

<https://doi.org/10.1016/j.infsof.2023.107249>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Making existing software quantum safe: a case study on IBM Db2

Lei Zhang¹, Andriy Miransky¹, Walid Rjaibi², Greg Stager², Michael Gray³, and John Peck⁴

¹Department of Computer Science, Toronto Metropolitan University, Toronto, Canada

²IBM Canada Lab, Markham, Canada

³IBM, Bundall, Australia

⁴IBM, Austin, USA

{leizhang, avm}@torontomu.ca, {wrjaibi, gstager}@ca.ibm.com, mickgray@au1.ibm.com, johnpeck@us.ibm.com

Abstract

The software engineering community is facing challenges from quantum computers (QCs). In the era of quantum computing, Shor’s algorithm running on QCs can break asymmetric encryption algorithms that classical computers practically cannot. Though the exact date when QCs will become “dangerous” for practical problems is unknown, the consensus is that this future is near. Thus, the software engineering community needs to start making software ready for quantum attacks and ensure quantum safety proactively.

We argue that the problem of evolving existing software to quantum-safe software is very similar to the Y2K bug. Thus, we leverage some best practices from the Y2K bug and propose our roadmap, called 7E, which gives developers a structured way to prepare for quantum attacks. It is intended to help developers start planning for the creation of new software and the evolution of cryptography in existing software.

In this paper, we use a case study to validate the viability of 7E. Our software under study is the IBM Db2 database system. We upgrade the current cryptographic schemes to post-quantum cryptographic ones (using Kyber and Dilithium schemes) and report our findings and lessons learned.

We show that the 7E roadmap effectively plans the evolution of existing software security features towards quantum safety, but it does require minor revisions. We incorporate our experience with IBM Db2 into the revised 7E roadmap.

The U.S. Department of Commerce’s National Institute of Standards and Technology is finalizing the post-quantum cryptographic standard. The software engineering community needs to start getting prepared for the quantum advantage era. We hope that our experiential study with IBM Db2 and the 7E roadmap will help the community prepare existing software for quantum attacks in a structured manner.

1 Introduction

The field of quantum computing is still young, but it has been evolving rapidly during the last decade. With the vast increase in computing power, quantum computers (QCs) promise to revolutionize many fields, including artificial intelligence, medicine, and space exploration [1]. However, they can also be abused to break most common encryption algorithms, which modern cryptography depends upon today to ensure digital information’s safety and privacy. Thus, we position that the software engineering community should start thinking about the impact of quantum computing on cybersecurity and the best practices to address these concerns. First, let us look at the evolution of QCs.

1.1 The timeline of QCs

The field of quantum computing is young. In 1982, Richard Feynman introduced the idea of simulating quantum physics with computers [2]. It took a while to implement an actual QC. A partnership between academia and IBM created the first working QC with two qubits¹ in 1998 [3]. Then, it took the company 18 years to develop a 5-qubit QC, which became accessible to the public in 2016 [4].

At present, a few QCs are commercially available. D-Wave started selling adiabatic QC² in 2011, with the current offerings having > 5000 qubits [7]. Recently, D-Wave announced that it would follow the lead of IBM, Google, and others to develop an all-purpose quantum computer using a “gate-model”, which has a similar design to those from its rivals [8].

QCs are also available via fully-managed Cloud services. IBM gave access to 20- and 50-qubit gate-based superconducting QCs to academic and industrial partners to explore practical applications in 2017 [9], and a 65-qubit machine was offered in 2020 [10]. For non-commercial use, IBM offers 5- and 15-qubit QCs via IBM Q Experience online platform based on IBM Cloud (along with local- and Cloud-based simulators) [11]. Rigetti offered 8-qubit superconducting QC in 2017 [12]. Google built a 72-qubit gate-based superconducting QC in 2018 [13]. IonQ introduced ion-trapped 11-qubit QC in 2019 [14]. Honeywell created ion-trapped 10-qubit QC in 2020 [15]. Xanadu offered 8- and 12-qubit photonic QCs in 2020 [16, 17]. Microsoft provides access to a simulator³ of a topological QC via Microsoft Quantum Development Kit [18] (and is planning to give access to an actual QC in the future).

Despite the rapid development of QCs, we conjecture that QCs will not replace CCs in the short run. Instead, QCs will be integrated into a System of Systems via private or public Cloud service [20, 21, 22]. Aggregated Cloud services are starting to appear as well. For example, Amazon Web Services started offering access to QCs from various vendors via its Braket service in 2019 [23]. Currently, it offers D-Wave adiabatic 2048- and 5640-qubit QCs, IonQ trapped-ion-based 11-qubit QC, and Rigetti 32-qubit superconducting QC [24].

To summarize, many formidable competitors are scaling up various QC architectures, bringing us closer to the day when QCs can solve practical problems. So let us look at when it may happen.

1.2 Quantum advantage

In the future, a large QC can solve some problems that a CC practically cannot, which is called quantum advantage [2] (often used interchangeably with the term of quantum supremacy). Quantum advantage was demonstrated on a superconducting QC in 2019 [25]. Another demonstration of quantum advantage was done on a photonic QC in 2020 [26], although the setup used in the experiment may be difficult to scale up or generalize [27].

But when will QCs start solving real-world problems? Using existing QCs, quantum chemists are already able to improve simulations of small chemical systems [28] and some large ones, albeit with approximations [29]. However, quantitative financiers will need a machine with $\approx 7.5K$ logical

¹Qubit is the basic unit of quantum information, on which QC operate. At any point of execution, the state of a classical computer (CC) is given by a vector of bits taking the values of 0 and 1. The state of a QC is, however, given by a vector of qubits and bits.

²The debate about adiabatic QC being a “true” QC is ongoing [5]. A hybrid of adiabatic and gate-based QC is promising [6], but no commercial implementation is available.

³A QC can be simulated on a CC [11, 18]. A quantum simulator interprets a mathematical function as part of a physical model [19]; however, it will not yield performance improvement that a QC would provide, as the underlying host system of the simulator is still based on bits rather than qubits (a basic unit of quantum information). Thus, one needs a real QC to reap performance benefits.




Figure 1: Number of qubits (in millions) needed to break RSA 2048 based on the literature review in [31, Table II]. The x-axis denotes a year when an article was published.

qubits to price financial instruments [30]. Hackers will need a computer with 20M qubits to break the 2048-bit RSA key in less than a day [31].

To start addressing practical use-cases within the next decade, IBM stated that they “need to at least double the Quantum Volume (QV) of our quantum computing systems every year.” [32] So far, IBM is on track, demonstrating the QV of 64 on a QC with 27 qubits in 2020 [10]. By 2023, IBM plans to ship a QC of 1,121 qubits (with the expectation of proportional QV growth) [10].

1.3 Quantum advantage: impact on cybersecurity

The power of quantum computing will threaten modern cybersecurity platforms by speeding up (1) factorization of integers, solving the discrete logarithm problem, and the elliptic-curve discrete logarithm problem (using Shor’s algorithm [33]); as well as (2) the search in a set (with the help of Grover’s algorithm [34]). Both tasks are foundational for modern encryption algorithms.

While the existing QCs are not ready to use Shor’s algorithm for production problems, many brilliant people (both on the hardware and the algorithm sides) are racing to get us there. For example, Gidney et al. [31] combine a set of clever tricks and techniques to implement Shor’s algorithm using modern QC architecture. They show that one needs 20M qubits to break the 2048-bit RSA key in less than a day. They compare their solution against an existing one by Fowler et al. [35], which was introduced in 2012. This baseline approach required one billion qubits. Thus, in seven years, algorithms designers reached a remarkable 165x [31, Table II] improvement in the algorithm implementation. This dynamic is represented graphically in Fig. 1. The 20M qubits is not an optimal boundary, which implies that in the future, improvement schemes may be derived, further reducing the hardware requirements.

To improve QCs, electrical engineers and physicists will continue working on enhancing coherence, noise reduction, and error correction solutions. Theoretically [36], a QC with 4099 ideal, i.e., perfectly stable, qubits can break RSA 2048 in 10 seconds. Such stable qubits will appear in fault-tolerant

quantum computers (FTQCs), which will employ error-correcting schemes that will allow us to operate on logical/ideal qubits (constructed from groups of physical qubits) [37]. Small FTQCs may appear within ten years [38]. Moreover, new contenders in this field (such as photonic QC developed by Xanadu [39] or ionic QCs developed by IonQ [14] and Honeywell [15]) may introduce new and efficient computing architectures that will speedup arrival of FTQCs. IonQ, for example, already produced a QC with 20 ideal qubits [40]. This paper will use the term FTQC to refer to a QC that is powerful enough to break existing encryption schemes.

1.4 Quantum advantage: call to action

Since FTQCs are years away, can we still rely on existing encryption schemes without upgrading them? Unfortunately, the short answer is “No.” To arrive at this conclusion, let us examine the relationships between three time intervals:

- (1) Time (denoted by X) needed to deploy quantum-safe cryptographic solutions;
- (2) Time (denoted by Y) required to maintain the security of your encrypted data; and
- (3) Time (denoted by Z) when a quantum computer capable of breaking existing encryption schemes, such as FTQC, will appear.

Suppose $X > Z$. In this scenario, the deployment of quantum-safe cryptography (i.e., X) will take longer than the emergence of FTQCs (i.e., Z). QCs capable of breaking public-key encryption schemes may be developed within the next twenty years [41]. While this feels like a long time, historically, it takes almost two decades to deploy public-key cryptography infrastructure [41]. Thus, even if we begin planning to switch from existing encryption algorithms to quantum-safe ones today, upgrading the security features will take a lot of time and effort.

What if $X \leq Z$? It implies that we have updated our encryption schemes proactively, which is an essential topic of this paper.

Now, suppose $Y > Z$. In this scenario, we should maintain the security of the encrypted data (i.e., Y) past the arrival of FTQCs (i.e., Z). Malicious entities can, for example, harvest sensitive data that must remain confidential for many decades (e.g., state secrets) and then use that data to perform a harvest-then-decrypt attack. Therefore, an adversary can harvest sensitive data communications today and — when FTQCs are available — use that computing power to break today’s non-quantum-safe encryption to access such sensitive data.

In the case of $Y \leq Z$, the data becomes stale and non-sensitive before the FTQCs arrive. For the time being, most data types fall into this category. However, new data are being generated every day, while $Z \rightarrow 0$. Thus, more data types will be at risk as time passes.

In summary, due to the uncertainties inherent in quantum threats and the ambiguous timeline, the sooner we start working on upgrading to quantum-safe encryption schemes, the sooner we will be able to mitigate this risk.

Now is the time to start planning the evolution of encryption schemes (e.g., identification of vulnerable hardware and software components and design of replacement schemes). It may take many months (if not years) to identify all problematic hardware and software components in an enterprise environment and devise corrective actions. An organization that begins working on these tasks now may be able to complete them by the time the security standards are finalized.

1.5 Quantum-safe: existing solutions

Existing approaches proposed to ensure information security in the era of quantum computing have two trends. The first approach is quantum cryptography, which harnesses quantum mechanics to

protect confidential internet communications. For example, quantum key distribution (QKD) is the most popular approach, where data is transferred via no-change and no-cloning photons instead of bits [42]. However, the downside of quantum cryptography is that it needs quantum channels for sending quantum bits among different locations [43]. The second approach is post-quantum cryptography (PQC) — the focus of this paper — which is about preparing for the era of quantum computing by updating existing mathematical algorithms and encryption standards with classical cryptographic algorithms that are secure against both QCs and CCs [44, 45, 46].

The standard of PQC is still in development. Based on the timeline of the National Institute of Standards and Technology (NIST), we expect to see a draft of the standard between 2022 and 2024 [47]. Nevertheless, it is still valuable to plan to upgrade existing cryptographic systems, particularly doing so while adhering to crypto-agility practices (as recommended by NIST [48]), so that any algorithm chosen today can be changed in the future (if necessary) without incurring a considerable cost. While we need to wait for one to three years to see the standard, we already have a good understanding of the classes of algorithms that will appear in this standard. In July 2020, NIST selected 15 candidate algorithms (7 finalists and 8 alternates) in Round 3 of the competition [47].

This paper extends our previous publication [1], which introduced the 7E roadmap for software developers to address encryption-related challenges associated with quantum advantage. The **contributions** of this paper are

- (1) Validating the 7E roadmap by conducting a case study on IBM software,
- (2) Describing our PQC implementation for the IBM Db2 software⁴,
- (3) Summarizing challenges and lessons learned based on our experience, and
- (4) Logging our best practices and improving the 7E roadmap.

These contributions may be helpful to practitioners planning PQC upgrades and academics working in the area of software maintenance and evolution.

The rest of this paper is structured as follows. Section 2 analyzes the impact of quantum computing on existing software systems. Section 3 recaps the original 7E roadmap. We propose the enhanced 7E roadmap in Section 4. The enhancements are based on the case study, which is described in the following sections. Section 5 presents the industrial settings of our case study. Section 6 illustrates the implementation of PQC and the performance evaluation. Section 7 discusses the challenges we met and our takeaway messages for other researchers and practitioners. We discuss threats to validity and our concerns in Section 8. Finally, Section 9 concludes the paper.

2 The impact of quantum computing on existing systems

It is essential to understand that quantum computing will affect encryption schemes differently depending on the class of encryption algorithm. In this section, we first elaborate on the asymmetric (public-key) encryption algorithms, which are used in many areas ranging from the Transport Layer Security (TLS) protocol (used to safeguard data passed between two systems on the Internet) to Pretty Good Privacy software used to encrypt and decrypt a file and safely transfer it between computers.

We then proceed to the symmetric (private-key) encryption algorithms, such as Advanced Encryption Standard (AES), used to protect sensitive data. There exist numerous use-cases, ranging

⁴Our PQC schemes are generic and applicable to other software products.

Table 1: Effective security strength of key encryption algorithms as per [49]

Encryption algorithm	Key size (bits)	Effective security level on CCs (bits)	Effective security level on QCs (bits)
RSA 1024	1024	80	0
RSA 2048	2048	112	0
ECC 256	256	128	0
ECC 384	384	256	0
AES 128	128	128	64
AES 256	256	256	128

from encrypting a file archive (e.g., implemented in 7z and WinZip software) to encrypting computers’ disks (e.g., using Apple MacOS FileVault and Symantec Endpoint Encryption).

2.1 Asymmetric encryption

Asymmetric encryption algorithms, which are based on factoring large integers (e.g., Rivest-Shamir-Adleman — RSA), discrete logarithms (e.g., Elliptic Curve Cryptography — ECC, and Diffie-Hellman key exchange — DH), or similar approaches (see [49] for review) will need to be replaced by quantum-safe alternatives [49]. Effective security strength, shown in Table 1, suggests that the strength of the RSA and ECC is somewhat weaker or comparable to AES on a CC, but is extremely weak on a QC. This is because Shor’s algorithm can perform integer factorization in polynomial time; so what requires thousands of years with classical computers would only take days/hours on a large-scale quantum computer. This, of course, assumes that a large-scale quantum computer with the required number of qubits and degree of coherence exists, which is not the case right now.

2.2 Symmetric encryption

Unlike asymmetric encryption algorithms, symmetric encryption algorithms do not face an existential threat: one needs to perform a brute-force attack to break it. However, on a CC generation of n keys require $O(n)$ operations, while on a QC it can be done using $O(\sqrt{n})$ operations, thanks to Grover’s algorithm. Thus, a large quantum computer running Grover’s algorithm could provide a quadratic improvement in brute-force attacks on symmetric encryption algorithms, such as AES. This translates into a need to double key size to support the same level of protection. For AES specifically, this means using 256-bit keys to maintain today’s 128-bit security strength,⁵ as depicted in Table 1.

2.3 Challenges on existing systems

If the legacy system is well-designed and actively maintained, then the solution is straightforward: one can replace an existing asymmetric algorithm with a new one (or increase a key size of a symmetric one) while ensuring that the existing data can be migrated to a new format. However, it may require downtime to re-encrypt existing data with a quantum-safe algorithm.

Often, altering existing (legacy) systems to address the security concern may be challenging. Legacy systems frequently lack adequate information or support to be maintained or upgraded. The root causes of these issues are numerous. For example, the system developers may be unavailable (e.g., because they left the company or retired), source code or documentation may be lost, or build platforms for the source code may be sunset. To make matters worse, the encryption-related code may be spread or cloned among multiple software components (due to bad design), making

⁵That is, an n -bit AES cipher provides a security level of $n/2$ because $\sqrt{2^n} = 2^{n/2}$.

alterations even more challenging. These root causes make it extremely difficult and expensive to upgrade such a system to the newest security protocols, sometimes making the replacement the only feasible option.

2.4 Threats to the existing data

As discussed in Section 1.4, sensitive encrypted data (e.g., protected by encrypted network packets of TLS or an encrypted disk) are vulnerable to the harvest-then-decrypt attack, meaning one can leverage QC power to break the asymmetric encryption used by TLS or brute-force access to the encrypted disk and recover the sensitive data. While not much can be done about the protocols involving asymmetric algorithms; for the symmetric ones, we can increase the length of the key right away to counter the brute-force attack [50]. We can also encrypt archived data (e.g., stored on backup devices) with a quantum-safe algorithm.

Another example is a blockchain platform using proof-of-work algorithms (e.g., based on SHA-256). The security of blockchain platforms relies on a digital signature, which is based on either Elliptic Curve Digital Signature Algorithm [51] or RSA algorithms; both are vulnerable to QCs. Kiktenko et al. [52] proposed a quantum-secured blockchain framework that utilizes QKD techniques via an experimental fiber network (the cost of the network is not disclosed). A more efficient approach may be to switch to a proof-of-stake algorithm, which does not require solving computationally intensive mathematical puzzles, from a proof-of-work algorithm.

3 Our method

We proposed a **7E** roadmap for software developers in [1], summarizing steps to address encryption-related challenges associated with quantum advantage and giving developers a structured way to prepare for the quantum advantage era. The 7E roadmap is inspired by the lessons learned from the Y2K era and is amalgamated with modern knowledge. We recap key features of the seven steps of 7E for completeness below. More detailed steps can be found in Section 4.

- (1) **Engage** executives and senior management so that they can sponsor the initiative.
- (2) **Examine** existing products and their cybersecurity components to identify and locate the issues, review the document and programs and assess the problems.
- (3) **Evolve**: design a new software with crypto-agility (as per NIST recommendations [48]) in mind so that quantum-safe algorithms can be added to the software later on.
- (4) **Educate** the programmers and designers to ensure that everyone is “on the same page” because (in most cases) the security-related component is coupled with the remaining software components.
- (5) **Estimate** the impact of potential problems and the cost of alternatives to prioritize the problems.
- (6) **Execute** the new cybersecurity policy.
- (7) **Essay** the new cybersecurity policy. Keep monitoring the performance and the robustness of your new cybersecurity policy in production to make sure that the challenges associated with quantum advantage were addressed; adjust the policy if needed.

We keep the generality of the 7E roadmap so that it serves as a general guideline and can be generalized in other applications. However, while the 7E roadmap shares some common characteristics with other software evolution guidelines (especially, ones in regard to Y2K), we highlight two distinguishing characteristics of 7E roadmap for quantum-safe cryptography:

- (1) **Crypto-agility**, which is important for PQC evolution because NIST is still standardizing post-quantum cryptography [47], and a FTQC has not yet been developed. Thus, the current implementations of PQC may be changed in the future. Crypto-agility provides the ability to evolve the cybersecurity components without significant impact on the rest components of your system;
- (2) **Education**, as the challenges of quantum attacks are new to managements and developers. Thus, knowledge transfer is important to make sure everyone is on the same page.

Additional details about how 7E differs from generic guidelines for maintaining and evolving a software product are provided in Section 7.

4 Enhanced 7E Roadmap

To validate the 7E roadmap, we conduct a case study on IBM Db2 database system (hereon abbreviated to Db2). Based on our case study (details of which will be given in Section 5-onward), we incorporate our experience⁶ and findings into the original 7E roadmap and propose an enhanced 7E roadmap. The key enhancements are highlighted in **bold** text. In addition, we find out the importance of education. Thus, we bring Step 4 of “Educate” to Step 2 in the enhanced version.

- (1) **Engage executives and senior management** so that they can sponsor the initiative. It is important to get acknowledgement from the decision makers in your company or organization. Moreover, executives and senior management can assess security concerns from a broader perspective. They will also sponsor the allocation of the human resources (e.g., the experts within the organization who will drive the project forward). Quantum computing is an emerging technology, and it may take time for executives to learn its growing powers and realize the side effect of these powers on the existing cryptography schemes. To educate the management, you can use formal presentations or reports and incorporate their feedback later.

In addition to management, we also need to engage multiple experts from different domains⁷, because the evolution of the cybersecurity component often involves other parts of the system. For our project to be successful, multiple experts had to collaborate.

- (2) **Educate the programmers and designers** to ensure that everyone is “on the same page” because (in most cases) the security-related component is coupled with the remaining software components. This implies that the whole development organization needs to be aware of the challenges of quantum attacks. The application of PQC is still in its infancy. Thus, security professionals may not have sufficient knowledge of PQC to perform the evolution of cryptographic schemes confidently. This step is complementary to Step 1, focusing on training technical staff rather than the management.

⁶Please note that not all of our experience will necessarily apply to other cases. However, by sharing as much information as possible, we allow practitioners and researchers to decide which approaches are best for them.

⁷In our case, we have a CTO of Data Security (IBM), Senior Software Developer (IBM Db2), Architects (IBM Security), and Software Engineering Researchers (Ryerson University).

We transfer knowledge and brainstorm within and outside this research team through work sessions, conferences, seminars, publications, etc. Researchers and practitioners can use these activities to estimate the impact of PQC evolution on IBM Db2 software. In addition, discussions during these events generate new ideas and questions⁸.

- (3) **Examine** existing products and their cybersecurity components to identify and locate the issues, review the document and programs and assess the problems. For legacy systems, there exist difficult scenarios, such as lack of documentation, source code, or build infrastructure (as discussed in Section 2.3). Identify existing data (if any) that may require protection.

Note that the evolution of PQC may also affect other components of the existing system. As an example, our work mainly altered the user authentication component of IBM Db2. Nevertheless, there are some cases where we need to make changes outside of the user authentication component. For example, some checks of encryption key sizes have been found in other parts of Db2.

- (4) **Evolve:** design a new software with crypto-agility (as per NIST recommendations [48]) in mind so that quantum-safe algorithms can be added to the software later on. Conceptually, this should be achieved by designing the systems in such a way that an existing encryption scheme can be easily replaced with a new one. Moreover, the systems should be able to recognize and “translate” multiple encryption schemes.

For example, the encryption component can be designed to be plug-and-play so that an existing encryption algorithm can be replaced (if it is discovered to be vulnerable) with a robust one. This will save costs in the future when the standards of PQC are finalized, and our software has to be updated with these algorithms. **We achieve crypto-agility to a certain degree by replacing fixed key size with dynamic key size in our case study. At the same time, we acknowledge the challenges to redesign the whole encryption component of existing software to implement crypto-agility.**

The challenges (identified in the previous steps and Section 2.3) may affect the easiness of design decisions. For example, if documentation is missing, one may first have to reverse engineer the architecture from the code by extracting caller-callee function pairs and identifying security-related ones.

For cryptography that involves multiple business partners, achieving crypto-agility requires that all business partners update hardware and software promptly. Moreover, all the partners should disclose crypto-related information, e.g., security certificates and protocols, to each other. Not only that, multi-organizational interactions can be enabled by institutionalizing open standards, preferably with open source reference implementations of those standards. Conceptually, this process would be akin to the creation of other open standards, where an open standards organization would be formed to drive the process.

- (5) **Estimate** the impact of potential problems and the cost of alternatives to prioritize the problems. The findings from Steps 2 and 3 should help to estimate the cost. Rate the cost of potential solutions in terms of human and time resources. Work first on the systems that handle critical data first. The definition of critical will vary from industry to industry. A representative example is a system handling personal data, such as financial transactions or health records. **It is always a good idea to prepare some buffer time in your**

⁸For example, choosing the right version of Kyber, so that performance and security are balanced.

project management. Though our project involved multiple experts and had a clear roadmap, we still underestimated the time needed to complete it by about 50% due to the challenges discussed in Section 7.

Upgrade of the cryptographic schemes involving symmetric encryption will typically be cheaper than the asymmetric one. The former will usually require increasing the length of the key and increasing the storage space for the keys. The latter may also require replacing the code needed to do the encryption.

- (6) **Execute the new cybersecurity policy.** Select and adopt appropriate solutions based on requirements, budgets, and priorities. As discussed in Section 2.3, practitioners can execute the new cybersecurity policy in different ways. For newly-built systems, PQC may be adopted (see Sections 2.1 and 2.2). For legacy systems, the software and associated hardware may have to be altered (see Section 2.3). For existing data, an intermediate solution — e.g., re-encrypting the existing data with a quantum-safe cryptographic algorithm — may be applied (see Section 2.4). **We choose the latter because we work on the existing Db2 software. The challenges and take-away messages from our execution can be found in Section 7.**
- (7) **Essay the new cybersecurity policy.** Keep monitoring the performance and the robustness of your new cybersecurity policy in production to make sure that the challenges associated with quantum advantage were addressed; adjust the policy if needed. Experiences and lessons learned from one project may also apply to another one. These lessons could serve as a building block to a general theory of making PQC evolution agile and smooth. **We transfer our knowledge of PQC and the experience with Db2 both inside and outside IBM via inter-project discussions, conferences, seminars, publications, etc. We hope our roadmap and experience can help more practitioners get prepared for the year of quantum advantage (also known as Q2K).**

5 Industrial Settings

We choose Db2 because it complies with various security standards to protect sensitive data. In addition, it is a mature product (initially released in 1987, 34 years ago). However, it is under active development, with new features being added regularly. Db2 codebase consists of tens of millions of lines of C/C++ code (note that the C/C++ codebase was started in the early 1990s). Let us briefly discuss Db2 features that require cryptographic algorithms (that are vulnerable to quantum attacks) and identify how to replace them with PQC algorithms.

Db2 consists of a Db2 server and data server clients. A Db2 server is an enterprise-grade relational database management system that delivers data to clients. A Db2 client is an application that runs Db2 and SQL commands against the server.

The first step of accessing a Db2 database is authentication. Authentication is the process by which a system verifies a user’s identity. According to IBM Db2 documentation version 11.5 [53], Db2 requires two items to authenticate a user, i.e., a user ID and a password. The user’s identity is verified through the correct user ID and its corresponding password.

To keep in-transit data safe, Db2 builds an encrypted communication between the user and Db2 using TLS protocol. TLS uses both asymmetric cryptography (e.g., DH) and symmetric cryptography (e.g., AES) for encryption.⁹ DH is a method to exchange cryptographic keys over a

⁹Technical note: the implementations of DH and TLS are independent. DH is adopted at the Distributed Relational

public channel securely. Db2 adopts DH to exchange a shared secret key between the user and Db2. The key is then used to encrypt subsequent communications with symmetric cryptography. Figure 2 illustrates a simplified process of DH key agreement between two parties for authentication.

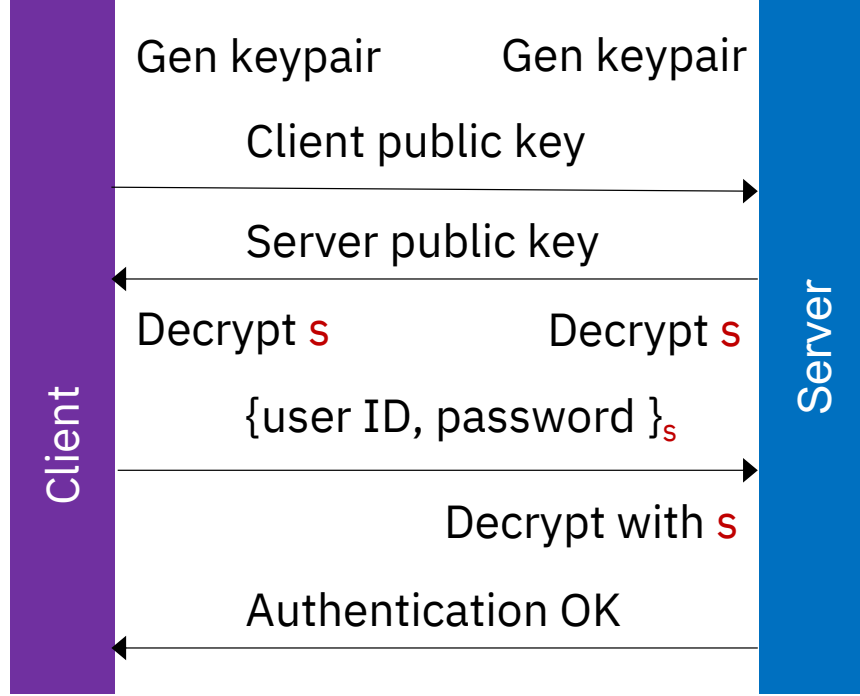


Figure 2: DH key exchange method establishes a shared secret between two parties for authentication. Both the client and the server generate a pair of keys, i.e., private key and public key. They then exchange their public keys. The client and the server independently compute the shared secret (s) using the other's public keys and their own private keys. The client encrypts user ID and password using s and sends them to the server. After decrypting the user ID and password using s , the server finishes the authentication of the client.

Key exchange algorithm. DH key exchange relies on the assumption that discrete logarithm problem is hard to solve on classical computers (as discussed in Section 2.1). However, the discrete logarithm problem can be efficiently solved using Shor's algorithm on quantum computers. In other words, DH key exchange is vulnerable to quantum attacks. To tackle this problem, we need to find a quantum-safe replacement.

We choose Kyber [56], which is one of the algorithms in the Cryptographic Suite for Algebraic Lattices (CRYSTALS), as the successor of DH. Kyber is a key encapsulation algorithm based on the hardness of solving the learning-with-errors problem over module lattices [57]. Kyber is one of the four public-key encryption and key-establishment algorithms in the round three finalists announced by NIST [58]. The procedure of Kyber key exchange in user authentication can be seen in Fig. 3.

Digital signature and TLS. Since DH key exchange does not provide authentication, it is vulnerable to man-in-the-middle attacks. For this reason, it is implemented alongside RSA digital signature for authentication in TLS. However, as discussed in Section 1, RSA is also vulnerable to quantum attacks. Thus, we need to replace RSA with a quantum-safe digital signature scheme.

Database Architecture (DRDA) layer [54], i.e., the application layer (layer seven) in Open Systems Interconnection (OSI) model [55]. TLS is implemented at the transport layer (layer four) in OSI and is independent of the DRDA layer.




Figure 3: Kyber key encapsulation mechanism between two parties for authentication. The client generates a pair of private key and public key, and the server generates a random number as the shared secret (s). The client sends its public key to the server. The server uses the client's public key to encrypt s into a ciphertext and sends the ciphertext back to the client. The client derives s from the ciphertext using its own private key. Then, the client encrypts its user ID and password using s and sends them to the server. After decrypting the user ID and password using s , the server finishes the authentication of the client.

We choose Dilithium [59], also part of CRYSTALS, as our new digital signature scheme. Similar to Kyber, Dilithium is also based on the hardness of lattice problems over module lattices. Dilithium is one of the three finalists of digital signature algorithms submitted to NIST in round three [58].

TLS 1.3 offers significant improvements over its predecessor TLS 1.2; the key feature that interests us is its support of post-quantum authentication within IBM Global Security Kit (GSKit) [60]. For this reason, we upgrade¹⁰ TLS 1.2 in Db2 to TLS 1.3.

6 Experiments

In this section, we discuss the implementation of Kyber, Dilithium, and TLS 1.3 in Db2.

6.1 GSKit

IBM GSKit provides libraries and utilities for both general purpose cryptography and Secure Sockets Layer (SSL) or TLS communication. More specifically, GSKit consists of two separate packages, containing different functionalities [61] (1) GSKit Crypt, which contains the cryptographic

¹⁰To do the upgrade, we integrated Dilithium into TLS 1.3.

```

//Keypair generation API (client)
int crypto_kem_keypair (unsigned char *pk,
                        unsigned char *sk,
                        KYBK k);

//Key encryption API (server)
int crypto_kem_enc (unsigned char *ct,
                   unsigned char *ss,
                   const unsigned char *pk,
                   KYBK k);

//Key decryption API (client)
int crypto_kem_dec (unsigned char *ss,
                   const unsigned char *ct,
                   const unsigned char *sk,
                   KYBK k);

```

Listing 1: Kyber APIs in GSKit.

algorithms that the package of GSKit SSL depends on; (2) GSKit SSL, which includes the essential runtime support to enable security calls and the use of the TLS protocol.

In our experiments, we focus on two things (1) use of the general-purpose cryptographic library to evolve the algorithm used in Db2 user authentication towards quantum safety (i.e., we upgrade from DH to Kyber), and (2) use of the TLS implementation to evolve the Db2 TLS capability towards quantum safety (i.e., we migrate from TLS 1.2 to TLS 1.3).

6.2 Implementation of Kyber

As can be seen in Fig. 3, Kyber key encapsulation scheme shares a secret key between the server and the client in three key steps as follows.

- (1) Client generates a private/public Kyber key pair and sends the public key to the server.
- (2) Server takes a public key and uses a random seed to generate a shared key and ciphertext.
- (3) Client uses private key and ciphertext to get the shared key.

After these three key steps, we can use Kyber shared secret as a symmetric key (e.g., an AES 256 bit key) to protect channel between client and server.

GSKit implements Kyber and provides three APIs corresponding to the three key steps above. They can be found in Listing 1.

The parameter **pk** is public key, **sk** is secret key, **ct** is ciphertext, **ss** is shared secret, and **k** is Kyber key, which indicates the security level of Kyber [57].

6.3 Dilithium and TLS 1.3

TLS, which is the successor of SSL, provides communications security over a computer network. Statistics show that over 90% of the web traffic is based on SSL/TLS-based HTTPS protocol [62]. Db2 uses TLS to secure all communications between servers and clients. TLS protocol requires digital signatures (in most cases, based on asymmetric cryptography).

In TLS, symmetric cryptography, i.e., Data Encryption Standard (DES) or AES, is used to encrypt the data transmitted. The keys for this symmetric encryption are uniquely generated for

Table 2: Summary stats of response time.

Algorithm	Avg. response time (ms)	St. Dev. (ms)
Kyber	162.514	15.281
DH	163.552	12.222

each connection and are based on a shared secret, which was negotiated using a public-key-agreement protocol — DH. The identity of the communicating parties are authenticated using asymmetric cryptography, i.e., RSA.

As discussed in Section 5, we need to replace RSA with quantum-safe cryptography, i.e., Dilithium. To enable a post-quantum digital signature scheme, we need to upgrade TLS to the latest version — 1.3 — from the existing version — 1.2. GSKit has implemented PQC algorithms as an extension of TLS 1.3 protocol.

The upgrade of TLS consists of two steps. In the first step, we need to change Db2 codebase so that Db2 will use TLS 1.3 provided by GSKit. This change is associated with two settings: (1) enable TLS 1.3 with PQC (on both client and server sides) for the PQC handshake to take place, (2) enable PQC certificates in TLS 1.3 for a digital signature on the server side. In the second step, we need to update the certificate chain using the GSKit command tool `gsk8capicmd` for key and certificate management. We create a traditional certificate chain and an attached post-quantum certificate chain to create the post-quantum certificates for TLS 1.3.

To verify the implementation of Dilithium and TLS 1.3, we print out the HTTPS connection reports and the TLS handshake transcript. The HTTPS connection reports verify two things: (1) TLS 1.3 has been enabled, and (2) quantum-safe certificates are used. To confirm the result, the TLS handshake transcript is used to verify that PQC, i.e., Dilithium, has been enabled (instead of RSA).

6.4 Performance evaluation

We evaluate the performance difference between Kyber and DH on our test bed, which has 32 CPUs (1.2 GHz) and 256 GB of memory. The performance is evaluated by running the command of `db2 connect` from a client to connect to a database on the server, which triggers TLS communication and user authentication. To evaluate the performance of Kyber, we adopt the Monte Carlo method to repeat this test 1,000 times with both Kyber and DH, respectively. The distributions of the timing data are shown in Fig. 4 and summary stats — in Table 2. The average response time with Kyber is 162.514 ms, and the standard deviation is 15.281 ms; the average response time with DH is 163.552 ms, and the standard deviation is 12.222 ms. Based on the QQ-plots and Shapiro-Wilk normality test, the distributions of timings are non-normal. The difference between the distributions are statistically significant: the one- and two-sided Wilcoxon rank-sum tests yield p -value of ≈ 0.0002 and ≈ 0.0004 , respectively. From a practical perspective, we can conclude that the performance of Db2 is not affected by the PQC upgrade — the average response time even decreases by 0.635% after we migrate from DH to Kyber (which is vital to performance-critical systems).

7 Challenges and take-away messages

We list the challenges that we met during the development and the lessons learned from this project below.




Figure 4: Distribution of response time.

- (1) Lack of documentation, ambiguity of documentation, and obsolete documentation (corresponding to Step 2 in 7E — Exam). Because of the passing of time, mature systems often lack development documentation, or the documentation is not up-to-date, which slows down the pace of the learning curve for the cryptography system. Even if the document exists, it can contain errors.

For example, we uncovered errors in the API document, where the socket handle parameter should have been an environment handle.

Another example relates to outdated documents. In one case, the instructions in the document we received were written for 32-bit operating systems, but we were running on the 64-bit system. As a result, the commands provided in the document were inoperable, and we had to explore and create new instructions tailored for the modern operating system.

- (2) Legacy/fragile development environment (corresponding to Step 6 in 7E — Execute). Because of the long development history of Db2 (34 years) and its mission-critical feature, Db2 project management requirements often lead to a limited number of unified legacy solutions. Many of them are command-line tools. Compared to graphical user interface, complex command-line interface has a steeper learning curve.
- (3) Large codebase (corresponding to Step 2 in 7E — Exam). A large codebase (in our case, tens of millions of lines of C/C++ code) with a long history of development often has three drawbacks during maintenance or evolution, i.e., (a) lack of comments, (b) complex structures, and (c) long compilation time. Let us elaborate on these three items now. First, the lack of well-written comments (detailing the whole design) can make code evolution much more challenging. Second, as the size of the codebase increases, the project becomes more and more

complex. Consequently, the project becomes less readable and understandable, making it more complicated to apply new changes to the system [63]. Finally, because of the codebase size, compilation can take a long time to finish. For instance, a complete build (i.e., recompiling the whole codebase from scratch) for Db2 takes about two to three hours on our test bed. This is an architectural necessity in Db2 because the crypto code resides in the components that many other higher-level components depend on.

- (4) Distributed teams (corresponding to Step 1 in 7E — Engage). Upgrading a complex system often involves collaborations among multiple teams or organizations. In this project particularly, we require expertise in two different areas, i.e., database engine (expert in the security component) and cryptography (research team and development team). Since the component that we need to upgrade in Db2 is mainly security related (i.e., the implementations of DH and TLS in Db2), we require expertise in Db2 security. During the PQC implementation, communication and collaboration among multiple geo-distributed teams (with different backgrounds) are required. A breakdown or lag in communication could cause delays in the development.
- (5) Technical debt because of hard coding (corresponding to Step 3 in 7E — Evolve). Technical debt is a concept in software engineering that reflects the extra work caused by previous work when choosing an easy solution instead of applying the best overall solution [64]. Hard coding is one of the most common decisions that lead to technical debt. Let us consider the case that we met in the evolution of encryption schemes — hard-coded cryptographic key lengths. Whenever you are required to increase the encryption key length (e.g., from 128-bit key to 256-bit key), to increase the level of cryptographic complexity, you are exposed to the risk of exceeding the bound of the key length (usually implemented in an array), leading to memory corruption and system crash.¹¹ In our project, not only the key lengths are hard-coded, but there also exist multiple and inconsistent key length checks, e.g., an “if” condition where key length has to be smaller than a specific constraint. In such a case, we need to loosen the constraints. However, those key length constraints are hard to trace because the key names can be changed, and the constraints are not necessarily at the same security layer (e.g., cryptographic protocol level vs. algorithm level).
- (6) Legacy encryption scheme on the client-side (corresponding to Step 3 in 7E — Evolve). As mentioned before, Db2 has been evolving for 34 years. Due to its long history, Db2 may still need to support older database client drivers that some customers may still be using in legacy environments. For example, the database security administrator may configure the Db2 server to allow DES instead of AES, typically in a tightly controlled environment. For various reasons (e.g., compatibility and/or cost issues), those customers keep using legacy software. To support those customers, not only do we need to maintain the compatibility of the security component for old versions of the software, but we also need to support customers when transitioning to new versions of the software.
- (7) Cryptographic APIs are in flux (corresponding to Step 6 in 7E — Execute). Both Kyber and Dilithium are still in development (NIST has not finalized the PQC standardization). To the best of our knowledge, the only commercial product that adopts Kyber (based on the version submitted to NIST competition Round 2) is Amazon Web Services Key Management Service [65]. At IBM, the applications of Kyber and Dilithium are not yet commercially supported. The Application Programming Interfaces (APIs) are not released for production

¹¹In C programming language, the compiler does not automatically check array bounds. Your array index may point to a virtual memory that is not mapped to a physical memory, which can cause a crash.

teams. Moreover, for security purposes, the API functions to access the PQC functionality are deliberately hidden by the developers of the cryptographic libraries (as they want to make sure that beta-level code does not accidentally end up in the production code), which makes it difficult for developers of the main product to integrate the functionality. In our case, we required technical support from the cryptography research team and GSKit development team to apply the APIs of Kyber and Dilithium to Db2. We first spent our time with the developers of the cryptographic libraries in online meetings trying to understand how to access the API functions. Then, we wrote the wrappers implementing the tricks needed to access the hidden functions. Once the cryptographic library APIs become officially available, the developers of the main product will have to replace the customized wrappers with the direct calls to the functions (which will not be a risky procedure, as long as the API calls do not change drastically).

- (8) Underestimation of sizing even though correct experts are brought in (corresponding to Step 5 in 7E — Estimate). As a cutting-edge and ever-changing technology, the application of PQC is still in its infancy, and developers have a steep learning curve for the application of PQC. In our project, though we involved multiple top experts at IBM, we realized that we had underestimated the size of the project (in terms of work time) during the development. While we hope that readers will be more accurate in their sizing estimates, we know that project size underestimation is a common issue in Software Engineering. We recommend that the readers err on the side of caution when estimating the amount of effort and resources required for PQC evolution.
- (9) An iterative and incremental approach is preferred to the waterfall-like model (corresponding to Step 2 in 7E — Exam). During the development, we realized that the uncertainties associated with our new code changes often require more changes in other components of the system. For instance, a small component change may lead to significant impact on the performance and/or functionality of the whole system. The reasons behind this are numerous: e.g., lack of understanding, lack of documentation, or code complexity. Thus, we advise adopting iterative and incremental methods in PQC evolution. In our case, we separate the PQC-related code into several modules and add these modules incrementally. We test the cryptographic functions and performance incrementally during the development.
- (10) Education is essential (corresponding to Step 4 in 7E — Educate). The education went both ways. The research part of the team educated the technical part of the team about the PQC-related challenges. In contrast, the technical experts educated researchers on the issues specific to the upgrade of the Db2 source code.
- (11) Maintain the PQC infrastructure (corresponding to Step 7 in 7E — Essay). Once the PQC is in place, the next step is to maintain the PQC infrastructure and prevent developers from introducing new security vulnerabilities. To achieve this goal, one can implement a cryptographic algorithm/code scanning policy to monitor future changes of cryptographic algorithms and ensure no quantum-unsafe algorithms are accidentally added to the code base. Human and technical resources should be allocated to the task by the management.

Take-away messages. Besides crypto-agility, we also mentioned other lessons learned in the previous section. From this experience, we summarize some takeaway messages. We would like to share these messages with practitioners who want to upgrade the current cryptographic algorithms to quantum-safe algorithms in their existing systems. We hope they will benefit from our suggestions. Our messages are as follows.

- (1) Prepare ahead and have a clear roadmap and timeline. As mentioned in Section 3 and Section 7 (Step 8), follow the 7E steps to update your existing cryptography and prepare for additional time for the development.
- (2) Get support from management. As mentioned in Section 3 (Step 1), PQC is a new emerging technology, it is important to educate your management and colleagues so that they are aware of the potential risk of data breach because of quantum attacks. It is more important to receive support from management so that the upgrade of the cryptography can be successfully proceeded with the collaboration from all the departments that need to be involved.
- (3) Collaborate with multiple departments. As mentioned in Section 7 (Step 4), the upgrade of cryptography can affect other components of the system or even cause failure of the system. Thus, we suggest that all stakeholders should be engaged in the upgrade as soon as possible to lower any risk to the system. Moreover, those experts will help to achieve the success of the project.
- (4) Prepare to pay the hard-coding debt. As mentioned in Section 7 (Step 5), technical debt can be a block in the way of your project. A thorough exam of the existing code is necessary to clean up the blocks. This is also one of the potential reasons that can cause delays in your project.
- (5) Document the development for future maintenance and evolution. As mentioned in Section 3 (Step 7) and Section 7 (Step 10), NIST is still working on the standardization of PQC. This implies that PQC will evolve. Thus, it is a good idea to keep all the records of your development to simplify future changes.
- (6) Plan for crypto-agility. The reason behind this is the same as the previous one — the evolution of PQC. As mentioned in Section 3 (Step 3) and Section 7 (Step 5), we need to design the new cryptography with crypto-agility (such as dynamic key sizes). Paying off technical debt at this stage will improve the productivity in future development.
- (7) Measure performance impact. As mentioned in Section 3 (Step 7), we need to assess the new cryptographic system’s performance and robustness. For performance-centric software systems, our finding shows that a PQC upgrade does not increase the time required to exchange cryptographic keys (and, on average, may reduce the timing slightly).

8 Discussions

Threats to validity. We cannot generalize our findings to all software products [66]. However, Db2 can be treated as a “critical case” (in a case study sense of the term [67]), showing the potential viability of 7E and highlighting some of the potential pitfalls. Db2 uses widespread encryption schemes, thus, our experience may be useful to other developers of software products. The 7E roadmap can be readily applied to other projects; thus, we encourage researchers and practitioners to apply 7E and our lessons learned to other instances of PQC evolution and are looking forward to their findings.

In the last 12 months, we have witnessed increased attention to the risk posed by quantum computing. Some leading institutions have published guidance to help organizations transition to quantum-safe encryption. While our 7E guidance predates these recent publications, it is worth noting that in essence our 7E and this recently published guidance are very much in synchronization. For instance, IBM proposes a framework that consists of four pillars — preparation, discovery,

transformation, and observability [68]. Preparation maps to our “Engage” and “Educate” principles, discovery maps to our “Examine” principle, transformation maps to our “Evolve”, “Estimate”, and “Execute” principles, and lastly, observability maps to our “Essay” principle. Similarly, the US Cybersecurity & Infrastructure Security Agency shares similar guidance that can also be grouped into three pillars — preparation, discovery, and transformation [69]. While the naming may be different, the essence of our 7E is very much in synchronization with the guidance recently emerging. Consequently, these guiding documents may serve as indirect evidence of the generalizability and applicability of our findings.

Further discussions. As we expected, the introduction of the encryption schema is not a trivial process. Some of the pain-points are anticipated and are typical for the evolution of a large-scale legacy product. They are associated with difficulties in comprehending a massive codebase, as well as fragile source code management and build environments.

However, some of the other problems are unexpected and are not typical. One example is the need to implement customized wrappers to call the PQC API functions because they are still in flux (see Challenge 7 in Section 7). Another example is the essentiality of education for people who are not aware of the urgency to ensure data security against quantum attacks (see Challenge 10 in Section 7).

It is not easy to follow the crypto-agility paradigm strictly in existing legacy systems. In the previous section, we have shown that we improved the crypto-agility of Db2 by introducing flexible key sizes. However, we have not solved this problem completely. For example, the authentication code in Db2 is written in the C programming language for performance reasons. C language lacks object-oriented (OO) concepts, such as abstraction. As a result, it is difficult to implement more complex objects on top of an abstraction. Though one can introduce OO principles by rewriting the code to C++, it will be expensive, time-consuming, and may lead to performance degradation, violating essential performance requirements.

9 Conclusions

This paper continued our work on the 7E roadmap for PQC evolution. While leveraging 7E, we conducted a case study to make IBM Db2 database management software quantum-safe. As a result, we validated the 7E roadmap throughout the project. Moreover, we pinpointed operational details needed to execute 7E, recorded lessons learned, delivered takeaway messages, and proposed an enhanced 7E for practitioners so that they could benefit from our experience and use it in their endeavours. Our data may also be of interest to academics as a building block in the theory of evolution of large mature software products.

For future work, we would like to apply the 7E roadmap and our best practices to other systems that are vulnerable to quantum attacks. At the same time, we would like to dig deeper into the topic of crypto-agility with PQC. One possible solution is to implement a hybrid classical-quantum cryptography. Finally, we would like to invite other practitioners and researchers to share their experiences with the usage of the 7E roadmap.

References

- [1] Lei Zhang, Andriy V. Miranskyy, and Walid Rjaibi. Quantum advantage and the Y2K bug: A comparison. *IEEE Softw.*, 38(2):80–87, 2021. doi: 10.1109/MS.2020.2985321. URL <https://doi.org/10.1109/MS.2020.2985321>.

- [2] Richard P Feynman. Simulating physics with computers. *International journal of theoretical physics*, 21(6):467–488, 1982.
- [3] Isaac L. Chuang, Neil Gershenfeld, and Mark Kubinec. Experimental implementation of fast quantum searching. *Phys. Rev. Lett.*, 80:3408–3411, Apr 1998. doi: 10.1103/PhysRevLett.80.3408. URL <https://link.aps.org/doi/10.1103/PhysRevLett.80.3408>.
- [4] IBM Corporation. IBM News room - 2016-05-04 IBM Makes Quantum Computing Available on IBM Cloud to Accelerate Innovation - United States, 2016. URL <https://www-03.ibm.com/press/us/en/pressrelease/49661.wss>.
- [5] Tameem Albash, Victor Martin-Mayor, and Itay Hen. Temperature scaling law for quantum annealing optimizers. *Phys. Rev. Lett.*, 119:110502:1–110502:7, Sep 2017. doi: 10.1103/PhysRevLett.119.110502. URL <https://link.aps.org/doi/10.1103/PhysRevLett.119.110502>.
- [6] Rami Barends, Alireza Shabani, Lucas Lamata, Julian Kelly, Antonio Mezzacapo, Urtzi Las Heras, Ryan Babbush, Austin G Fowler, Brooks Campbell, Yu Chen, et al. Digitized adiabatic quantum computing with a superconducting circuit. *Nature*, 534(7606):222–226, 2016.
- [7] D-Wave Systems Inc. D-Wave Announces General Availability of First Quantum Computer Built for Business, September 2020. URL <https://www.dwavesys.com/company/newsroom/press-release/d-wave-announces-general-availability-of-first-quantum-computer-built-for-business/>.
- [8] Sean Silcoff. Once a pioneer, quantum computer developer d-wave will start making same types of machines as rivals - the globe and mail. <https://www.theglobeandmail.com/business/article-quantum-computing-pioneer-d-wave-expanding-to-develop-same-type-of/>, 2021.
- [9] IBM Corporation. IBM Announces Collaboration with Leading Fortune 500 Companies, Academic Institutions and National Research Labs to Accelerate Quantum Computing - Dec 13, 2017, 2017. URL <https://newsroom.ibm.com/2017-12-13-IBM-Announces-Collaboration-with-Leading-Fortune-500-Companies-Academic-Institutions-and-National-Research-Labs-to-Accelerate-Quantum-Computing>.
- [10] IBM Corporation. Ibm’s roadmap for scaling quantum technology — ibm research blog, 2020. URL <https://www.ibm.com/blogs/research/2020/09/ibm-quantum-roadmap/>.
- [11] IBM Corporation. IBM Q Experience, 2021. URL <https://quantumexperience.ng.bluemix.net/qx/experience>.
- [12] Rigetti & Co, Inc. Rigetti computing, 2021. URL <https://www.rigetti.com/what>.
- [13] Google LLC. Google AI Blog: A Preview of Bristlecone, Google’s New Quantum Processor, 2018. URL <https://ai.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html>.
- [14] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. C. Pienti, M. Chmielewski, C. Collins, K. M. Hudek, J. Mizrahi, J. D. Wong-Campos, S. Allen, J. Apisdorf, P. Solomon, M. Williams, A. M. Ducore, A. Blinov, S. M. Kreikemeier, V. Chaplin, M. Keesan, C. Monroe, and J. Kim. Benchmarking an 11-qubit quantum computer. *Nature*

- Communications*, 10:5464, Nov 2019. ISSN 2041-1723. doi: 10.1038/s41467-019-13534-2. URL <https://doi.org/10.1038/s41467-019-13534-2>.
- [15] Honeywell International Inc. Quantum computer, 2021. URL <https://www.honeywell.com/us/en/company/quantum/quantum-computer>.
 - [16] Xanadu Quantum Technologies. Xanadu releases world’s first photonic quantum computer in the cloud, 2020. URL <https://www.newswire.ca/news-releases/xanadu-releases-world-s-first-photonic-quantum-computer-in-the-cloud-818969456.html>.
 - [17] Xanadu Quantum Technologies. Cloud platform - xanadu, 2021. URL <https://www.xanadu.ai/cloud>.
 - [18] Microsoft Corporation. Quantum computing | Microsoft, 2021. URL <https://www.microsoft.com/en-us/quantum/>.
 - [19] Tomi H. Johnson, Stephen R. Clark, and Dieter Jaksch. What is a quantum simulator? *EPJ Quantum Technology*, 1(1):10, Jul 2014. doi: 10.1140/epjqt10.
 - [20] Andriy Miranskyy and Lei Zhang. On testing quantum programs. In Anita Sarma and Leonardo Murta, editors, *Proceedings of the 41st International Conference on Software Engineering: New Ideas and Emerging Results, ICSE (NIER) 2019, Montreal, QC, Canada, May 29-31, 2019*, pages 57–60. IEEE / ACM, 2019. doi: 10.1109/ICSE-NIER.2019.00023. URL <https://doi.org/10.1109/ICSE-NIER.2019.00023>.
 - [21] Andriy Miranskyy, Lei Zhang, and Javad Doliskani. Is your quantum program bug-free? In *ICSE-NIER 2020: 42nd International Conference on Software Engineering, New Ideas and Emerging Results*, pages 29–32, Seoul, South Korea, 2020. ACM. doi: 10.1145/3377816.3381731. URL <https://doi.org/10.1145/3377816.3381731>.
 - [22] Andriy Miranskyy, Lei Zhang, and Javad Doliskani. On testing and debugging quantum software. *arXiv preprint arXiv:2103.09172*, 2021.
 - [23] Amazon. Introducing Amazon Braket, a service for exploring and evaluating quantum computing, 2019. URL <https://aws.amazon.com/about-aws/whats-new/2019/12/introducing-amazon-braket/>.
 - [24] Amazon. Amazon braket pricing, 2021. URL <https://aws.amazon.com/braket/pricing/>.
 - [25] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie

- Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. doi: 10.1038/s41586-019-1666-5.
- [26] Han-Sen Zhong, Hui Wang, Yu-Hao Deng, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Jian Qin, Dian Wu, Xing Ding, Yi Hu, et al. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463, 2020. doi: 10.1126/science.abe8770.
- [27] Charles Q. Choi. In the Race to Hundreds of Qubits, Photons May Have ”Quantum Advantage”. *IEEE Spectrum*, 2021. URL <https://spectrum.ieee.org/tech-talk/computing/hardware/race-to-hundreds-of-photonic-qubits-xanadu-scalable-photon>.
- [28] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, 2017.
- [29] Anton Robert, Panagiotis Kl Barkoutsos, Stefan Woerner, and Ivano Tavernelli. Resource-efficient quantum algorithm for protein folding. *npj Quantum Information*, 7(1):1–5, 2021.
- [30] Shouvanik Chakrabarti, Rajiv Krishnakumar, Guglielmo Mazzola, Nikitas Stamatopoulos, Stefan Woerner, and William J Zeng. A threshold for quantum advantage in derivative pricing. *arXiv preprint arXiv:2012.03819*, 2020.
- [31] Craig Gidney and Martin Ekerå. How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum*, 5:433, 2021. doi: 10.22331/q-2021-04-15-433. URL <https://doi.org/10.22331/q-2021-04-15-433>.
- [32] Jerry Chow and Jay Gambetta. Quantum Takes Flight: Moving from Laboratory Demonstrations to Building Systems — IBM Research Blog, 1 2020. URL <https://www.ibm.com/blogs/research/2020/01/quantum-volume-32/>.
- [33] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, October 1997. ISSN 0097-5397. doi: 10.1137/S0097539795293172. URL <http://dx.doi.org/10.1137/S0097539795293172>.
- [34] Lov K. Grover. A fast quantum mechanical algorithm for database search. In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 212–219, Philadelphia, PA, 1996. ACM. doi: 10.1145/237814.237866. URL <https://doi.org/10.1145/237814.237866>.
- [35] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86:032324, Sep 2012. doi: 10.1103/PhysRevA.86.032324. URL <https://link.aps.org/doi/10.1103/PhysRevA.86.032324>.
- [36] Andreas Baumhof. The deal with quantum computing and cryptography, 2019. URL <https://www.infosecurity-magazine.com/opinions/quantum-computing-cryptography-1-1-1/>.
- [37] Rui Chao and Ben W Reichardt. Fault-tolerant quantum computation with few qubits. *npj Quantum Information*, 4(1):1–8, 2018.
- [38] Jaime Sevilla and C Jess Riedel. Forecasting timelines of quantum computing. *arXiv preprint arXiv:2009.05045*, 2020.

- [39] Xanadu Quantum Technologies. Xanadu, 2021. URL <https://www.xanadu.ai>.
- [40] Thomas Lubinski, Sonika Johri, Paul Varosy, Jeremiah Coleman, Luning Zhao, Jason Necaie, Charles H Baldwin, Karl Mayer, and Timothy Proctor. Application-oriented performance benchmarks for quantum computing. *arXiv preprint arXiv:2110.03137*, 2021.
- [41] National Institute of Standards and Technology. Post-quantum cryptography. <https://csrc.nist.gov/projects/post-quantum-cryptography>, 2021.
- [42] Valerio Scarani, Helle Bechmann-Pasquinucci, Nicolas J. Cerf, Miloslav Dušek, Norbert Lütkenhaus, and Momtchil Peev. The security of practical quantum key distribution. *Rev. Mod. Phys.*, 81:1301–1350, Sep 2009. doi: 10.1103/RevModPhys.81.1301. URL <https://link.aps.org/doi/10.1103/RevModPhys.81.1301>.
- [43] Michele Mosca. Cybersecurity in an era with quantum computers: will we be ready? *IEEE Security & Privacy*, 16(5):38–41, 2018.
- [44] Daniel J. Bernstein and Tanja Lange. Post-quantum cryptography. *Nat.*, 549(7671):188–194, 2017. doi: 10.1038/nature23461. URL <https://doi.org/10.1038/nature23461>.
- [45] Fang Song. A note on quantum security for post-quantum cryptography. In Michele Mosca, editor, *Post-Quantum Cryptography - 6th International Workshop, PQCrypto 2014, Waterloo, ON, Canada, October 1-3, 2014. Proceedings*, volume 8772 of *Lecture Notes in Computer Science*, pages 246–265. Springer, 2014. doi: 10.1007/978-3-319-11659-4_15. URL https://doi.org/10.1007/978-3-319-11659-4_15.
- [46] Johannes Buchmann, Denis Butin, Florian Göpfert, and Albrecht Petzoldt. Post-quantum cryptography: State of the art. In Peter Y. A. Ryan, David Naccache, and Jean-Jacques Quisquater, editors, *The New Codebreakers - Essays Dedicated to David Kahn on the Occasion of His 85th Birthday*, volume 9100 of *Lecture Notes in Computer Science*, pages 88–108. Springer, 2016. doi: 10.1007/978-3-662-49301-4_6. URL https://doi.org/10.1007/978-3-662-49301-4_6.
- [47] National Institute of Standards and Technology. Post-quantum cryptography—workshops and timeline, 2021. URL <https://csrc.nist.gov/Projects/post-quantum-cryptography/workshops-and-timeline>.
- [48] Lily Chen, Stephen Jordan, Yi-Kai Liu, Dustin Moody, Rene Peralta, Ray Perlner, and Daniel Smith-Tone. Report on post-quantum cryptography. Technical Report NISTIR 8105, US Department of Commerce, National Institute of Standards and Technology, 2016.
- [49] Vasileios Mavroeidis, Kamer Vishi, Mateusz D. Zych, and Audun Jøsang. The impact of quantum computing on present cryptography. *CoRR*, abs/1804.00200, 2018. URL <http://arxiv.org/abs/1804.00200>.
- [50] Walid Rjaibi, Sridhar Muppidi, and Mary O’Brien. Wielding a double-edged sword, 2018. URL <https://www.ibm.com/thought-leadership/institute-business-value/report/quantumsecurity>.
- [51] Don Johnson, Alfred Menezes, and Scott A. Vanstone. The elliptic curve digital signature algorithm (ECDSA). *Int. J. Inf. Sec.*, 1(1):36–63, 2001. doi: 10.1007/s102070100002. URL <https://doi.org/10.1007/s102070100002>.

- [52] Evgeniy O. Kiktenko, Nikolay O. Pozhar, Maxim N. Anufriev, A. S. Trushechkin, R. R. Yunusov, Y. V. Kurochkin, A. I. Lvovsky, and Aleksey K. Fedorov. Quantum-secured blockchain. *CoRR*, abs/1705.09258, 2017. URL <http://arxiv.org/abs/1705.09258>.
- [53] IBM. Db2 11.5 - IBM documentation. <https://www.ibm.com/docs/en/db2/11.5>, 2021.
- [54] IBM Corporation. Distributed relational database architecture - ibm documentation. <https://www.ibm.com/docs/en/db2/11.5?topic=systems-distributed-relational-database-architecture-drda>, 2021.
- [55] Baha Hebrawi. *Open systems interconnection: upper layer standards and practices*. McGraw-Hill, Inc., 1993.
- [56] Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - kyber: A cca-secure module-lattice-based KEM. In *2018 IEEE European Symposium on Security and Privacy, EuroS&P 2018, London, United Kingdom, April 24-26, 2018*, pages 353–367, London, UK, 2018. IEEE. doi: 10.1109/EuroSP.2018.00032. URL <https://doi.org/10.1109/EuroSP.2018.00032>.
- [57] CRYSTALS Team. Kyber. <https://pq-crystals.org/kyber/>, 2020.
- [58] National Institute of Standards and Technology. Post-quantum cryptography. <https://csrc.nist.gov/Projects/post-quantum-cryptography/round-3-submissions>, 2021.
- [59] CRYSTALS Team. Dilithium. <https://pq-crystals.org/dilithium/>, 2020.
- [60] IBM Corporation. IBM global security kit (gskit) - IBM documentation. <https://www.ibm.com/docs/en/informix-servers/14.10?topic=protocol-global-security-kit-gskit>, 2021.
- [61] IBM Corporation. Introduction to global security kit installation - IBM documentation. <https://www.ibm.com/docs/en/db2/11.5?topic=encryption-global-security-kit-installation-gskit>, 2021.
- [62] Google LLC. HTTPS encryption on the web — Google transparency report. <https://transparencyreport.google.com/https/overview?hl=en>, 2019.
- [63] Capers Jones. *Applied software measurement*. McGraw-Hill Education, 2008.
- [64] techopedia. What is technical debt? — definition from techopedia. <https://www.techopedia.com/definition/27913/technical-debt>, 2017.
- [65] Alex Weibel. Round 2 post-quantum TLS is now supported in AWS KMS. <https://aws.amazon.com/blogs/security/round-2-post-quantum-tls-is-now-supported-in-aws-kms/>, 2020.
- [66] Roel J. Wieringa and Maya Daneva. Six strategies for generalizing software engineering theories. *Sci. Comput. Program.*, 101:136–152, 2015. doi: 10.1016/j.scico.2014.11.013. URL <https://doi.org/10.1016/j.scico.2014.11.013>.
- [67] Robert K. Yin. *Case Study Research: Design and Methods*. Applied Social Research Methods. SAGE Publications, 2009. ISBN 9781412960991.

- [68] IBM Institute for Business Value. Security in the quantum computing era. <https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/quantum-safe-encryption>, 2022.
- [69] Cybersecurity & Infrastructure Security Agency. Prepare for a new cryptographic standard to protect against future quantum-based threats. <https://www.cisa.gov/news-events/alerts/2022/07/05/prepare-new-cryptographic-standard-protect-against-future-quantum-based-threats>, 2022.