Designing and Delivering a Post-Quantum Cryptography Course

Thomas J. Borrelli
tjbcis@rit.edu
Department of Computer Science
Rochester Institute of Technology
Rochester, New York, USA

Monika Polak
mpolak@ur.rochester.edu
Department of Computer Science
University of Rochester
Rochester, New York, USA

Stanisław Radziszowski
spr@cs.rit.edu
Department of Computer Science
Rochester Institute of Technology
Rochester, New York, USA

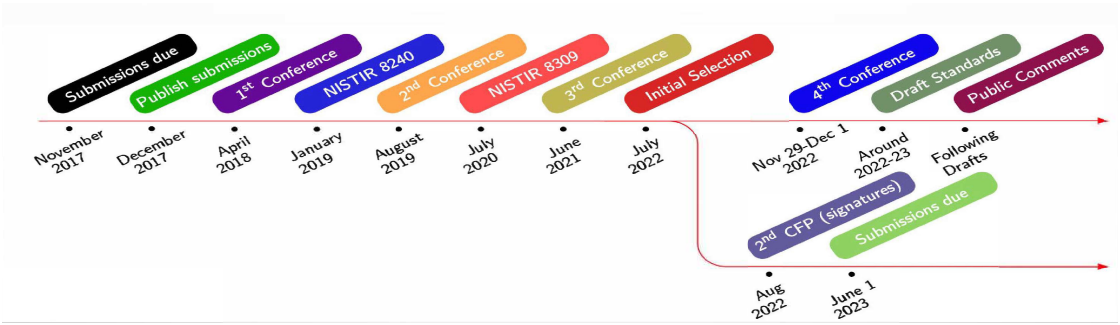

Figure 1: Timeline of the NIST Post-Quantum Cryptography initiative, from a presentation by Dustin Moody [31]

ABSTRACT

The security of many commonly used cryptographic protocols, especially public-key cryptosystems, would be compromised if general-purpose, large-scale, fault-tolerant quantum computers become a reality. In this paper we present our experience developing and launching a course in Post-Quantum Cryptography (PQC). PQC refers to cryptographic systems that are secure against both quantum and classical computers. Such systems may be achieved through classical (i.e. non-quantum) means.

Because of progress in the design of quantum computers and the ongoing National Institute of Standards and Technology (NIST) process to develop post-quantum cryptographic systems, we realized that there is a need to design a course that covers the consequences of developments in quantum computing (QC), the threats of QC to currently used cryptographic schemes, and how to mitigate those threats by developing quantum-resistant schemes. We designed a new course that is attracting students interested in the future of cryptography and computing security/cybersecurity. We first offered it as an MS-level graduate course, also open to upper-level undergraduates, in Spring 2022, and followed with the second offering in Spring 2023. The course covers the PQC algorithm design process, the consequences of QC, some computationally hard problems and then discusses selected proposals for post-quantum cryptosystems designed to be resistant to known classical and quantum attacks. Three main types of such designs include lattice-based, code-based, and hash-based schemes. These three types are used both for key encapsulation methods (KEM) and digital signatures (DS), and more generally for encryption and authentication.

SIGCSE 2024, March 20–23, 2024, Portland, OR, USA

ACM ISBN 979-8-4007-0423-9/24/03…$15.00
https://doi.org/10.1145/3626252.3630823

CCS CONCEPTS

• **Security and privacy** → Cryptography.

KEYWORDS

post-quantum cryptography, course design, quantum-resistant cryptography, curricula initiative

ACM Reference Format:
Thomas J. Borrelli, Monika Polak, and Stanisław Radziszowski. 2024. Designing and Delivering a Post-Quantum Cryptography Course. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2024), March 20–23, 2024, Portland, OR, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3626252.3630823

1 INTRODUCTION

We first overview the current state of development of QC and its consequences. We discuss the NIST standardization process and increased interest in PQC. This provides a good background of our motivation and a need for a PQC course.

Section 2 is devoted to the course description that allows enough detail for adoption by others. We also describe the motivating context for developing the course. Section 3 contains the course schedule.

In Section 4 we reflect on what did or didn't work, and how we improved the course in the second run. We provide grade distribution and SmartEvals data. We decided to add a separate Section 5 to discuss challenges. The paper ends with conclusions and future work in Section 6.

1.1 QC engenders PQC

Quantum computing (QC) has been a topic of interest especially since the early 1980s when Richard Feynman observed that some quantum mechanical effects could not be efficiently simulated on a classical computer, but that maybe a new type of computer could be built that specifically took advantage of quantum mechanical effects [39]. Since then, progress in this area has been slow but steady and consistent. In recent years, however, there are many indications that this progress has accelerated significantly, which in turn greatly increased attention of cryptographers to post-quantum cryptography (PQC). This attention quickly extended to the much broader cybersecurity community.

There has been much anticipation [33], excitement [2] and even hype [44] [16] around recent developments in QC. In 2019, Google announced that they had achieved so-called *Quantum Supremacy* [3], the point at which a quantum computer outperforms any classical computer. Google claimed to have solved a problem related to quantum circuits in just over 3 minutes, that would take a classical computer 10,000 years. IBM countered by pointing out that the problem was selected to give QC an advantage and that *Summit*, the fastest supercomputer at the time [43], could solve this problem in 2.5 days by using an alternative algorithm. Thus Google's claims seem to be not as strong as stated [9]. The current phase of development in QC is often referred to as *noisy intermediate-scale quantum* (NISQ) since robust error detection and correction is something currently lacking from today's quantum computers.

Despite these setbacks, significant progress is being made. If progress continues in this fashion, as seems likely, many of the commonly-used cryptographic systems will need to be reevaluated.

Protocols such as the Rivest, Shamir, Adleman (RSA) cryptosystem, Digital Signature Algorithm (DSA), and Elliptic-Curve DSA (ECDSA) c.f. [37, 42] are all *extremely* vulnerable to attacks on a quantum computer due to Shor's algorithm [41]. Shor's algorithm can be used to efficiently factor (for RSA) or efficiently find the period (for DSA and ECDSA). Shor's algorithms effectively means that these protocols can be broken in polynomial time on a quantum computer (BQP), whereas no known classical algorithm can solve these in polynomial time.

The need for protocols that are resistant to QC will necessitate a migration to new or upgraded secure cryptographic algorithms. We originally named the course *Post-Quantum Cryptography* to match the industry-standard term as well as the name of the algorithm selection process run by NIST [11]. Unfortunately, this name seems to cause some confusion among those who are under the impression that it means strictly *after* the arrival of quantum computers. To address this concern we modified the course description (see Section 2.3 below). Moving forward the name of the course will be *Quantum-Resistant Cryptography*, which the authors feel is a more appropriate descriptive term.

This paper builds on the work of the second author at the Research Institute of Science and Technology (RISAT) conference in Summer 2022 [38].

1.2 NIST PQC standardization process

With the looming threat of QC in mind, the National Institute for Standards and Technology (NIST) has recognized the need for replacement algorithms for RSA, DSA, ECDSA, and related protocols. Michele Mosca suggested in 2015 [32] that if we call D the time that we wish currently encrypted data to be secure for, T the time it takes to adopt PQC standards, and Q the time it takes to build a QC to break said standards, if $D + T > Q$, this is a concern. This observation has come to be known as "Mosca's Inequality." Therefore, now is the time to develop and test PQC standards [23].

NIST has initiated a PQC process to solicit, evaluate, and standardize one or more quantum-resistant public-key and digital signature cryptographic algorithms. This is an ongoing process that gained a lot of attention [11] [12]. The timeline of the process is presented in Figure 1, and finalized standards are expected to be ready by 2024.

Symmetric algorithms such as the AES currently do not appear to be strongly affected by the arrival of QC as Shor's algorithm [41] is not currently known to be applicable to achieve the speed-up that would render them vulnerable. Grover's algorithm [17], however, may apply but it would yield merely a *quadratic* not an *exponential* speed-up. In practical terms this means that AES-128, instead of having a search space size of 2^{128} would have an effective search space size of 2^{64}. Moving from AES-128 to AES-256 means that with Grover's algorithm, the effective search space would be 2^{128} (i.e. comparable to the current level of AES-128 in the absence of QC).

As of this writing November 13th, 2023), there is one public-key/KEM candidate algorithm that has been approved by NIST for standardization, CRYSTALS-Kyber. On the Digital Signature (DS) side, there are three approved algorithms: CRYSTALS-Dilithium, FALCON and, SPHINCS+. Draft standards of each of the above are available (except for FALCON which is slated to be released in 2024) [34–36]. There are also three additional KEM algorithms being further considered: BIKE, HQC and Classic McEliece [12].

SIKE had also been considered up through July of 2022, however, a critical vulnerability making it susceptible to attacks from classical, let alone quantum computers, was reported in a preprint paper by Castryck and Decru released in August of 2022 [8].

The authors are aware of the development of a few courses on QC (such as those reported in two SIGCSE 2023 papers [15] [27]), but none so far on PQC. Our initiative and the resulting course described in this paper seems to be the first one to be documented and disseminated.

1.3 General audience interest in PQC

In the last couple of years, almost every issue of one of the top magazines in computing, *Communications of the ACM* (CACM), contains at least one (and often several) items on post-quantum cryptography. These news items and general audience articles usually only just touch the technical issues, but they build up interest and often excitement of the readership. They are written by specialized writers in the area or experts from academia. A representative sample of such CACM items is as follows: overview of PQC by a science writer Don Monroe, 2023 [30]; editor Leah Hoffman interview with a top expert in QC Scott Aaronson, 2021 [19]; Lance

Fortnow's essay on 50 years of the P=NP question and how PQC is now modifying the overall perspective of complexity theory, 2022 [14]; an article by Torsten Hoefler, Thomas Häner and Matthias Troyer, whose title in itself, *Disentangling Hype from Practicality: On Realistically Achieving Quantum Advantage*, says quite a lot, 2023 [18]; an article by Brian LaMacchia titled *Security, the Long Road Ahead to Transition to Post-Quantum Cryptography*, 2022 [24]; and a review article by Petros Wallden and Elham Kashefi, 2019 [45]. There is no lack of PQC coverage in other venues, including: Daniel Bernstein and Tanja Lange in *Nature*, 2017 [5]; Kristin Lauter in the *Notices of the American Mathematical Society*, 2020 [25]; and an arXiv preprint by Matt Campagna, Brian LaMacchia and David Ott, *Post Quantum Cryptography: Readiness Challenges and the Approaching Storm*, 2020 [7].

Many readers could wish to understand more about PQC from such CACM items without much additional reading, but there seem to be no easy shortcuts. For our students, some of whom are among these readers, we think that our PQC course can give them great background for proper understanding of the summative articles. All this together can enhance student's interest in the area, to the point that in the future some of them may venture to write articles about PQC themselves.

2 COURSE OVERVIEW

Our main motivation to design the course was to offer our students a course that can give them an additional competitive advantage in the job market and bring them to (or very near) the cutting-edge of research being done in the field of PQC. There is ongoing progress on designing better quantum computers (e.g. IBM, Google, Microsoft, D-Wave) [10, 21, 22, 28]. If a general-purpose, large-scale, fault-tolerant quantum computers become a reality, then the current NIST public-key cryptographic standards will be vulnerable to attacks. Space is somewhat limited, so we invite you to view the full course syllabus that was used most recently in the Spring of 2023 (2022-23 AY) [6].

2.1 Pedagogical Approach

Whenever one is teaching a course, whether it's a new course or one that has been delivered many times, it is important to know what students have as background prior to the course and to teach in such a way that students will be able to follow the lectures. To this end, our course was taught using a "bottom-up" approach with many basic examples to help students work through the underlying concepts of the lecture of the day. These smaller examples could then be expanded upon and used as building blocks to create larger and more complex scenarios. This approach seemed more appropriate than taking a "top-down" approach of looking at a whole system and attempting to dissect it to its principal components. However, a top-down approach was successfully employed by graduate students in their graduate presentations, however.

The authors co-taught the course, each instructor taking a different topic or week of material. In this way we were able to take advantage of our combined expertise. For example, Polak has much experience with code-based cryptography, Borrelli has a background in physics which helps in explaining some of the quantum mechanics concepts underlying Quantum Computing, and Radziszowski has a significant experience with cryptography in general and lattice-based cryptography specifically.

We offered the course twice: in Spring 2022 and Spring 2023. The class capacity was 25. After the one week drop-add period, we had enrollment of 14 and 21 students, respectively.

2.2 Textbooks and Tools

There was no single textbook appropriate for this course, so we ended up utilizing material from several sources [4, 33, 37, 42]. In particular, the book [20] by Hoffstein, Pipher and Silverman [1] had an *excellent* chapter on lattice-based cryptography, while [4] provided good background in hash-based and code-based cryptography. We also provided students with an excellent primer on QC from Rieffel and Polak (no known relation to this paper's author), *An Introduction to Quantum Computing for Non-Physicists* [39].

At about half way through the course we introduced students to SageMath [40], a free open-source programming framework built on Python, which has implementations of many of the common cryptographic algorithms that are utilized in this course such as the gcd, Extended Euclidean Algorithm, LLL, CRT, and other mathematical algorithms. Many of the graduate students ended up utilizing SageMath in their presentations at the end of the term.

2.3 Course Description

The following course description is available for our students:

Post-Quantum Cryptography (PQC) refers to cryptographic systems that are secure against both quantum and classical computers. Such systems may be achieved through classical (i.e. non-quantum) means. The security of many commonly used cryptographic protocols (especially public-key cryptosystems) would be compromised if general-purpose, large-scale, fault-tolerant quantum computers become a reality. This course covers the consequences of Quantum Computing and why it poses a threat to currently used cryptographic systems, and then discusses potential Post-Quantum cryptosystems designed to be resistant to such attacks. Students should have background in cryptography such as that obtained by taking CSCI-462 (Introduction to Cryptography) or CSCI-662 (Foundations of Cryptography) or similar courses, or permission of the instructor.

Learning outcomes. Students will be able to:

- explain the need of post-quantum cryptography,
- describe and implement some post-quantum encryption algorithms and explain their security,
- explain the limitations of existing quantum computer models and how they affect classical cryptography, and
- explain the principles of post-quantum cryptography.

This course can be taken by undergraduates or graduate students having taken the introductory-level cryptography course. Linear algebra, which is a prerequisite for the intro cryptography course, will also have been taken by students prior to this course. As a result of taking this course, students will be able to conduct research at or near the cutting-edge of Post-Quantum/Quantum-Resistant Cryptography work and will have an advantage in the job market.

[1] Also the authors of the NTRU cryptosystem, a 3rd round finalist in the NIST PQC standardization process

Table 1: Topic Schedule (2023)

Week & Topics	
1	Introduction, consequences of QC, Review compromised cryptographic protocols
2	Consequences of quantum computing, NIST competition updates, overview of NP-completeness and NP-hardness
3	Quantum computing overview
4	Quantum algorithms including Shor's and Grover's
5	Preliminaries on lattices, shortest vector problem
6	NIST competition, basics of error correcting codes, Goppa codes
7	Goppa bounded decoding, syndrome decoding
8	Midterm Exam, lattice-based cryptography, the Lenstra-Lenstra-Lovász (LLL) algorithm
9	The Goldreich-Goldwasser-Halevi (GGH) and Hermite Normal Form (HNF) public-key cryptosystem
10	Lattice-based cryptography, NTRU preliminaries, NTRU and FALCON cryptosystems
11	Code-based cryptography, McEliece public-key cryptosystem, decoding attacks against McEliece
12	Overview of multivariate cryptography and hash-based cryptography
13 -14	Graduate student's presentations

2.4 Grade Components

The course had the following grade components for undergraduate students: homework assignments 31% attendance & participation 7%, midterm exam 31%, final exam 31%. For graduate students grade components were: homework assignments 25%, attendance/participation 5%, midterm exam 25%, final exam 25%, paper/presentation 20%. The presentation component was added to differentiate the level of difficulty and amount of work required by graduate students.

There were eight homework assignments that were typically of the "problem-set" format but some also required to write short programs or scripts. Eight homeworks were given and students had at least one week to work on them. No homework was due the week of the midterm exam nor the week leading up to the graduate student's presentations.

Since the course is offered for undergraduate and graduate students, we introduced an additional grade component for graduate students: graduate student's presentation and paper. Graduate students write a 5-10 page paper on the selected KEM or DS cryptosystem. In addition, graduate students present on the same topic. Presentations were between 25 and 35 minutes long and covered the details of the selected KEM or DS cryptosystem along with a (short) worked-through example demonstrating its use. Graduate student presentations went well during both offerings of this course, and there were two student presentations during the 2nd offering that were exceptional.

3 COURSE SCHEDULE

We planned topics for a 14 week semester with two 75 minute classes per week (see Table 1). We decided that we should start with reviewing widely used public-key algorithms based on factorization and discrete logarithm problems (RSA, DHKE, ECDHKE, DSA). Students would have been familiar with most of these concepts from the prerequiste course, but we felt it was worth review since many of the topics in later weeks depend upon this understanding. We then discussed consequences of current development of quantum computers and explained why we need new standards. We emphasized good student understanding of the ongoing PQC developments and how computationally hard problems (e.g. shortest vector problem or computational syndrome decoding) can be used to develop secure public-key algorithms: key exchange mechanism (KEM) and digital signatures (DS). We also covered some computational complexity and how QC are thought to be able to solve all problems in P, but *not* problems that are NP-Complete. The problems feasible to be solved on QC are represented by the computational complexity class BQP (Bounded-error Quantum Polynomial time, see Figure 2). We decided to focus on lattice-based and code-based public-key algorithms due to many submissions for the NIST process of these types.

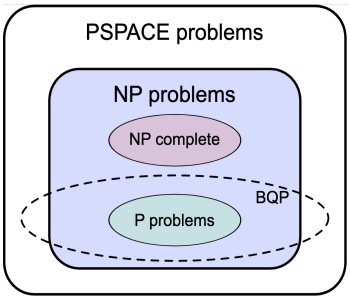

Figure 2: Conjectured relationship among some computational complexity classes neighboring BQP [29]

During week 3, we presented a basic overview of QC such as qubits, the Bloch sphere and basic quantum circuits (e.g. X, Y, Z, I). This week culminated in an example based on quantum dense coding. Week 4 was spent discussing Shor's algorithm [41] and utilized Scott Aaronson's excellent not-too-technical primer [1] to discuss the Quantum Fourier Transform. Grover's algorithm was also discussed. Week 5 laid the foundations for further discussion of lattices and lattice-based cryptography. Week 6 discussed recent developments in the NIST PQC standardization process and continued with the basics of error correcting codes, laying the foundation for code-based cryptosystems. Week 7 continued with error correcting codes.

In week 8 we had a midterm examination. Next, we introduced the famous LLL algorithm of Lenstra-Lenstra-Lovász [26] and discussed cyclotomic polynomials in the 2nd class of the week. Week 9 continued the theme of lattice-based cryptosystems with the Goldreich-Goldwasser-Halevi (GGH) cryptosystem. Week 10 discussed NTRU which was a Round 3 finalist in the NIST PQC standardization process, but did not proceed to the 4th round. However, the DS cryptosystem FALCON is based upon NTRU. Week 11 proceeded with code-based cryptography and the McEliece public-key cryptosystem. We also discussed decoding attacks against McEliece. In week 12, we discussed multivariate cryptography and hash-based cryptography. Finally, the last two weeks of class were devoted to graduate student presentations. The detailed course schedule, also available in our syllabus, is shown in Table 1.

4 DISCUSSION

As with any new course, there are some things that we felt worked well and some things that could be improved upon. While the sample sizes for both offerings of the course in Spring 2022 and Spring 2023 were relatively small (14 and 21, respectively), we feel as though there are important takeaways. According to combined data (2022/2023) that we collected from SmartEvals, most students considered that the amount of work was adequate and they consider the course very valuable (see Table 2).

4.1 What worked well and what to improve

We offered the course in Spring 2022 for the first time. Based on student feedback and our direct observations, we feel that the following went well:

- Good selection of topics
- Students liked the course overall
- Use of SageMath
- We had a smart group of students
- Good homework assignments design

After the first offering of the course we decided that the following needs improvement:

- Change the order of topics - do hard problems and related cryptosystems together
- Some in-class examples were too long
- Incorporate SageMath earlier in the course
- This course is not a great choice for inadequately prepared students, there is no way to prevent enrollment
- Not many ready examples or assignments that are publicly available for students for extra reading
- Update references

In 2023 we changed the order of topics (see Table 1) and we incorporated SageMath earlier in the course, which allowed us to save some time by showing fewer details of computations. We decided to drop the required textbook and used mixed resources as mentioned in Section 2.2.

4.2 Final grade distributions

We consider the course successful also from student's perspective. The majority of students received A grade. Final grades distributions can be found in Figures 3 and 4. The average grade in 2022 was 91.11 % with 90.97 % median. 85.42 % was the average grade in 2023 with increased median of 91.03 %. The grade of F in the Spring of 2023 was for a student who missed the graduate presentation as well as final exam.

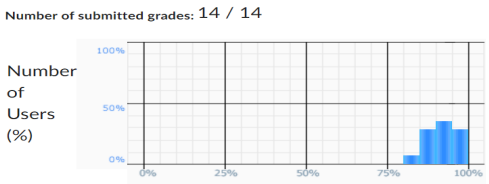

Figure 3: Grade distribution - Spring 2022

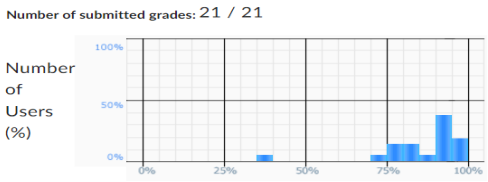

Figure 4: Grade distribution - Spring 2023

In regards to the point that the course is not a great choice for inadequately prepared students, the first two times this course was offered it was as a seminar course. Thus, enforcing the prerequisites had to be done manually. Moving forward, the permanent version of the course will automatically enforce prerequisites so the expectation is that this will be less of an issue.

5 CHALLENGES

5.1 Background of faculty and students

It is obvious that the ideal background required to understand in sufficient depth the subjects involved in this project is very broad. Several of its components are also often difficult even for people seasoned in cryptography and mathematics. In addition, topics on two extremes of the theory-practice spectrum are involved: some knowledge of quantum physics and mathematics behind it and the very practical but complex realities of cybersecurity industry. These extremes could hardly be further apart, yet understanding of both is needed to develop a reasonable expertise in PQC.

It is very difficult to find single faculty with sufficient background to develop and deliver such course. That's why there are three of us, so we can complement and support each other when numerous problems of understanding *what* and *how* emerge.

It is even harder to find students with appropriate background for this course. Technically, we require mainly a cryptography course as a prerequisite, but really more than that is needed, even if it is difficult to fully specify what it is. Certainly, there are many computing students who do not lack enthusiasm and often have a cheerful approach that they can handle it somehow. They are very curious about what the whole PQC buzz is, and they want to be part of it. They also think, and we support it wholeheartedly, that a computing security professional of the future should know PQC.

Table 2: Main data from SmartEvals, Spring 2022 and Spring 2023, (N=14)

Question Text	Str Disagree	Disagree	Neutral	Agree	Str Agree
Course objectives are valuable	0%	0%	0%	43%	57%
Student adequately prepared for course	7%	0%	0%	21%	71%
Student learned something of value	0%	0%	0%	14%	86%
Course well organized	0%	7%	7%	21%	64%
Course advanced student understanding	0%	0%	0%	21%	79%
Would recommend course to others	0%	0%	0%	29%	71%
Question Text	Sign. Less	Less	Adequate	More	Sign. More
Amount of work required	0%	7%	71%	7%	14%

Overall, this optimistic approach also had a very positive impact on faculty delivering this course. It is a challenge to address all the above issues together.

5.2 Level of details

Choosing the right level of detail to be presented in the course for each of the given topics is a challenge. On the QC side, take for example the Shor's [41] algorithm. It involves an important part of QC we have to understand both for integer factorization and discrete logarithms. Now, on top of standard cryptography of RSA, DSA and ECDSA, we have to discuss the quantum version of Fourier transform, provided the students already can handle linear algebra of qubit superpositions and entanglement. Furthermore, all of that has to fit in a couple of classes. After all is well planned, what should we do if we find out that some students have never heard of Fourier transform?

On the PQC side, take one of the frontrunner candidates for signatures in the NIST PQC standardization process, FALCON. This acronym's origin is based on the full name of this signature, *Fast Fourier lattice-based compact signatures over NTRU*. Should we stop here in the course and insert a couple hours-long fascinating story of NTRU? In the near future, FALCON must be presentable to security professionals in about an hour or two, yet presenting this with enough background material on FFT, hardness of lattices, and signatures over NTRU with cyclotomic modulus, remains a significant challenge. The authors feel as though we have struck a good balance of accessibility and level of academic rigor. At the same time, students these days like to explore the actual source code. We better get ready to explain how different software components of FALCON reference implementation correspond to the mathematical and cryptographic primitives involved.

5.3 Migration to PQC schemes

NIST is issuing warnings about a quickly approaching unavoidable task of migration to PQC protocols [13]. Calling it a task is a great understatement: most likely it will be a difficult decades-long process involving both the cybersecurity industry and academia. When PQC is needed, many, if not most, entities may not be ready to properly implement the transition. Thus, as urged by NIST, preparations for this transition must start now, even if the entire process may take many years. We should be essentially done with the transition when the power of existing QC approaches the point of legitimately threatening classical cryptography. How much time do we have? Nobody knows for certain, but more and more people predict that it will come sooner rather than later.

The industry must first realize which parts of their existing cybersecurity systems will be immune to QC as is, which may be adjusted by replacement of carefully chosen components, and which have to be eliminated entirely and replaced by new PQC framework. Current staff in industry may have hard time to embark on the implementation of these tasks. In contrast, the students who studied in a PQC course like ours, hopefully will take on the challenge. There are probably a number of new PhD graduates who have already studied PQC, however there's not a sufficient number of them to complete this migration alone. A large number of computing professionals with BS/MS degrees will be needed, and our course may be an important step to satisfy this demand.

6 CONCLUSION AND FUTURE WORK

Based on student feedback, and discussion with students and administration, we are convinced that the course was a great success as well as a much needed addition to the current security curriculum. In the ever-changing world of Quantum Computing, new security protocols will likely be necessary and increasingly relevant.

We anticipate continuing to co-teach this class at RIT, and starting next year (2023-24 AY) a version at the University of Rochester as well. It will also be interesting to see the final results of the PQC standardization process and recommendations from NIST over the next few years.

Moving forward, we note that this course will be called *Quantum-Resistant Cryptography*, which the authors feel is a more appropriate term and less apt to cause confusion than *Post-Quantum Cryptography*. We anticipate that more students will enroll in the course over time. Due to the ongoing NIST process, the course content will need to be modified to reflect the final recommendations.

ACKNOWLEDGMENTS

The authors would like to thank Professor Mohan Kumar for his ideas and support, as well as Professor Matthew Fluet for his support at the department and college curriculum committees and Professor Zack Butler for comments on an earlier draft of this paper. We'd like to acknowledge Professor Eustrat Zhupa for his feedback and moral support.

REFERENCES

[1] Scott Aaronson. 2007. Shor I'll do it. Retrieved June 1, 2023 from https://scottaaronson.blog/?p=208

[2] Scott Aaronson. 2013. *Quantum Computing Since Democritus.* Cambridge University Press, Cambridge, UK.

[3] Frank Arute, Kunal Arya, Ryan Babbush, et al. 2019. Quantum supremacy using a programmable superconducting processor. *Nature* 574 (October 2019), 505–510. https://doi.org/10.1038/s41586-019-1666-5

[4] Daniel J. Bernstein, Johannes Buchmann, and Erik Dahmen. 2009. *Post-Quantum Cryptography.* Springer, Berlin.

[5] Daniel J. Bernstein and Tanja Lange. 2017. Post-quantum cryptography. *Nature* 549 (October 2017), 188–194. https://doi.org/10.1038/nature23461

[6] Thomas J. Borrelli, Monika Polak, and Stanisław P. Radziszowski. 2023. Topics in Theory: Post-Quantum Cryptography / Quantum-Resistant Cryptography (CSCI-769). Retrieved June 6, 2023 from https://cs.rit.edu/~tjb/pqc/syllabus_2225.html

[7] Matt Campagna, Brian LaMacchia, and David Ott. 2020. *Post Quantum Cryptography: Readiness Challenges and the Approaching Storm.* Computing Community Consortium (CCC). Retrieved June 11, 2023 from https://arxiv.org/abs/2101.01269

[8] Wouter Castryck and Thomas Decru. 2022. An efficient key recovery attack on SIDH. Cryptology ePrint Archive, Paper 2022/975. https://eprint.iacr.org/2022/975

[9] Adrian Cho. 2019. *IBM casts doubt on Google's claims of quantum supremacy.* IBM. Retrieved June 7, 2023 from https://www.science.org/content/article/ibm-casts-doubt-googles-claims-quantum-supremacy

[10] D-Wave. 2023. *Unlock the Power of Practical Quantum Computing Today.* D-Wave. Retrieved June 2, 2023 from https://www.dwavesys.com/

[11] National Instutute for Standards and Technology (NIST). 2022. *Post-Quantum Cryptography Standardization.* NIST. Retrieved May 31, 2023 from https://csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization

[12] National Instutute for Standards and Technology (NIST). 2022. *PQC Standardization Process: Announcing Four Candidates to be Standardized, Plus Fourth Round Candidates.* NIST. Retrieved May 31, 2023 from https://csrc.nist.gov/News/2022/pqc-candidates-to-be-standardized-and-round-4

[13] National Instutute for Standards and Technology (NIST). 2023. *Migration to Post-Quantum Cryptography: Preparation for Considering the Implementation and Adoption of Quantum Safe Cryptography (Preliminary Draft), SP 1800-38.* NIST. Retrieved June 5, 2023 from https://csrc.nist.gov/publications/detail/sp/1800-38/draft

[14] Lance Fortnow. 2022. Fifty Years of P vs. NP and the Possibility of the Impossible. *Commun. ACM* 65, 1 (January 2022), 76–85. https://doi.org/10.1145/3460351

[15] Adrian German, Marcelo Pias, and Qiao Xiang. 2023. On the Design and Implementation of a Quantum Architectures Knowledge Unit for a CS Curriculum. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2023).* Association for Computing Machinery, New York, NY, USA, 1150–1156. https://doi.org/10.1145/3545945.3569845

[16] Roger A. Grimes. 2020. *Cryptography Apocalypse: Preparing for the Day When Quantum Computing Breaks Today's Crypto.* John Wiley & Sons, Inc., Hoboken, NJ.

[17] Lov K. Grover. 1996. A Fast Quantum Mechanical Algorithm for Database Search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (Philadelphia, Pennsylvania, USA) *(STOC '96).* New York, NY, USA, 212–219. https://doi.org/10.1145/237814.237866

[18] Torsten Hoefler, Thomas Häner, and Matthias Troyer. 2023. Disentangling Hype from Practicality: On Realistically Achieving Quantum Advantage. *Commun. ACM* 66, 5 (May 2023), 82–87. https://doi.org/10.1145/3571725

[19] Leah Hoffmann. 2021. Q&A Exploring the Promise of Quantum Computing. *Commun. ACM* 64, 12 (December 2021), 119–120. https://doi.org/10.1145/3490319

[20] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. 2014. *An Introduction to Mathematical Cryptography* (2nd ed.). Springer, Berlin, Germany.

[21] IBM. 2022. *Secure your enterprise for the quantum era.* IBM. Retrieved June 2, 2023 from https://www.ibm.com/quantum

[22] Alphabet Inc./Google. 2023. *Explore the possibilities of quantum.* Google. Retrieved June 2, 2023 from https://quantumai.google/

[23] Ilyas Khan. 2017. Mosca's inequality - why it matters. LinkedIn. Retrieved June 14, 2023 from https://www.linkedin.com/pulse/moscas-inequality-why-matters-ilyas-khan-ksg

[24] Brian LaMacchia. 2022. Security, the Long Road Ahead to Transition to Post-Quantum Cryptography. *Commun. ACM* 65, 1 (1 2022), 28–30. https://doi.org/10.1145/3498706

[25] Kristin Lauter. 2020. How to Keep Your Secrets in a Post-Quantum World. *Notices of the American Mathematical Society* 67, 1 (January 2020), 22–29. https://doi.org/10.1090/noti2004

[26] A.K. Lenstra, H.W. Lenstra Jr., and L. Lovász. 1982. Factoring polynomials with rational coefficients. *Math. Ann.* 261, 4 (Dec 1982), 515–534. https://doi.org/10.1007/BF01457454

[27] Jonathan Liu and Diana Franklin. 2023. Introduction to Quantum Computing for Everyone: Experience Report. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1* (Toronto ON, Canada) *(SIGCSE 2023).* New York, NY, USA, 1157–1163. https://doi.org/10.1145/3545945.3569836

[28] Microsoft. 2023. *Azure Quantum Accelerating scientific discovery.* Microsoft. Retrieved June 2, 2023 from https://azure.microsoft.com/en-us/solutions/quantum-computing/#overview

[29] mike1024. 2007. BQP complexity class diagram. Retrieved June 13, 2023 from https://commons.wikimedia.org/wiki/File:BQP_complexity_class_diagram.svg

[30] Don Monroe. 2023. Post-Quantum Cryptography. *Commun. ACM* 66, 2 (February 2023), 15–17. https://doi.org/10.1145/3575664

[31] Dustin Moody. 2022. *Looking into the future.* NIST. Retrieved June 6, 2023 from https://csrc.nist.gov/csrc/media/Presentations/2022/nist-pqc-looking-into-the-future/images-media/session-1-moody-looking-into-future-pqc2022.pdf

[32] Michele Mosca. 2015. Cybersecurity in an era with quantum computers: will we be ready? Cryptology ePrint Archive, Paper 2015/1075. https://eprint.iacr.org/2015/1075

[33] Michael A. Nielsen and Isaac L. Chuang. 2000. *Quantum Computation and Quantum Information.* Cambridge University Press, New York.

[34] National Institute of Standards and Technology. 2023. FIPS 203 (Initial Public Draft) Module-Lattice-Based Key-Encapsulation Mechanism Standard. Retrieved November 13, 2023 from https://csrc.nist.gov/pubs/fips/203/ipd

[35] National Institute of Standards and Technology. 2023. FIPS 204 (Initial Public Draft) Module-Lattice-Based Digital Signature Standard. Retrieved November 13, 2023 from https://csrc.nist.gov/pubs/fips/204/ipd

[36] National Institute of Standards and Technology. 2023. FIPS 205 (Initial Public Draft) Stateless Hash-Based Digital Signature Standard. Retrieved November 13, 2023 from https://csrc.nist.gov/pubs/fips/205/ipd

[37] Christof Paar and Jan Pelzl. 2010. *Understanding Cryptography.* Springer, Berlin.

[38] Monika Polak. 2022. *Designing a post-quantum cryptography course.* RIT. Retrieved June 6, 2023 from https://sites.google.com/risat.org/alb/conf/algebra-2022/talks?authuser=0

[39] Eleanor G. Rieffel and Wolfgang Polak. 2000. An Introduction to Quantum Computing for Non-Physicists. arXiv:quant-ph/9809016 [quant-ph]

[40] SageMath.org. 2023. *SageMath - Open-Source Mathematical Software System.* SageMath. Retrieved June 7, 2023 from https://www.sagemath.org/

[41] P.W. Shor. 1994. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science.* IEEE, Santa Fe, NM, USA, 124–134. https://doi.org/10.1109/SFCS.1994.365700

[42] Douglas R. Stinson and Maura B. Paterson. 2019. *Cryptography: Theory and Practice* (4th ed.). CRC Press, Boca Raton, FL.

[43] Top500.org. 2019. *November 2019, Top500 The List.* Top500.org. Retrieved June 14, 2023 from https://top500.org/lists/top500/2019/11/

[44] Mark van Rijmenam. 2022. How Quantum Computing Will Change The World. Retrieved June 7, 2023 from https://www.thedigitalspeaker.com/quantum-computing-change-world/

[45] Petros Wallden and Elham Kashefi. 2019. Cyber Security in the Quantum Era. *Commun. ACM* 62, 4 (April 2019), 120–129. https://doi.org/10.1145/3241037

Received 18 August 2023; accepted 02 October 2023