerate the adoption of quantum-resistant methods in advance of any cyber threats posed by quantum computation. Spearheaded by MITRE and SandboxAQ, the coalition brings together tech industry leaders, cryptographic researchers, and engineers including from PQShield, IBM, Microsoft, Comcast, CryptoNext Security, PQSecure, QuSecure, SecureG and University of Waterloo.

The coalition surveyed international regulatory bodies that are defining the quantum-resistance requirements to be followed by technology vendors used in national or government security systems. They created a reference of international PQC requirements and are identifying alignment and misalignment areas, which could pose challenges for international vendor compliance and interoperability.

The global coalition's new website, www.pqcc.org, serves as a central resource for the broader cyber community, providing reference materials, tools, and guidance related to PQC. The coalition will soon publish new materials to help accelerate crypto inventories.

Coalition participants are collaborating on four workstreams regarding standards, education, implementation, and agility. Together, these collective efforts help to ready the nation and the world for the post-quantum transition.

Participants in the coalition will attend the Quantum World Congress in McLean, Virginia, next week and lead the Real-World PQC Summit co-located with Real World Crypto Symposium in Sofia, Bulgaria next March.

Organizations or individuals that wish to participate in the coalition can visit http://www.pqcc.org/ or contact pqcc-registration-list@mitre.org.

21.Quantum Technology Algorithm Trends – Tidying Up

by Amara Graps
https://quantumcomputingreport.com/quantum-technology-algorithm-trends-tidying-up/

45+ years ago, the availability of Apple I, II, Commodore, Kim and other personal computers, inspired a generation of hobbyists, who relished accessing the computer's 6502 chip. I worked in a group at NASA-Ames, where we kept a half-dozen extra Apple IIs for parts to maintain the operations of our 'portable' infrared astronomy, data-acquisition system. One of the astronomy team members wrote the code closest to the Apple II's hardware, which was in 6502 assembly-language; code that looked similar to this. 6502-Hobbyists should not throw out their old Apple II 6502 Assembly Language manuals, either, as they are apparently a collector's item.

Today, quantum computer programmers, upon accessing a QPU, require a similar, raw programming style, that reminds me of those old days. Along with the 'assembly-language code', i.e., circuit programming, are their collection of algorithms, and a variety of languages and programming environments. From GQI's Outlook Report: *Quantum Algorithms Outlook*

> *It is not possible to divorce algorithms from the specs of the hardware on which they are intended to run. Academics have been developing quantum algorithms for 30 years. Early quantum hardware varied from non-existent to very poor. […]*

Inside Presentation *GQI Quantum Software State of Play* you'll see GQI's Framework for how GQI views Algorithm classes. GQI has incorporated the best conceptual ideas from the literature and will be expanding the GQI tool in the future. I'd like to step through the classification thinking in the Algorithm community first and provide some useful references for you.

Algorithms Organization Challenge

The software tools for writing today's quantum algorithms have emerged as ubiquitously as the devices. When Kumar et al., 2022, in *Futuristic view of the Internet of Quantum Drones*, summarized the current programming tools for the quantum subset of Internet-of-Things, the authors required two pages of illustrations (see their Figures 14 and 15).

Dalzell et al.,2023 in *Quantum Algorithms: A Survey of Applications and End-to-End Complexities* addressed the algorithm organization challenge by providing an application focus woven in a Wiki format. Additional hyperlinks in the header of every page provide a means to quickly jump to sections and subsections. Nevertheless, for a document, which is 337 pages and ~1000 references, the hyperlinks provide a smart navigational solution, not a simplifying organizational structure.

Algorithms Re-Organization

This year, we see researchers approaching the algorithm maze with a purposeful classification mindset.

Arnault et al., 2024 in A typology of quantum algorithms looked at algorithmic trends in the NISQ era that might reveal the useful building blocks that can aid the discovery of quantum advantage. The authors categorized 133 quantum algorithms, based on their ability to solve basic mathematical problems, their practical applications, the primary subroutines they use, and other factors. Their aim was the discovery of relationships, dependencies or patterns among the algorithms and the criteria used.

Their dependence network offers a comprehensive picture of every layer of subroutines. It reveals also, to some degree, the history of ideas guiding the building of various algorithms, therefore providing an insight into the logical links among them. You can also download this figure and more from their paper from their Github site.

Arnault et al., 2024 saw changes in the development of the algorithms, with some mathematical classes and application domains becoming more and more prominent. Notably, there was an increase in the themes of “Operational Research,” “First-Principle Quantum Simulation,” and “Machine Learning & Data Science.” In Arnault’s Appendix A: Algorithms Classification table, there were 50 classified as potential candidates for using a NISQ processor, while the remaining candidates require a fully error-corrected (FTQC) machine. See also their Figure 6: Percentage of NISQ/ FTQC algorithms per Mathematical Class and Figure 7: Percentage of NISQ/ FTQC algorithms per Application Domain.

GQI’s Doug Finke says in a recent Laser Focus article that of Arnault et al.’s 133 algorithms, “a total of 50 were classified as potential candidates for using a NISQ processor, while the remaining candidates require a LSQ machine. It’s possible one of these NISQ algorithms can indeed provide a usable commercial quantum production before the FTQC quantum computers are available.”

You’ll recognize some of these same classifications in GQI’s own Algorithm Framework, where GQI has incorporated similar classifications and more to fit into NISQ and FTQC (fault-tolerant quantum computing) eras. GQI’s method of assessing quantum era and their uncertainties is here.

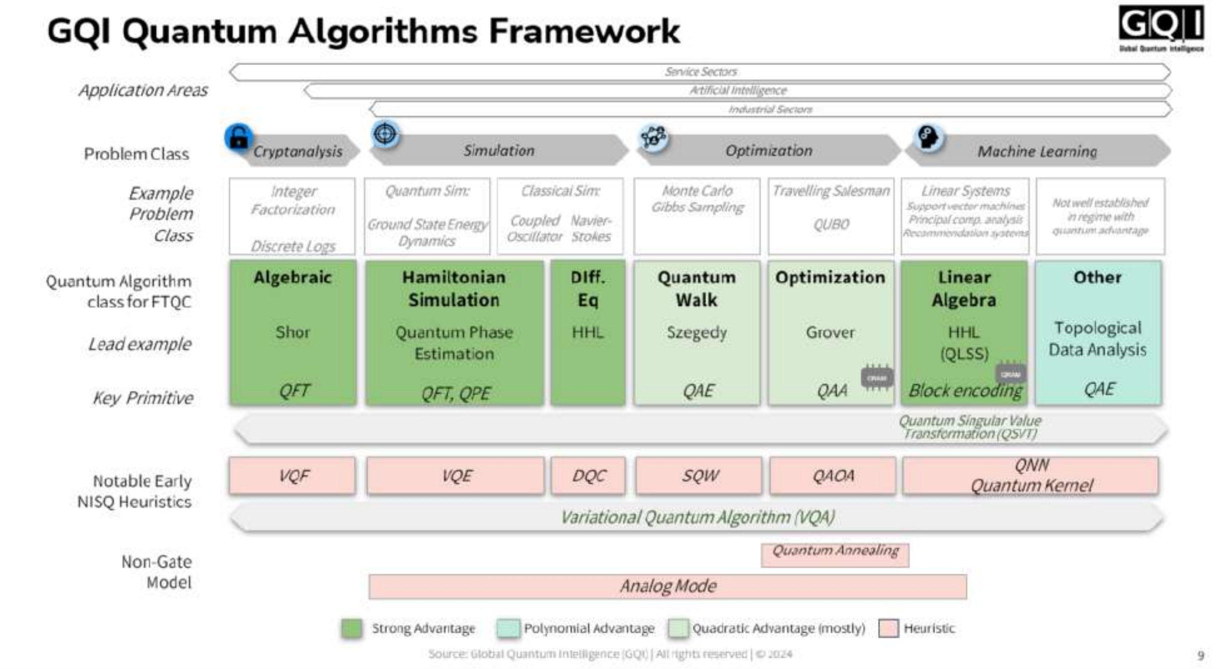

GQI Quantum Algorithms Framework

Application Areas: Service Sectors; Artificial Intelligence; Industrial Sectors

Problem Class	Cryptanalysis	Simulation		Optimization		Machine Learning	
Example Problem Class	Integer Factorization Discrete Logs	Quantum Sim: Ground State Energy Dynamics	Classical Sim: Coupled Oscillator Navier-Stokes	Monte Carlo Gibbs Sampling	Travelling Salesman QUBO	Linear Systems Support vector machines Principal comp. analysis Recommendation systems	Not well established in regime with quantum advantage
Quantum Algorithm class for FTQC	Algebraic	Hamiltonian Simulation	Diff. Eq	Quantum Walk	Optimization	Linear Algebra	Other
Lead example	Shor	Quantum Phase Estimation	HHL	Szegedy	Grover	HHL (QLSS)	Topological Data Analysis
Key Primitive	QFT	QFT, QPE		QAE	QAA	Block encoding	QAE
	Quantum Singular Value Transformation (QSVT)						
Notable Early NISQ Heuristics	VQF	VQE	DQC	SQW	QAOA	QNN Quantum Kernel	
	Variational Quantum Algorithm (VQA)						
Non-Gate Model				Quantum Annealing			
	Analog Mode						

🟩 Strong Advantage 🟦 Polynomial Advantage 🟩 Quadratic Advantage (mostly) 🟥 Heuristic

Source: Global Quantum Intelligence (GQI) | All rights reserved | © 2024

Slide from Presentation *GQI Quantum Software State of Play* showing GQI’s own Algorithm Framework that incorporates algorithmic themes for different hardware eras. *If you are interested to learn more, please don’t hesitate to contact info@global-qi.com.*

Now that you’ve seen useful classification of these algorithms, let’s see more tidying-up.

Programming Order: Abstracting the principles

Di Matteo et al., 2024, in An Abstraction Hierarchy Toward Productive Quantum Programming presented quantum computer programming in a similar level of abstraction as that for High Performance Computers; a three-layer abstraction hierarchy that incorporates:

1. a programming model is used to map algorithms onto software,
2. an execution model to understand how software executes, and
3. a hardware model to define an abstract machine that represents physical hardware.

Some Examples

- **Programming model**
    - High-level quantum subroutines
    - Ansatz with hybrid optimization loop
    - Unitary operations
    - Repeat-until-success
- **Execution Model**
    - Logical (non-hardware-native) circuits
    - Error mitigation
    - Error correction
    - QRASP
- **Hardware model**
    - Hardware-native circuits
    - Pulses

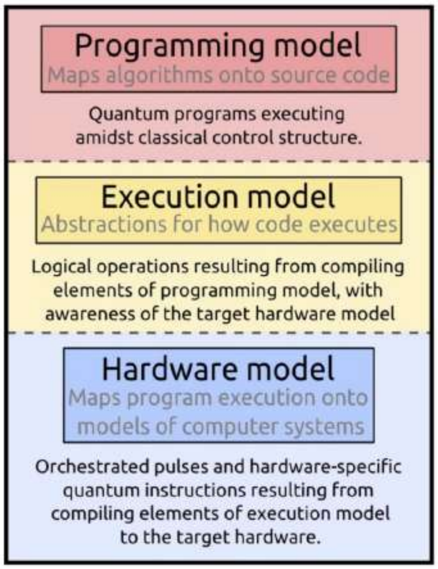
Programming model
Maps algorithms onto source code
Quantum programs executing amidst classical control structure.
Execution model
Abstractions for how code executes
Logical operations resulting from compiling elements of programming model, with awareness of the target hardware model
Hardware model
Maps program execution onto models of computer systems
Orchestrated pulses and hardware-specific quantum instructions resulting from compiling elements of execution model to the target hardware.

Figure. A proposed abstraction hierarchy derived from the three-level (Programming, Execution, Hardware) model. The state of today's quantum computing hardware and software can lead to overlapping, blurry boundaries from, Di Matteo et al., 2024, in An Abstraction Hierarchy Toward Productive Quantum Programming

The authors provide two case studies of eigenvalue estimation using their methodology,

VQE with error mitigation and
QPE (Quantum Phase Estimation) with error correction.

With a hope (and a call for discussion) that they've provided a layered complexity of abstractions, which correspond to the levels of software development today. These abstractions would support the evolution of software, without the need to rethink hardware.

22.PQShield builds NIST-ready PQC in silicon

by Matthew Stubbs
https://pqshield.com/pqshield-builds-nist-ready-pqc-in-silicon/

With the release of the first three NIST PQC standards, the world of post-quantum cryptography (PQC) is rapidly shifting from a focus on preparation, to a much sharper focus on compliance.

"For the first time, everyone has to look at how they adopt," explains Graeme Hickey, PQShield's VP, Engineering. "It's no longer a PoC or a research item; it's now something that's a must-do, and I think we're going to see an increase in interest from industry and companies looking to work out how to solve the post-quantum problem."

Many of these first adopters will of course be at the root of the supply chain, such as semiconductor manufacturers who provide the boards and systems other companies rely on. Having PQC in these components is critical for the rest of the supply chain, and it's clear already from many of our customers in this sector that the quantum shift – from preparation to compliance – is already on.

That's one of the reasons why PQShield are excited to announce that we've designed and built our own fully functional PQC silicon test chip.

We think it's the first ever PQC-compliant silicon chip, and it's fully loaded with all the power of our PQPlatform IP, specifically focused on platform security for SoC semiconductor vendors with implementation security in mind. These are linked below:

- PQPlatform – Hash
- PQPlatform – Lattice
- PQPlatform – CoPro
- PQPlatform – SubSys

By building our own ASIC, we can now demonstrate and test the IP in the way that a customer will use it. We're able to evaluate its operation in real detail, looking at how to be compliant against the standards, and perform security testing in a way that would be practically difficult on a customer product deployed in the field.

In addition, the test chip gives us the ability to deep-dive into areas like power utilization, performance, and critically, the efficiency of Side Channel Analysis (SCA) countermeasures. As Graeme explains:

> *"It's much easier to do this type of evaluation on a real product, as opposed to a pre-silicon simu-*

lation, or using FGPA, and it's far more reflective of what a final product would look like."

Another benefit is the test chip's flexibility. The chip itself is a hardware/software co-design IP, and it allows us to update the firmware directly in order to modify the algorithms it supports. That means that it can be configured to be more standards-compliant in future, or perhaps even more secure from side-channel attack as analysis continues. The test chip has completely programmable firmware – a feature that will certainly be useful as NIST standardization continues, especially with the outstanding Round 4 KEMs, as well the on-ramp for new, wider digital signature schemes, such as our own side-channel secure, masking-friendly RACCOON scheme.

RACCOON has been specifically included in our silicon and this is likely to be the first time there has ever been an implementation in silicon that can actually be tested.

As the quantum shift continues, the focus on compliance with the NIST standards will sharpen, particularly as industry regulation follows. Our engineering and research teams have worked hard to ensure that PQShield has a tool that can allow us to test, evaluate and configure real-world post-quantum cryptography, keeping us all one step ahead of the quantum threat.

23.YubiKeys are vulnerable to cloning attacks thanks to newly discovered side channel

by Dan Goodin
https://arstechnica.com/security/2024/09/yubikeys-are-vulnerable-to-cloning-attacks-thanks-to-newly-discovered-side-channel/

The YubiKey 5, the most widely used hardware token for two-factor authentication based on the FIDO standard, contains a cryptographic flaw that makes the finger-size device vulnerable to cloning when an attacker gains temporary physical access to it, researchers said Tuesday.

The cryptographic flaw, known as a side channel, resides in a small microcontroller used in a large number of other authentication devices, including smartcards used in banking, electronic passports, and the accessing of secure areas. While the researchers have confirmed all YubiKey 5 series models can be cloned, they haven't tested other devices using the microcontroller, such as the SLE78 made by Infineon and successor microcontrollers known as the Infineon Optiga Trust M and the Infineon Optiga TPM. The researchers suspect that any device using any of these three microcontrollers and the Infineon cryptographic library contains the same vulnerability.

Patching not possible

YubiKey-maker Yubico issued an advisory in coordination with a detailed disclosure report from NinjaLab, the security firm that reverse-engineered the YubiKey 5 series and devised the cloning attack. All YubiKeys running firmware prior to version 5.7—which was released in May and replaces the Infineon cryptolibrary with a custom one—are vulnerable. Updating key firmware on the YubiKey isn't possible. That leaves all affected YubiKeys permanently vulnerable.

"An attacker could exploit this issue as part of a sophisticated and targeted attack to recover affected private keys," the advisory confirmed. "The attacker would need physical possession of the YubiKey, Security Key, or YubiHSM, knowledge of the accounts they want to target and specialized equipment to perform the necessary attack. Depending on the use case, the attacker may also require additional knowledge including username, PIN, account password, or authentication key."

Side channels are the result of clues left in physical manifestations such as electromagnetic emanations, data caches, or the time required to complete a task that leaks cryptographic secrets. In this case, the side channel is the amount of time taken during a mathematical calculation known as a modular inversion. The Infineon cryptolibrary failed to implement a common side-channel defense known as constant time as it performs modular inversion operations involving the Elliptic Curve Digital Signature Algorithm. Constant time ensures the time sensitive cryptographic operations execute is uniform rather than variable depending on the specific keys.

More precisely, *the side channel is located in the Infineon implementation of the Extended Euclidean Algorithm, a method for, among other things, computing the modular inverse*. By using an oscilloscope to measure the electromagnetic radiation while the token is authenticating itself, the researchers can detect tiny execution time differences that reveal a token's ephemeral ECDSA key, also known as a nonce. Further analysis allows the researchers to extract the secret ECDSA key that underpins the entire security of the token.

In Tuesday's (03 Sep 2024) report, NinjaLab co-founder Thomas Roche wrote:

> *In the present work, NinjaLab unveils a new side-channel vulnerability in the ECDSA implementation of Infineon 9 on any security microcontroller family of the manufacturer.This vulnerability lies in the ECDSA ephemeral key (or nonce) modular inversion, and, more precisely, in the Infineon implementation of the Extended Euclidean Algorithm (EEA for short). To our knowledge, this is the first time an implementation of the EEA is shown to be vulnerable to side-channel analysis (contrarily to the EEA binary version). The exploitation of this vulnerability is demonstrated through realistic experiments and we show that an adversary only needs to have access to the device for a few minutes. The offline phase took us about 24 hours; with more engineering work in the attack development, it would take less than one hour.*
>
> *After a long phase of understanding Infineon implementation through side-channel analysis on a Feitian 10 open JavaCard smartcard, the attack is tested on a YubiKey 5Ci, a FIDO hardware token from Yubico. All YubiKey 5 Series (before the firmware update 5.7 11 of May 6th, 2024) are affected by the attack. In fact all products relying on the ECDSA of Infineon cryptographic library running on an Infineon security microcontroller are affected by the attack. We estimate that the vulnerability exists for more than 14 years in Infineon top secure chips. These chips and the vulnerable part of the cryptographic library went through about 80 CC certification evaluations of level AVA VAN 4 (for TPMs) or AVA VAN 5 (for the others) from 2010 to 2024 (and a bit less than 30 certificate maintenances).*

In an online interview, Roche elaborated:

> *Infineon produces "security microcontrollers" or "secure elements." You can find many of them out there. Some of them (and this is the case for YubiKey 5 Series) run the Infineon cryptographic library (that Infineon develops for their customers that do not want to develop their own).*
>
> *This cryptolibrary is highly confidential (even its API is secret, you need to sign an NDA with Infineon just to know the API). Nobody, but Infineon, knows the cryptolibrary details and notably its countermeasures choices.*
>
> *This cryptolibrary, as many others, implement the ECDSA (core crypto function of FIDO, but also used in many different applications/protocols). Inside the ECDSA scheme, there are several sub-functions calls, one of them is the modular inversion of the ECDSA ephemeral key. This is a very sensitive oper-*

> *ation: any information leaking about the ECDSA ephemeral key would eventually reveal the ECDSA secret key.*
>
> *In the Infineon cryptolibrary the modular inversion is not constant time: different ephemeral key will lead to different inversion execution time. When acquiring the electromagnetic radiation of a chip running this function one can extract tiny differences of execution times throughout the inversion computation. These small timing leakages allow us to extract the ephemeral key and then the secret key.*

The attacks require about $11,000 worth of equipment and a sophisticated understanding of electrical and cryptographic engineering. The difficulty of the attack means it would likely be carried out only by nation-states or other entities with comparable resources and then only in highly targeted scenarios. The likelihood of such an attack being used widely in the wild is extremely low. Roche said that two-factor-authentication and one-time password functionalities aren't affected: because they don't use the vulnerable part of the library.

Tuesday's report from NinjaLab outlines the full flow of the cloning attack as:

1. The adversary steals the login and password of a victim's application account protected with FIDO (e.g., via a phishing attack).
2. The adversary gets physical access to the victim's device during a limited time frame without the victim noticing.
3. Thanks to the stolen victim's login and password (for a given application account), the adversary sends the authentication request to the device as many times as is necessary while performing side-channel measurements.
4. The adversary quietly gives back the FIDO device to the victim.
5. The adversary performs a side-channel attack over the measurements and succeeds in extracting the ECDSA private key linked to the victim's application account.
6. The adversary can sign in to the victim's application account without the FIDO device and without the victim noticing. In other words, the adversary created a clone of the FIDO device for the victim's application account. This clone will give access to the application account as long as the legitimate user does not revoke its authentication credentials.

The list, however, omits a key step, which is tearing down the YubiKey and exposing the logic board housed inside. This likely would be done by using a hot air gun and a scalpel to remove the plastic key casing and expose the part of the logic board that acts as a secure element storing the cryptographic secrets. From there, the attacker would connect the chip to hardware and software that take measurements as the key is being used to authenticate an existing account. Once the measurement-taking is finished, the attacker would seal the chip in a new casing and return it to the victim.

The attack and underlying vulnerability that makes it possible are almost entirely the same as that allowed NinjaLab to clone Google Titan keys in 2021. That attack required physical access to the token for about 10 hours.

The attacks violate a fundamental guarantee of FIDO-compliant keys, which is that the secret cryptographic material they store can't be read or copied by any other device. This assurance is crucial because FIDO keys are used in various security-critical environments, such as those in the military and corporate networks.

That said, FIDO-compliant authentication is among the most robust forms of authentication, one that's not susceptible to credential phishing or adversary-in-the-middle attacks. As long as the key stays out of the hands of a highly skilled and well-equipped attacker, it remains among the strongest forms of authentication. It's also worth noting that cloning the token is only one of two major steps required to gain unauthorized access to an account or device. An attacker also must obtain the user password used for the first factor of authentication. These requirements mean that physical keys remain among the most secure authentication methods.

To uncover the side channel, the researchers reverse-engineered the Infineon cryptographic library, a heavily fortified collection of code that the manufacturer takes great pains to keep confidential. The detailed description of the library is likely to be of intense interest to cryptography researchers analyzing how it works in other security devices.

People who want to know what firmware version their YubiKey runs can use the Yubico Authenticator app. The upper-left corner of the home screen displays the series and model of the key. In the example below, from Tuesday's advisory, the YubiKey is a YubiKey 5C NFC version 5.7.0.

YubiKeys provide optional user authentication protections, including the requirement for a user-supplied PIN code or a fingerprint or face scan. For the cloning attack to work against YubiKeys using these additional measures, an attacker would need to possess the user verification factor as well. More information about using these additional measures to lock down YubiKeys further is available here.

A key question that remains unanswered at the moment is what other security devices rely on the three vulnerable Infineon secure modules and use the Infineon cryptolibrary? Infineon has yet to issue an advisory and didn't respond to an email asking for one. At the moment, there is no known CVE for tracking the vulnerability.

24.Samsung unveils the Galaxy Quantum5 – a quantum-powered secure smartphone

by Samsung
https://www.gsmarena.com/samsung_unveils_the_galaxy_quantum5___a_quantumpowered_secure_smartphone-news-64363.php

First and foremost, let's start by saying that the Galaxy Quantum5 was developed in collaboration with SK Telecom and is exclusive to the network. So, you won't be able to get one outside of South Korea. The other partnership that made the phone possible is with ID Quantique (IDQ). That's the company that provides the quantum cryptographic chip.

The chip is a so-called Quantum Random Number Generator (QRNG). It uses quantum physics to generate truly random numbers to be used in the encryption and decryption process of sensitive data like biometrics and passwords. Using truly random numbers from an independent chip limits the possibility of outside influence and tampering, that more traditional random number generation methods are susceptible to.

Other than the additional QRNG chip, the Galaxy Quantum5 is basically a Galaxy A55. Some specs highlights include an aluminum frame and glass front and back body with IP67 ingress protection, a 6.6-inch, 120Hz Super AMOLED display, stereo speakers, a 50MP main camera, 12MP ultrawide and 5MP macro on the back. There is an Exynos 1480 running the show with 8GB of RAM and 128GB of expandable storage. A 5,000 mAh battery is keeping the lights on with 25W fast charging.