QUANTUM THREAT TIMELINE REPORT 2023

Authors

Dr. Michele Mosca

Co-Founder & CEO, evolutionQ Inc.

Dr. Marco Piani

Senior Research Analyst, evolutionQ Inc.

GRI GLOBAL RISK INSTITUTE

evolution

DECEMBER 2023

Contents

Declaration on Potential Conflict of Interest

evolutionQ Inc. offers services and quantum-safe cybersecurity products to help clients deploy and manage quantum-safe technologies across their networks.

The basis of the report is a survey of leaders in the fields of quantum computing research and commercialization.

.

Summary

Cybersecurity protocols that are widely used today rely on computational challenges believed to be practically unsolvable with classical computers. We have known for decades that the advent of quantum computers would allow some of those challenges to be overcome in a meaningful timeframe, posing severe risks to cybersecurity.

To mitigate this threat, new classical and quantum-based cryptographic techniques considered or definitively known to be immune to quantum attacks can be used. Unfortunately, such a transition is no easy task. It involves creating and implementing new hardware and software, developing standards, and updating older systems. Crucially, a successful transition hinges on proactive technology lifecycle management, rather than reactive crisis management, and will take considerable time.

The urgency of moving to quantum-safe cryptography varies for each organization, based on its security needs and risk tolerance. This urgency can be gauged using three primary factors:

- the *shelf-life time*: how many years the data must remain secure for;
- the *migration time*: how many years it will take to securely upgrade the systems guarding that data;
- the *threat timeline*: the estimated time until potential adversaries gain access to quantum computers of cryptographic significance.

The mitigation of the quantum threat to cybersecurity requires a transition to quantum-safe cryptography that can be implemented safely only with enough time at disposal.

This report is part of a series that aims at providing an educated perspective of how far away the quantum threat is, by collecting and examining the perspectives of global experts from academia and industry, involved in diverse facets of quantum computing.

These experts – 37 this year – respond to a questionnaire crafted by evolutionQ Inc. and aimed at gleaning valuable insights on the cyber-risk posed by quantum cryptanalysis.

Additional value is provided by the serial nature of the reports, which allows one to gauge the dynamics of quantum computing research – e.g., is progress accelerating? – based on potential changes in the experts' estimates tracked survey after survey.

Predicting the pace at which a Cryptographically-Relevant Quantum Computer (CRQC) will be developed – let alone *when* it will be developed – is plagued by uncertainty. The reason is that building such a device requires continuously pushing the boundaries of science and engineering.

Despite existing challenges, the polled experts generally accept that a CRQC will eventually be built on the basis that no specific fundamental roadblock has been identified and that there has been steady progress. Quantum researchers and companies have identified and continued to achieve key milestones in their roadmaps to further scale the size and performance of current devices towards the level needed for cryptographic applications.

The expert responses we collected suggest that the quantum threat might rise to prominence much faster than many might anticipate. For example, almost half (46%) felt it was more than 5% likely already within a 10-year timeframe and more than a quarter of respondents indicated a likelihood of about 50% or more.

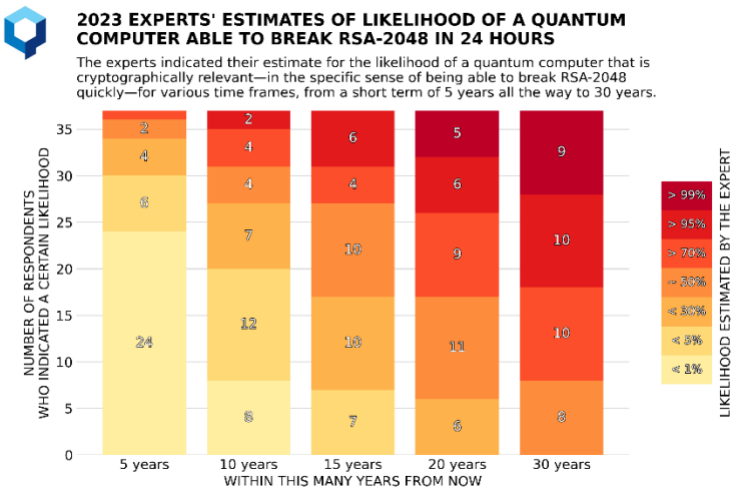

The responses can be averaged to produce an overall opinion-based estimated likelihood. We note that, depending on the risk tolerance and needs of companies and institutions, such estimates may correspond already to an intolerable risk that needs to be mitigated through immediate action.

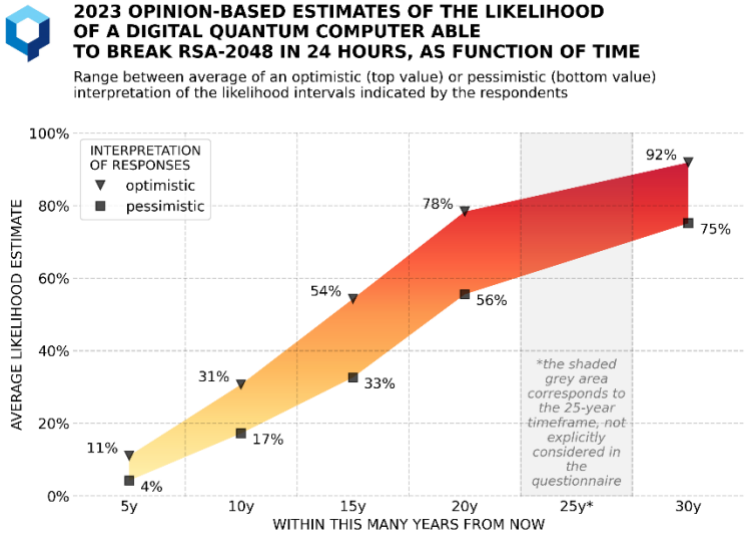

The main challenge in building a CRQC lies in the fragility of physical qubits, the building blocks of quantum computation. Quantum error correction (QEC) can harness multiple imperfect physical qubits to form stable logical qubits. Despite advancements in logical encoding, scaling up to numerous logical qubits needed for quantum cryptanalysis remains daunting. Many believe rapid strides in QEC theory and implementation might hasten the development of a CRQC. This suggests caution when relying on current best estimates for timing and planning the transition to quantum-safe cybersecurity.

On the other hand, economic uncertainties, high interest rates, the appeal of other disruptive technologies like artificial intelligence, and facing global issues like climate change may be slowing investments in quantum computing, which skyrocketed in recent years. There is also doubt about quantum computing devices achieving practical commercial benefits without a large error-corrected quantum computer. Coupled with a quantum computing "hype", these factors pose the risk of disillusion of investors. Importantly, slowing or even diminishing investments could defer the quantum threat timeline.

The experts' likelihood estimates for when a cryptographically relevant quantum computer will appear suggest that some companies might already be facing an intolerable risk requiring urgent action.

Independently of the exact time when a CRQC may become available, it is crucial to note that adversaries do not have to remain inactive while waiting for it: they can currently intercept, duplicate, and archive encrypted communications for eventual decryption later on – a so-called "Harvest Now, Decrypt Later (HNDL)" attack strategy. This is the rationale at the base of the aforementioned Mosca inequality, which takes into account the required shelf-life time of the data.

Those responsible for managing cyber-risk should not wait to act and solutions are available today. Given the recent advancements in quantum computing, the expert opinions collected in our survey, the momentum generated by the currently significant investments in the field, and the threat posed by the HNDL attack, there should be a conscious effort towards developing crypto-agility and building layered defenses against the quantum threat. This proactive approach can also help to mitigate the risks associated with a hasty transition to quantum-safe cryptographic tools and infrastructure.

From the threat timeline to the migration timeline

All organizations should evaluate their urgency in proceeding with migration to quantum-safe systems, depending on specific shelf-lives, migration times, and risk appetite. The Global Risk Institute and evolutionQ Inc. have already made available a quantum risk assessment methodology (Mosca and Mullholland 2017) on which such a process may be based.

1 Introduction

In this Introduction and in the Appendix, we provide some background information to understand both why and how quantum computers pose a threat to cybersecurity and why and how building such computers is an incredible scientific and technological challenge.

1.1 Quantum computing

Quantum mechanics serves as our most accurate framework for understanding the intricacies of nature at the microscopic scale, shedding light on the behavior of fundamental particles like electrons. Importantly, quantum phenomena are highly sensitive to disturbances. Interactions between a quantum system and its environment often diminish or entirely erase its quantum properties through a process known as *decoherence*. This phenomenon, along with the relevant physical scales involved, largely explains why quantum effects are not immediately apparent in our daily experiences.

Quantum computers leverage unique quantum phenomena to handle information in a manner radically different from current classical computers.

In computing, information needs a physical medium for storage and manipulation. In current and so-called classical computers, a standard bit—representing either a "0" or a "1"—is encoded in physical system akin to lightbulbs or switches, which are either "off" or "on". Can the principles of quantum mechanics be harnessed to store and process information in a fundamentally different manner? The concept of quantum computing arose from exploring this very question. Stemming from ideas initially proposed by physicist and Nobel laureate Richard Feynman (Feynman 1982), quantum computing aims to tackle complex problems in physics that are virtually intractable for classical computers (Nielsen and Chuang 2000).

The primary obstacle in advancing quantum computing is the unprecedented requirement to maintain and control quantum behaviour at a level that has never been attempted before in human history.

The foundational unit of quantum information in quantum computing is the *quantum bit*, or *qubit*. Unlike a traditional bit that stores either a 0 or a 1, a qubit can exist in a superposition of both states. This means both values can be thought of as "coexisting" and can be processed simultaneously. The monumental challenge in the field of quantum computing is to maintain and control these fragile quantum features. This involves mitigating and counteracting the effects of decoherence.

There are multiple approaches to building a quantum computer, varying both in the choice of physical substrate to create physical qubits—from superconducting circuits and trapped ions to quantum optics, among others—and in the strategies for implementing *quantum error correction* (QEC) all the way to so-called *fault tolerance*. These latter schemes are crucial for encoding quantum information in more robust *logical qubits*, rather than in the inherently imperfect physical qubits, thus allowing for reliable information processing. A key milestone along the path towards a quantum computer is that of demonstrating that QEC schemes allow one to go beyond the so-called "break-even" condition, that is, that encoded logical qubits perform better than the underlying physical qubits.

Once realized, quantum computers will not only fulfill Feynman's vision of simulating quantum systems but also, by cleverly leveraging quantum features like superposition through specialized algorithms, solve a range of mathematical, optimization, and search problems at speeds unattainable by classical computers (Nielsen and Chuang 2000).

For more details on physical implementations, QEC, and fault tolerance, please refer to the Appendix.

1.2 Quantum threat to cybersecurity

Commonly used public-key cryptographic algorithms, such as the Rivest–Shamir–Adleman (RSA) cryptosystem (Rivest, Shamir, and Adleman 1978), are based on mathematical challenges believed to be insurmountable for classical computers. RSA, for example, is predicated on the complexity of factorizing large composite numbers into their prime components.

Quantum computers have the potential to undermine these cryptographic systems. Specifically, RSA could be compromised through the use of Shor's quantum algorithm (Shor 1994). Additionally, Grover's algorithm (Grover 1996) allows a quantum computer to search a solution space consisting of 2^n values in roughly $2^{n/2}$ steps, thereby weakening symmetric-key cryptography.

Quantum computers pose a threat to cybersecurity because they can break or weaken widely used cryptographic schemes.

The advent of quantum computing poses a risk of catastrophic failures in cyber-systems, either through direct cryptographic attacks or by eroding trust. This looming threat of a Cryptographically-Relevant Quantum Computer (CRQC) can be mitigated through the adoption of quantum-safe cryptographic methods, which can either be conventional or quantum-based.

The first category involves employing cryptographic algorithms grounded on problems that are believed to be difficult even for quantum computers. Progress has been made in this area, evidenced by the US National Institute of Standards and Technology (NIST) selection of the first *post-quantum* cryptographic due in 2024 (NIST 2023). The second type of quantum-safe tools leverage quantum phenomena themselves, as in the case of quantum key distribution (Nielsen and Chuang 2000).

Transitioning to this new breed of quantum-safe cryptography is a complex and delicate process (Mosca 2013): it requires the development and deployment of hardware and software solutions, the establishment of standards, the migration of legacy systems, and more[1].

With the necessity to devote enough time to an orderly and safe transition to a 'post-quantum world', the urgency for any organization to complete the transition to quantum-safe cryptography for a particular cyber-system can be assessed by considering three simple parameters[2]:

[1] As an example of the needed 'migration time', it is worth stressing that the NIST selection process started in 2016 (NIST 2016).

[2] Often, these parameters have respectively been called *x*, *y*, *z* in literature; see e.g., (Mosca 2013). Here we adopt a more explicit notation.

- $T_{\text{SHELF-LIFE}}$ **(shelf-life time)**: the number of years the information should be protected by the cyber-system;
- $T_{\text{MIGRATION}}$ **(migration time)**: the number of years needed to migrate the system properly and safely to a quantum-safe solution;
- T_{THREAT} **(threat timeline)**: the number of years before the relevant threat actors will be able to break the quantum-vulnerable systems.

Figure 1 The timeline for the emergence of quantum computers capable of threatening cybersecurity needs to be juxtaposed with the combined time required for migrating to post-quantum security and the duration for which the data needs to be protected. See main text for details.

If $T_{\text{SHELF-LIFE}} + T_{\text{MIGRATION}} > T_{\text{THREAT}}$, that is, if the time required to migrate the system *plus* the time for which the information needs to be protected goes *beyond* the time when the quantum threat will become concrete, then an organization may not be able to protect its assets for the required $T_{\text{SHELF-LIFE}}$ years against the quantum threat (see Figure 1). This is the content of the *Mosca Inequality* (Mosca 2013).

Organizations need to assess $T_{\text{SHELF-LIFE}}$ and T_{THREAT}. The difference $(T_{\text{MIGRATION}})^{\text{MAX}} := T_{\text{THREAT}} - T_{\text{SHELF-LIFE}}$ is the **maximum available migration time**, that is, the maximum time organizations have at their disposal to safely realize the transition.

Rushing the process of migration to post-quantum cryptography might itself create security issues which could be exploited even by attackers who use only traditional methods.

It is essential to understand that a hasty migration to post-quantum security systems can introduce new vulnerabilities which could be exploited using conventional hacking methods. These vulnerabilities might arise from oversights, design flaws, or implementation errors. There could also be issues related to interoperability and backward compatibility, complicating the transition.

While the security shelf-life time $T_{\text{SHELF-LIFE}}$ is generally a business decision or dictated by regulations, assessing the threat timeline T_{THREAT} is more complex. There are many scientific and engineering obstacles to developing a quantum computer potent enough to crack existing cryptographic systems. These challenges suggest that CRQCs are likely years away, but breakthroughs, which are by nature unpredictable, could potentially fast-track development.

Financial investments into quantum computing and related technologies are another significant factor influencing the rate of progress. Funding has surged in recent years from diverse sources, including government bodies, established companies, and private investors backing new startups (Kung and Fancy 2021). This influx of capital underscores the urgency to consider a carefully managed transition to post-quantum cryptographic systems.

1.3 Quantum computing before achieving fault tolerance

Present leading quantum processors are composed of tens-to-hundreds of physical qubits and cannot sustain fault-tolerant quantum computation. Such systems are known as *noisy intermediate-scale quantum (NISQ) systems* (Preskill 2018).

NISQ devices, despite their current limitations, are a testament to our growing ability to manipulate quantum systems. There is a massive push to discover how these devices—or their imminent successors—can be beneficial even before fully developed quantum computers are at our disposal. Demonstrating their practical value would bolster and validate continued investments in this domain. There is also ongoing research focused on affirming that our advancements in quantum computation have expanded the boundaries of possible computations.

The term "quantum supremacy"[3] (Preskill 2018) broadly refers to a quantum device's capability to execute computations that are practically unattainable for classical computers, regardless of the computation's practical value. Establishing clear-cut criteria for quantum supremacy is challenging. This is because it is tough to determine that no classical method, even with the most advanced supercomputers or optimal classical algorithms, can achieve the same computation within a "reasonable" timeframe. Even if we restrict ourselves only to known classical algorithms, the goalpost for quantum supremacy keeps shifting as classical computers and their algorithms evolve. Google claimed to have reached quantum supremacy in 2019 (Arute et al. 2019) and the 2020 version of this report gathered expert views on the importance of this claim result (Mosca and Piani 2021). Since then, while there have been enhanced demonstrations of quantum supremacy, Google's initial claim faced challenges due to advancements in classical algorithms and computing.

There's undeniable excitement about the practical and business potential of "early-stage" quantum computers that aren't yet advanced enough to threaten cybersecurity. For those primarily wary of the cybersecurity risks posed by quantum computers, the interest in these nascent quantum applications may seem indirect. However, such applications would:

- offer tangible signs and early alerts for the impending quantum challenges to cybersecurity;
- increase the likelihood of consistent funding and resources for quantum computing research aimed at developing a digital quantum computer with cryptographic significance.

[3] This terminology is somewhat controversial because it recalls, e.g., racial supremacy, but it has been widely used in literature, in the same way in which, e.g., "air supremacy" may be used in warfare jargon. In our context, "quantum supremacy" indicates superiority of quantum computers over classical computers for some specific task(s), in some strictly technical sense. Nonetheless, also considering the controversy, the quantum computing community has often chosen to refer to the same superiority as "quantum primacy", "quantum advantage", or similar.

KEY POINTS

- Quantum computing is a new paradigm for computers that leverages properties from quantum mechanics, making it different and for some computational tasks much more efficient than traditional computing, using specially designed quantum algorithms. The fundamental unit of information is the quantum bit or qubit.
- A quantum computer able to run quantum algorithms for cryptanalysis poses a threat to many widely used cryptosystems; at sufficient scale, it constitutes a Cryptographically-Relevant Quantum Computer (CRQC).
- The development of a CRQC will require error correction to encode and manipulate logical qubits, thus overcoming the inherent fragility of quantum features.
- It is an enormous scientific and engineering challenge: it requires maintaining a high level of quality and control of physical qubits while scaling their number. Given that developing a CRQC means pushing the boundaries of science and engineering, estimating when a CRQC may finally be built is a very difficult task itself.
- Mitigating the cyber-risk posed by a CRQC requires moving to quantum-safe cryptographic tools. The Migration Time will be different for each organization, but such a change requires a significant transition time and must be done thoughtfully to avoid introducing additional vulnerabilities or implementation errors.
- In addition to the Timeline for a CRQC and the Migration Time, how urgent the transition is for a given organization depends on one more parameter: the shelf-life time for which the data needs to stay secure.
- If the time required to migrate the system plus the time for which the information needs to be protected goes beyond the time when the quantum threat will become concrete, then an organization may not be able to protect its assets for the required time against the quantum threat. This is the content of the Mosca Inequality.

2 Scope of this report

This document presents the results of a survey conducted by evolutionQ Inc., with the participation of 37 internationally leading experts on quantum computing. Following similar surveys conducted in the past four years, we asked the experts to complete an online questionnaire on the state of development of the field. For some, we gave the option to answer a key question via email. More details on the questions that were asked are available in Appendix A.3 .

We stress that we aim both to provide a snapshot of the experts' opinions and to identify potential trends in the evolution of such opinions in time. This evolution may be due to steady progress, to new key developments or challenges identified, and to any additional circumstances which may be considered as "external" to research per se yet still affect research activity, like funding levels.

> *"I love these reports, and frequently refer people to them. To me they represent the best current synthesis of expert opinion on timelines for fault-tolerant quantum computation."*
>
> RESPONDENT

In creating the questionnaire, we try to be concrete and specific when it comes to considering quantum computers as a threat to cybersecurity. For this reason, the most important question speaks explicitly of breaking RSA-2048, whose security is based on the difficulty of factoring a 2048-bit number.

Other approaches have been taken to try to gauge the timeline for the creation of a fault-tolerant quantum computer that may threaten cybersecurity. For example, in (Sevilla and Riedel 2020), the authors try to forecast future progress in the domain of quantum computing by extrapolating past progress in the field. They look at relevant metrics—roughly speaking, at how many effective logical qubits are available for computation. Sevilla and Riedel focus on superconducting implementations, and, similarly to what we do, on the task of breaking RSA-2048. Their estimates for when (superconducting) quantum computers could achieve such a feat are described by the authors themselves as "one piece of relevant evidence that can supplement expert opinion" and "more pessimistic but broadly comparable to those produced through the survey of experts in [(Mosca and Piani 2019)]". They also write that a CRQC could be built earlier than estimated by them, if progress is faster than what one can extrapolate from current trends. Such an extrapolation suffers at the very least from the fact that the field of quantum computing is relatively young, so that the progress achieved and tracked so far still covers only a limited temporal span.

Relevant indications about the quantum threat timeline come also from the roadmaps of companies working towards the realization of fault-tolerant quantum computers (see, e.g., the Google and the IBM roadmaps).

KEY POINTS

- This report is part of a series based on annual surveys to collect and analyze opinions of tens of leading experts in the field of quantum computing.
- The major goal of the report is to provide unique insight into the quantum threat timeline based on expert opinions, complementing other approaches and sources of information.

3 Participants

Since the inaugural survey in 2019, we have annually reached out to top international experts with the aim of garnering a diverse and insightful array of perspectives on the progress in the quantum computing field. Throughout the years, we have endeavored to maintain the original group of 2019 respondents to monitor shifts in their views. Additionally, we have approached other potential participants, chosen from an extensive list of over a hundred preeminent experts. Those who agreed to participate were requested to fill out the online survey.

Some candidate respondents we contacted did not reply to our invitation, while some others declined. Overall, in 2023 we were able to collect responses from 37 experts (see Appendix A.1 for a complete list).

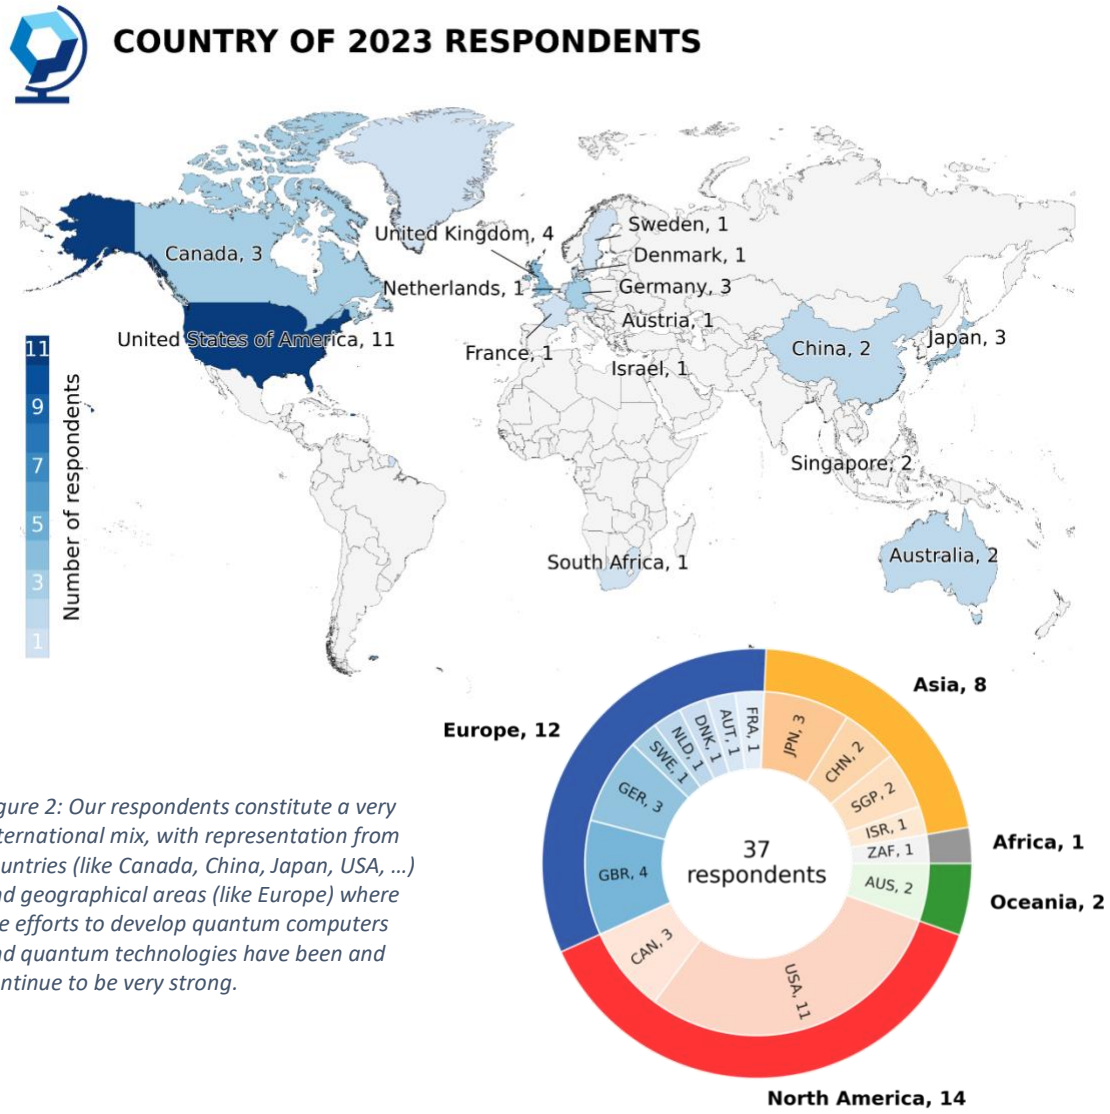

Figure 2: Our respondents constitute a very international mix, with representation from countries (like Canada, China, Japan, USA, …) and geographical areas (like Europe) where the efforts to develop quantum computers and quantum technologies have been and continue to be very strong.

Here we summarize graphically the composition of the group in terms of:

- country where they work (Figure 2),
- kind of activity they lead (Figure 4), and
- kind of organization they belong to (Figure 3).

The captions of the figures provide guidance in interpreting the presented statistics.

In essence, our respondent pool showcases a rich blend of expertise, national backgrounds, and representation from both academic and private sectors, aptly reflecting the multifaceted nature of the leading figures in the quantum computing community. Over the years, there has been a noticeable uptick in academics from our survey who also engage in corporate roles, signifying a heightened focus on the commercial aspects of quantum technologies and computing.

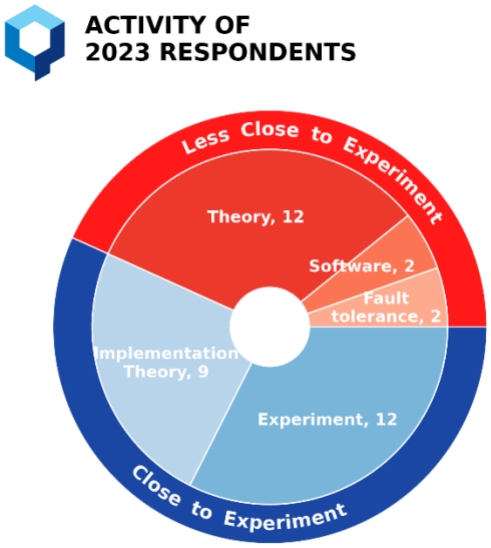

Figure 4 Our respondents cover a wide range of research activities. While the major division is between non-experimental research and experimental one, research that is not directly experimental can be very different. E.g., implementation theory focuses on guiding, supporting, and, in general, facilitating experimemental effort. Respondents are classified under simply "theory" if their more abstract activity is not specificically related to experiments or implementations, or to fault-tolerance, or to software development.

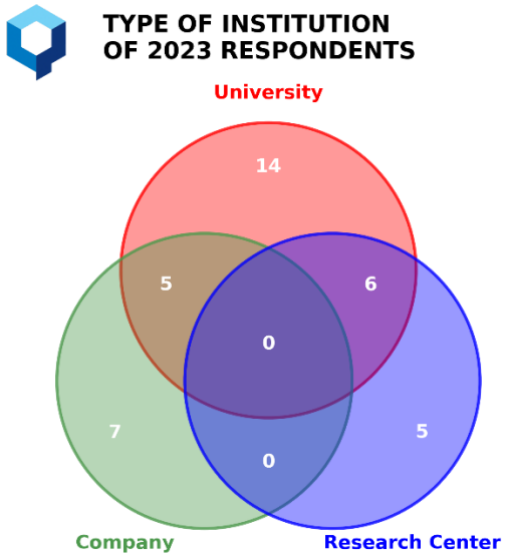

Figure 3 Most of the respondents work at universities, but many work at companies or research centres. Some researchers/academics may have some role in—or at least collaborate closely with—external companies. A larger fraction of our respondents has fallen in the latter category in the last reports, also because some past academic respondents have joined or founded companies.

KEY POINTS

- Our respondent pool showcases a rich blend of expertise, national backgrounds, and representation from both academic and private sectors, aptly reflecting the multifaceted nature of the leading figures in the quantum computing community.
- Thirty-seven respondents took part in the 2023 survey. Eighteen of these have taken part in all the five annual surveys run so far.

4 Survey results

We provide an aggregate quantitative analysis of the key responses about the following:

- the potential of various physical implementations/platforms for quantum computing (Section 4.1);
- the quantum threat timeline (Section 4.2);
- views on potential concerns regarding the realization of a cryptographically-relevant quantum computer in the relatively-near future (Section 0);
- the expected change in funding in support of quantum computing research (Section 4.7.1);
- the status and potential development of the so-called "quantum race" (Section 4.7.2);
- the potential sources of unexpected speed-up in the development of a cryptographically relevant quantum computer (Section 4.8).

"Predicting the future is hard, and as a crutch, people (especially scientists) have a bad habit of extrapolating forward linearly based upon past trends. This likely drives most of the respondents' ultra-long estimates in this survey. However, the historical pattern of transformative technologies is not like this at all, it is wildly nonlinear."

STEPHANIE SIMMONS
Photonic Inc.
& Simon Fraser University,
Co-Chair of Canada's National
Quantum Strategy Advisory Council

We also include:

- a selection of opinions about:
  - key recent research developments, as highlighted by the respondents;
  - near-future (that is, approximately, by mid-2024) developments that the respondents see as essential on the path to developing a fully scalable fault-tolerant quantum computer;
  - next milestones to track, not necessarily attainable by mid-2024;
- a collection of other notable remarks made by the respondents.

Comments by the respondents may be quoted with the respondents' permission. A quote may be attributed to the specific respondent or may be reported anonymously as coming from a "Respondent". Quotes may be lightly edited for conciseness and clarity.

Where we deem appropriate, we analyze shifts in the responses as compared to responses from the last four years. In the aggregated analysis of the responses, we indicate how many of the respondents (alternatively, what percentage of them) chose a specific answer among the many possible ones, when dealing with multiple choices. Not all the 37 respondents provided an input for all questions. Moreover, while the *number* of respondents has stayed relatively stable, there have been some changes in the *composition* of the pool of respondents. Finally, some questions might have been modified or tweaked in their wording from survey to survey, but we have intentionally kept the key question about breaking RSA-2048 the same.

These considerations suggest caution in interpreting any trend that may appear via a simple comparison with past responses, as it is challenging to disentangle confounding factors (see also the Appendix). Nonetheless, where we notice a trend that could potentially be significant, we point it out, and, where feasible and/or appropriate, we try to provide a rationale that may explain it.

4.1 Physical realizations

With respect to the physical realizations of quantum computers, we asked the respondents to indicate the potential of several physical implementations as candidates for fault-tolerant quantum computing. We specifically asked the respondents to consider the goal of implementing a quantum computer with ~100 logical qubits in the next 15 years.

The responses indicate a significant consensus that the present leading platforms are superconducting systems and trapped ions (Figure 5). This is consistent with the opinions collected in the preceding surveys.

Recent progress in quantum information processing with cold atoms and integrated photonics is reflected in the increase in the years in the number of experts who see such platforms as having potential, being very promising, or altogether being lead candidates. Hybrid implementations were mentioned under the "Other" category by several respondents.

> *"The industrial efforts in semiconductor qubits are continuing to expand successfully, and are accompanied by promising qubit performance. [..] Superconducting systems continue to grow and, importantly, many [alternative devices] are being developed."*
>
> RESPONDENT

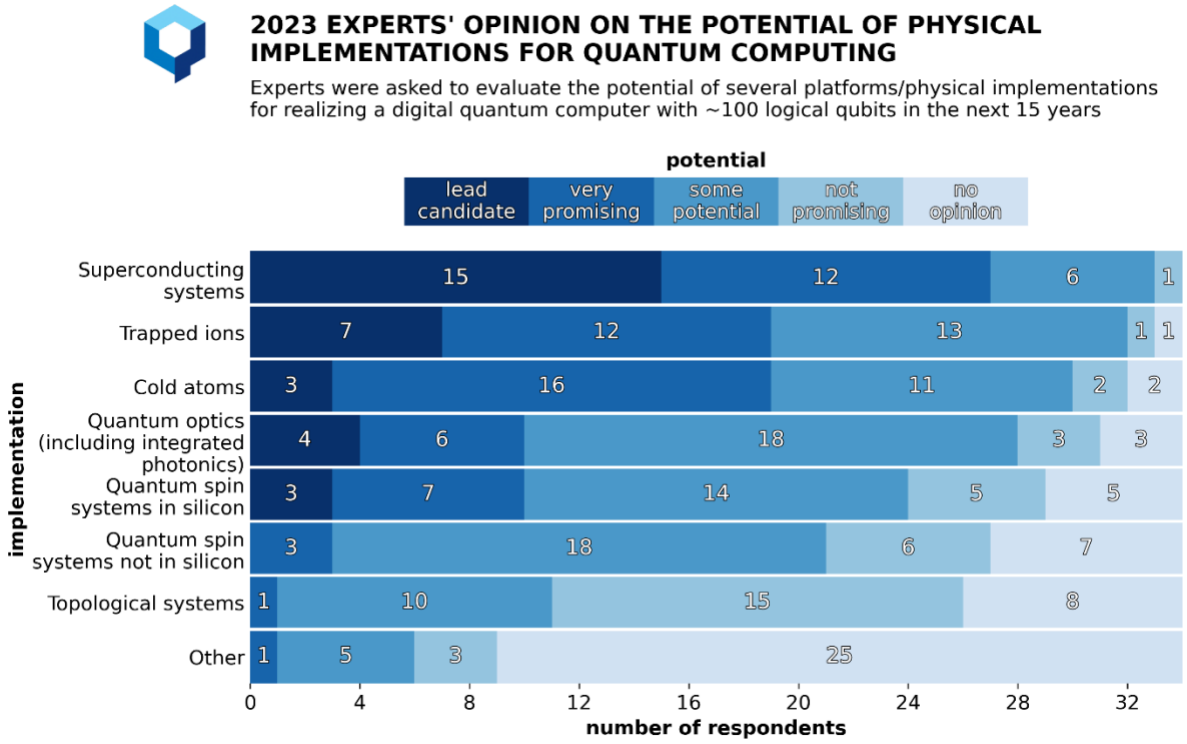

Figure 5: Similarly to previous years, superconducting-system implementations, followed by ion-trap implementations, are perceived as presently having some edge over other physical realizations. Nonetheless, many respondents point to other platforms as having high potential, all the way to being considered lead candidates.

Nicolas Menicucci, a professor at RMIT University, points out that implementations differ significantly in the nature of individual "low-level" qubits themselves:

> *Some architectures, such as trapped ions or the transmon qubits in superconducting architectures, have a natural interpretation as the qubits being the material systems themselves. [..] In such architectures, the meaning of a "logical qubit" is straightforward: It requires encoding logical quantum information in a multitude of physical qubits and doing measurements and operations that preserve and manipulate this encoded information. In contrast, [in] bosonic systems such as optics or microwave cavities [..] the qubits are not made of matter; they are created and manipulated by the material system [and t]here is a "level-0" question to be asked [..], which is how to encode the qubits.*

Menicucci stresses that this means that bosonic-system implementations work with fundamental "physical qubits" that can be already seen as simple "logical qubits" that include some 'built-in' form of error correction at the lowest possible level of encoding.

Simon Benjamin, a professor at the University of Oxford, writes:

> *Regarding ion traps, several of the lead research teams seem to be targeting very high-fidelity physical gates so as to run [without] error correction, and/or small [error-correcting] codes.*

In other terms, Benjamin points out that some researchers are not necessarily focusing on the goal of creating a digital quantum computer with information encoded in logical qubits – which is currently the only known path to efficiently breaking RSA-2048. They may rather try to create a high-quality programmable quantum device which could be useful in ways alternative to running ideal quantum algorithms, for example for quantum simulations or tackling optimization problems.

KEY POINTS

- Several physical implementations of quantum computers are presently being developed; they differ in the kind of physical system that constitutes the fundamental qubit. Each implementation has strengths and weaknesses, which become even more relevant when considering the need to scale to a large number of qubits.
- While certain implementations, like superconducting devices and trapped ions, may currently be considered as leading the efforts towards a CRQC, many other implementations are promising and showing progress.
- There might not be just one winner; different kinds of physical systems may end up being integrated in modular fashion to make the most of the advantages of different implementations.

4.2 Quantum factoring

In this survey, the most directly relevant information about the quantum threat timeline comes from the experts' assessment of the likelihood of realizing a quantum computer able to break RSA-2048 in a short time, in response to the following question (see also Appendix A.3):

Q: Please indicate how likely you estimate it is that a quantum computer able to factorize a 2048-bit number in less than 24 hours will be built within the next 5 years, 10 years, 15 years, 20 years, and 30 years.

Estimates on the practical requirements to achieve such a feat, also considering the imperfections of physical implementations, were presented for example in (Gheorghiu and Mosca 2017)[4] and in (Gidney and Ekerå 2021).

The primary findings of our yearly survey are illustrated in Figure 6, which provides the aggregate distribution of the responses of the experts[5]. It depicts the estimated increase of the likelihood of the quantum threat as we transition from the near future to the more distant one. Many participants in our annual surveys have emphasized the inherent challenges in making such predictions.

Some key features of the collection of likelihood estimates are summarized in Table 1.

> *"I see no chance of this happening within 5 years, and very low within 10 years. Beyond that point, I leave the odds at 50%, simply because it's not yet clear whether we will have physical qubit performance good enough to support the necessary error correction codes efficiently."*
>
> RESPONDENT

We note that there is large variability among the opinions of the experts: some lean towards optimism, while others are more cautious about the pace at which quantum computers will be developed. This is also illustrated in Figure 7 and Figure 8; in the latter, the individual pattern of responses for each expert is displayed.

For some respondents, their highest estimated likelihood for the quantum threat peaks before the 30y mark. For a subset of these, such maximum likelihood is less than the highest possible in our survey. Such perspectives can perhaps be seen as the acknowledgment of potential unforeseen technological hurdles or even insurmountable barriers (see also Section 0).

[4] Further updates consider recent developments and complement from a more technical perspective the present opinion-based series of reports (Gheorghiu and Mosca 2021).
[5] The same data are provided in a more data-sharing-friendly table in Appendix A.4 .

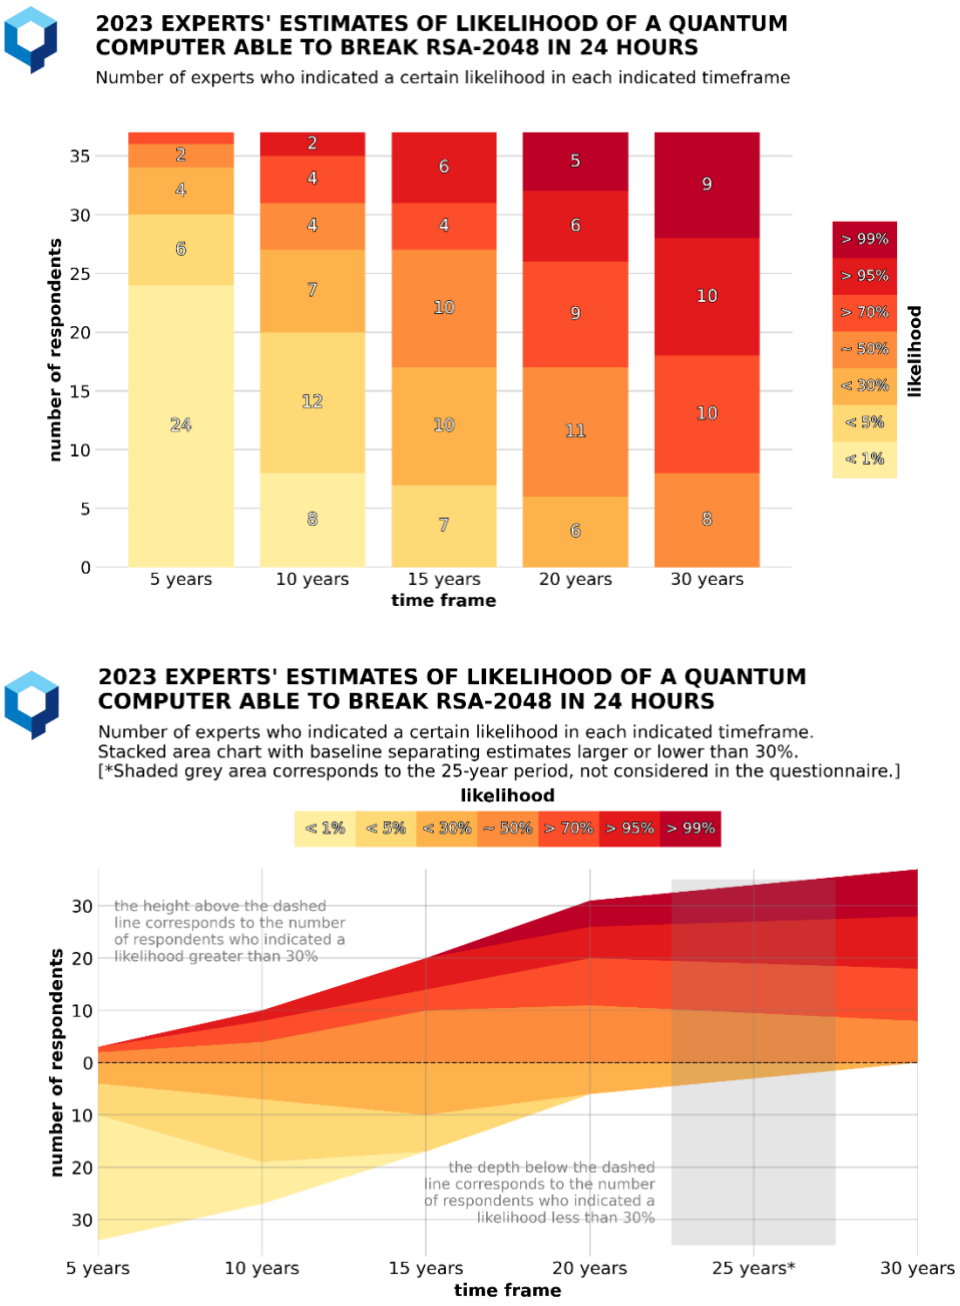

Figure 6 This figure illustrates the central information collected through our survey. The experts were asked to indicate their estimate for the likelihood of a quantum computer that is cryptographically relevant—in the specified sense of being able to break RSA-2048 in 24 hours—for various time frames, from a short term of 5 years all the way to 30 years.Top: stacked barchart with explicit indication of the number of experts estimating a certain likelihood. Bottom: stacked area chart conveying the same information, but allowing one to better appreciate the shift in likelihood estimates moving from short-term to long-term timeframes. Please note the inclusion of a dummy 25y timeframe.

TIMEFRAME	WHAT ONE MAY EXPECT BASED ON THE EXPERTS' OPINIONS ON THE LIKELIHOOD OF A QUANTUM COMPUTER ABLE TO BREAK THE RSA-2048 CRYPTOSYSTEM
NEXT 5 YEARS	Most experts (24/37) judged that the threat to current public-key cryptosystems in the next 5 years is "<1% likely". About a sixth of them (6/37) judged it relatively unlikely ("<5% likely"). The rest selected "<30%" or "about 50%" likely, with a single expert indicating a likelihood ">70%". Overall, *there seems to be a non-negligible chance of an impactful surprise within what would be considered a very short-term future.*
NEXT 10 YEARS	More than half of the respondents (20/37) still judged the event is less than 5% likely but more than a quarter of them (10/37) felt it was "about 50%" or more likely, suggesting *there is a significant chance that the quantum threat becomes concrete in this timeframe.*
NEXT 15 YEARS	A majority (20/37) of the respondents indicated "about 50%" likely or more likely, among whom ten indicated a ">70%" likelihood or higher. That is, w*ithin this timeframe, a majority of respondents assigns to the existence of cryptographically relevant quantum computer an about even likelihood or better.*
NEXT 20 YEARS	More than 83% (31/37) of respondents indicated "about 50%" or more likely, with 30% (11/37) pointing to ">95%" or ">99%" likely: *within this timeframe, the realization of the quantum threat appears to be seen as substantially more likely than not.*
NEXT 25 YEARS	We did not directly probe this timeframe in our questionnaire, as we believe the unavoidable uncertainty involved in the estimates does not warrant such a fine-grained distinction between what may happen between 20 years and 30 years from now. In some graphs, this timeframe may be included by showing interpolated values, for the sake of preserving a linear timescale.
NEXT 30 YEARS	Twenty-nine experts out of 37 (78%) indicated that the quantum threat has a likelihood of 70% or more this far into the future, with about a quarter of the experts (9/37) indicating a likelihood greater than 99%: *in general, there is a relatively low expectation of issues that would prevent a cryptographically-relevant quantum computer from being realized in the long run.*

Table 1 Summary analysis of the experts' likelihood estimates at the core of the present report.

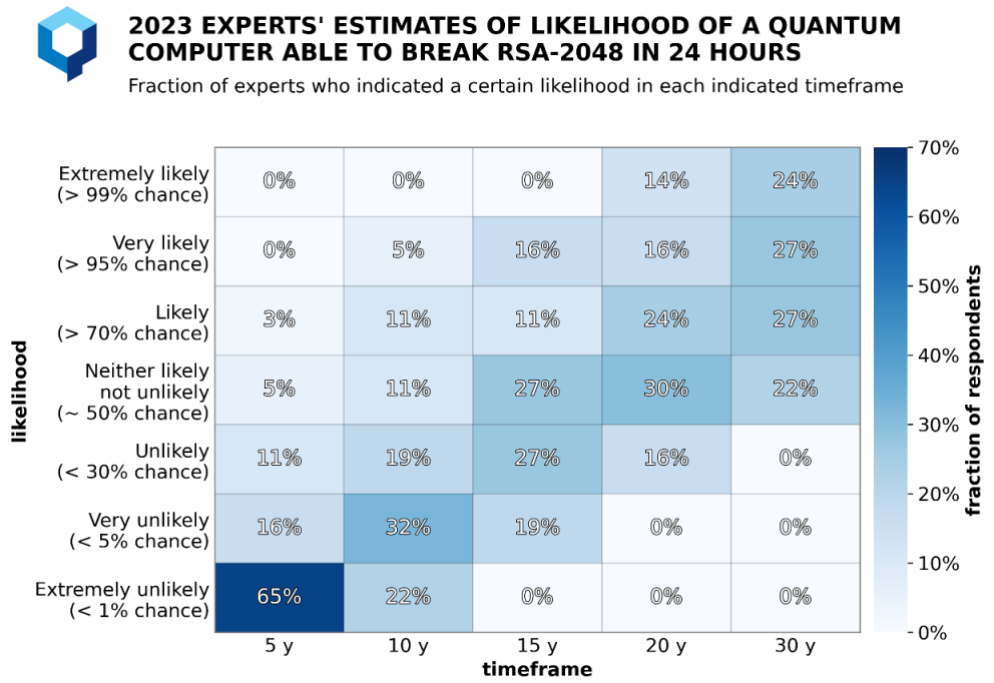

Figure 7 Heatmap and percentages for the distribution of the likelihood estimates of the 2023 survey, displaying the diversity in the opinion of the experts.

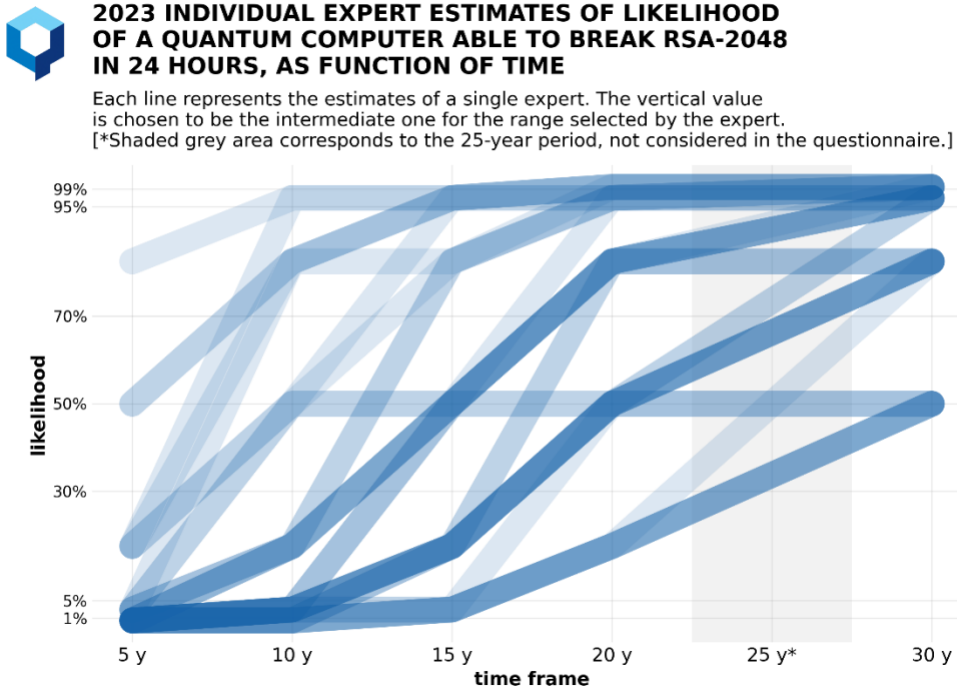

Figure 8 This figure illustrates the likelihood estimates of the individual experts, represented as curve growing in time. This plot allows one to appreciate not only the significant variance of the estimates for each timeframe considered, but also the diversity in how each expert estimates the likelihood will grow in time. One can nonetheless identify more common and more similar "trajectories" that are visually more opaque in this kind of plot.

Coarse-grained likelihood estimates

We aim to summarize succinctly the insight that the experts provided, to arrive at some single likelihood estimate. We will do this by averaging the estimates of the experts.

We may interpret the choice of one of the likelihoods, e.g., "*likely*", as the indication of a numerical probability in the range associated to it, i.e., in this case, a probability greater than 70% but less than 95%. We do not know what the best point estimate by each expert could have been. We take a conservative approach and consider the two extreme alternatives where each respondent is assigned either the higher or the lower of the extreme values of the range they picked. This can be roughly described as considering a "pessimistic interpretation" or, alternatively, "optimistic interpretation" of the answers' ranges. This approach allows us to calculate an average cumulative probability distribution for both interpretations. Had each respondent selected a precise estimate within the respective ranges, then the average estimate for the likelihood would sit in the range between the optimistic-interpretation and pessimistic-interpretation curves. In turn, the latter two curves provide what we may consider some notion of uncertainty about the average likelihood assigned by the experts, reflecting the width of the likelihood bins. An idea about the dispersion of the estimates is provided by Figure 7. In

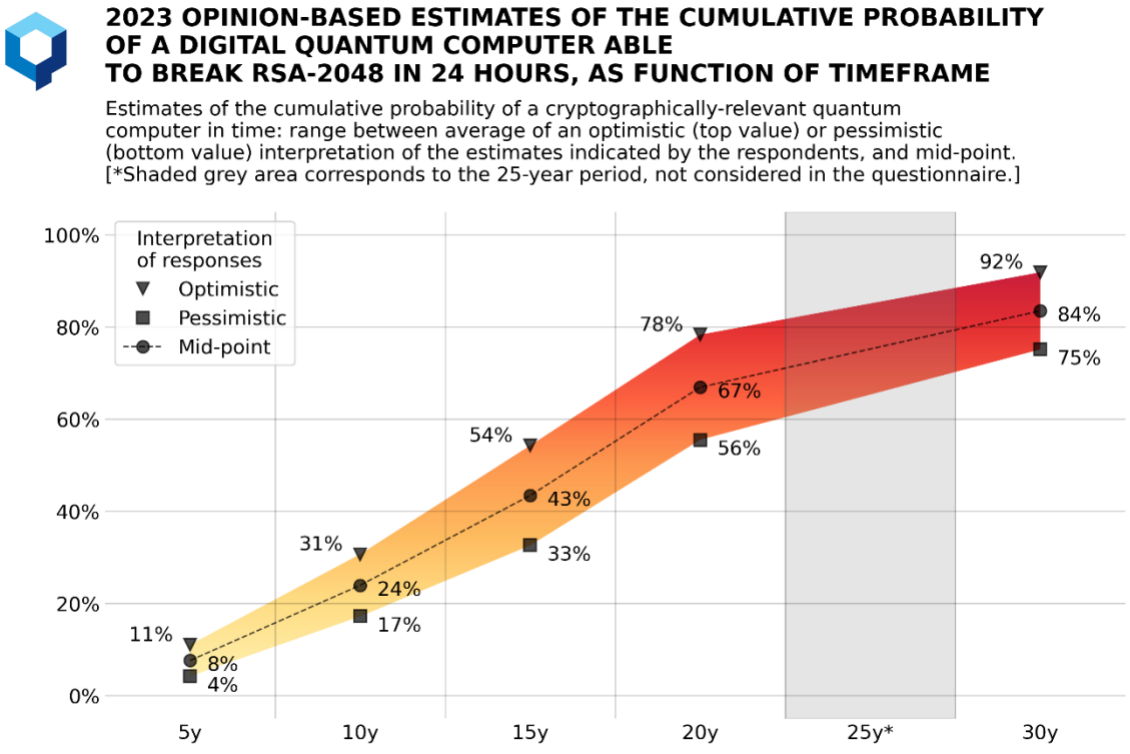

Figure 9 One way of reducing the set of likelihood estimates provided by the experts to some aggregate likelihood is that of interpreting optimistically or, alternatively, pessimistically, the answers of each respondent within the likelihood range they indicated, and averaging over the respondents. Note that, in line with the notion that all likelihood estimates are necessarily vague and imprecise and unable to really differentiate between 5-year intervals far in the future, we did not inquire about expectations for the 25-year timeframe; we introduced a dummy timeframe in the figure to restablish a linear scale on the horizontal temporal axis.

Figure 9, we also show a mid-point estimate, which should not be interpreted as best estimate. More details on the method are given in Appendix A.4 .

Even in a 'pessimistic' interpretation of expert likelihood estimates as the lowest compatible probability for a given likelihood range, the average probability associated to the disruptive quantum threat is already ~33% in the next 15 years.

In general, Figure 9 should be interpreted cautiously as it is a coarse-grained summary of our respondents' opinions but it offers valuable summary information. For example, even in a 'pessimistic' interpretation of responses as the lowest compatible probability for a given likelihood range, the average probability associated by the above-described analysis to the disruptive quantum threat is already ~17% in the next 10 years and growing quickly in the timeframes that follow. Still within a 'pessimistic' interpretation, the average estimated probability is ~33% by the 15-year mark, and ~56% by the 20-year mark. Nonetheless, outliers, particularly for the 5y and 10y timeframes, tend to skew the averages. For example, the heatmap of Figure 7 shows how the distribution at 5 years stands out as largely peaked at "<1%" but with some optimistic outlier estimates.

Comparison with previous years

Our series of surveys allows us to track changes in the likelihood estimates from survey to survey. We think this is useful for at least the two following reasons:

- it provides information on whether the sentiment expressed by the experts is becoming more pessimistic or more optimistic, as their opinions get affected by changing circumstances and recent progress; in turn, a change of sentiment may be interpreted in terms of a likely slowdown or speedup for the development of a quantum computer;
- it provides "differential" information that is conceivably less dependent on the baseline attitude of the pool of experts.

We stress that, while caution is already advisable when interpreting single-survey data, year-to-year comparisons carry additional risks. Among other factors, spurious signals may be introduced by changes in the composition of the pool of respondents, by year-to-year fluctuations in the responses – particularly relevant when dealing with small pools of respondents – and by the relatively wide and unequally spaced likelihood intervals we consider.

In Figure 10, we plot the distribution of the likelihood estimates, for each survey conducted so far – five surveys, from 2019 to 2023. We use distributions rather than the absolute number of respondents so that it is possible to compare surveys with different numbers of respondents. The top graph in Figure 10 considers all respondents for each survey, while the bottom graph is for the set of 18 respondents who have so far taken part in all surveys (see list of respondents in the Appendix).

In Figure 11, we plot the average likelihood intervals for each survey, similarly to what was done in Figure 9 for just the 2023 survey. Top and bottom graphs refer to all respondents and to the stable subset, respectively.

The experts' estimates are relatively consistent from survey to survey, particularly if the analysis is limited to the group of 18 respondents who have taken part in all the surveys so far. Nonetheless, at

face value, one may judge that the growth of the likelihood of a CRQC as a function of timeframe has been sightly pushed further into the future, especially considering the time passed between surveys. The opinions explicitly expressed by our respondents indicate that this may well be a real signal, at least when it comes to the expectations the experts have about future progress. That is, some experts declare that they have grown slightly more pessimistic. Some of the issues responsible for this pessimistic turn are related to the overall socio-political and economic situation (see Section 4.7). We recall that recent global events like the COVID-19 pandemic and the Russian invasion of Ukraine had the respondents indicate a potential slowdown in the surveys preceding this one.

These are comments by two of the core set of respondents:

> *I've become more pessimistic than last year in my near-term estimates. I believe that current technology (especially transmon-qubit [technology]) is going to hit a roadblock in scalability very soon. Successful factoring will require [..] a paradigm shift or revolutionary breakthrough in technological capability. I believe the chance of this happening within 30 years is high. – RESPONDENT*

> *I have given the same approximate probability estimates since 2019, which of course implies that the projected timeline is gradually pushed forward. This is intentional. – RESPONDENT*

Frank Wilhelm-Mauch, a professor at Saarland University and a long-time participant in our surveys, explicitly wrote about being “behind schedule”:

> *Compared to previous years' expectations we are slightly behind schedule as we have not broken even yet - but we understand error models a lot better now and the ball is now in the court of [Fault-Tolerant Quantum Computing] research to deal with these more advanced error models.*

On the other hand, some component of the “pessimistic” signal in the top graphs of Figure 10 and Figure 11 is likely to be spurious and due to changes in the composition of our pool of respondents; this is suggested by the fact that the bottom plots, for the stable subset of respondents, see a smaller “delay”, particularly in the long term.

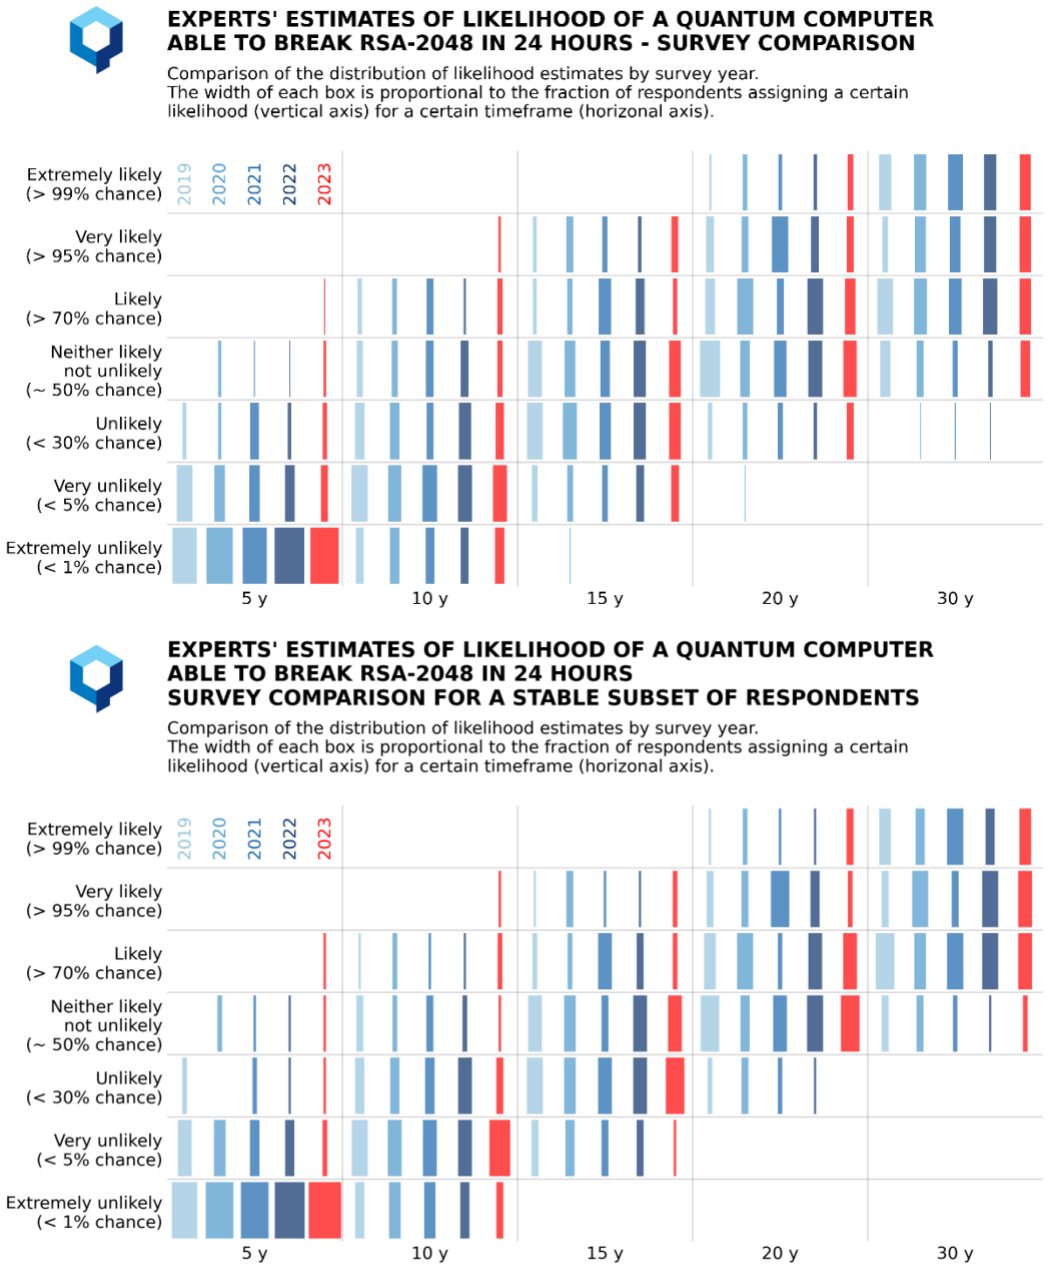

Figure 10 Distribution of the likelihood estimates for each survey conducted so far. Top: likelihood estimates for all the respondents to each survey. Bottom: likelihood estimates for the subset of respondents who took part in all the surveys so far.

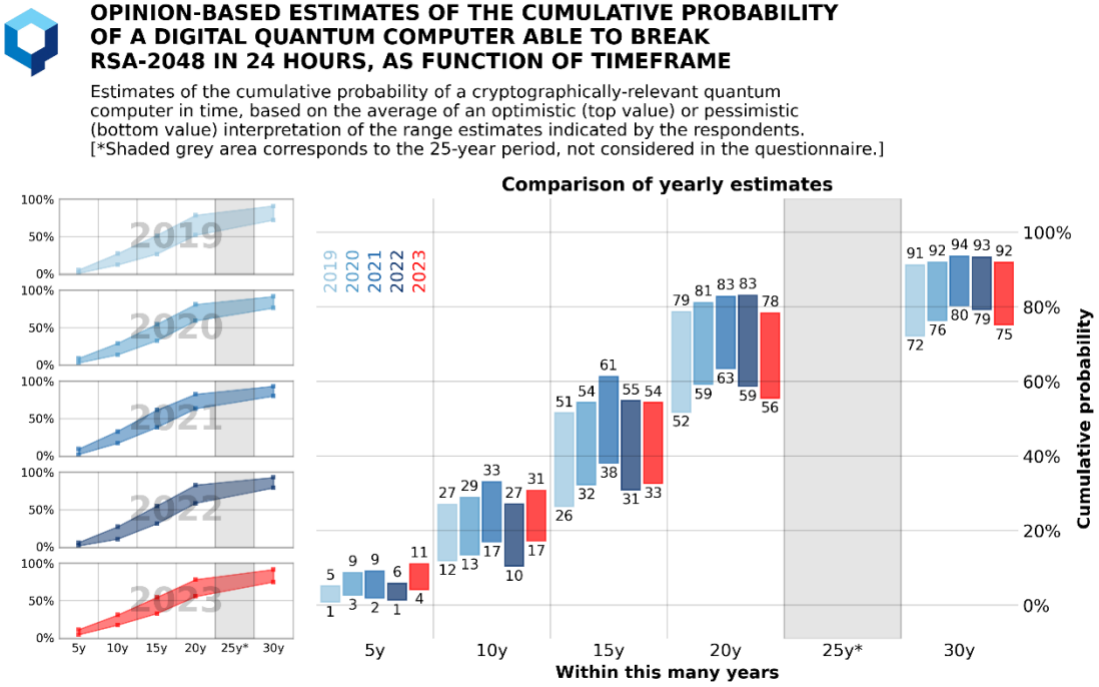

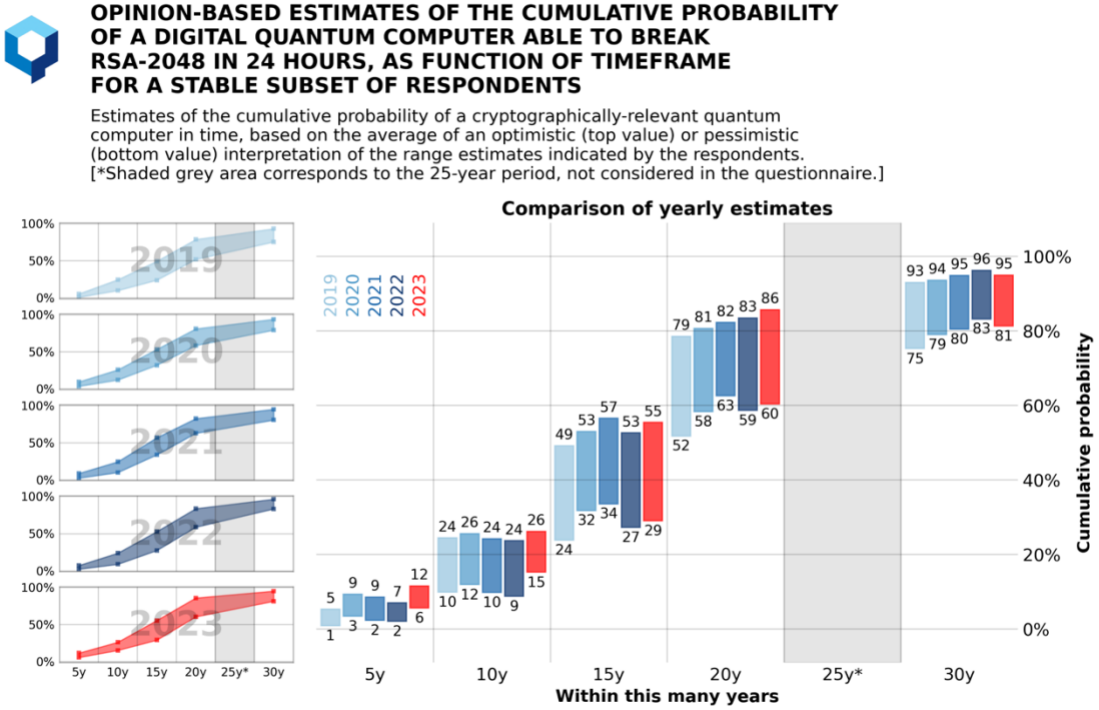

Figure 11 Evolution of the likelihood estimates by the experts in surveys about the quantum threat timeline conducted so far, for all respondents (top) and for a stable subset of respondents (bottom). For both top and bottom graphs: in the subgraphs on the left, probability estimates based on the optimistic or, alternatively, pessimistic interpretation of the responses for the 2019-2023 surveys (see Figure 9 for details for 2023); in the large graph on the right, survey by survey and timeframe by timeframe comparison of such estimates. Note the inclusion of a dummy 25-year timeframe (grey area).

KEY POINTS

- Year after year we have asked our pool of experts to provide their best likelihood estimate for when a quantum computer will be able to perform relatively quickly a specific cryptographically relevant task: breaking RSA-2048.
- The experts display a significant variety of opinions.
- Nonetheless, the estimates expressed suggest a substantial likelihood already in a 10-year timeframe, rapidly growing when moving to 15 and 20 years into the future.
- The opinions of the experts have fluctuated from survey to survey, but they have been and stayed roughly compatible with the above assessment. Such a consistency is stronger when one restricts the analysis to a set of 18 respondents who have taken part to all the five annual surveys conducted so far.

4.3 Potential Concerns

As reported in Section 4.2, 18 of the 37 experts have indicated less than a 95% chance that a CRQC will be built within the next 30 years; of these 18 experts, 8 have estimated a likelihood less than 70%. We would like to better understand the rationale behind such estimates. In general, various reasons why a cryptographically-relevant quantum computer may take 30 years or longer to be built (if ever) have been articulated.

> *"I do not think there is any fundamental physical limit for quantum computing. However, there remain many technical challenges"*
>
> RESPONDENT

We asked the experts to provide their opinion on the level of concern elicited by the following possible issues:

- new-physics phenomena, like a hypothesized random collapse of the wavefunction;
- yet unappreciated standard-physics phenomena that may disrupt quantum computation, like some unavoidable source of correlated noise;
- yet unappreciated fundamental trade-offs in controlling quantum features, that is, something akin to the uncertainty principle;
- excessive technical challenges / requirements not attributable to any of the above, which, despite no being fundamental limitations, would make the scaling to a fault-tolerant quantum computer practically impossible.

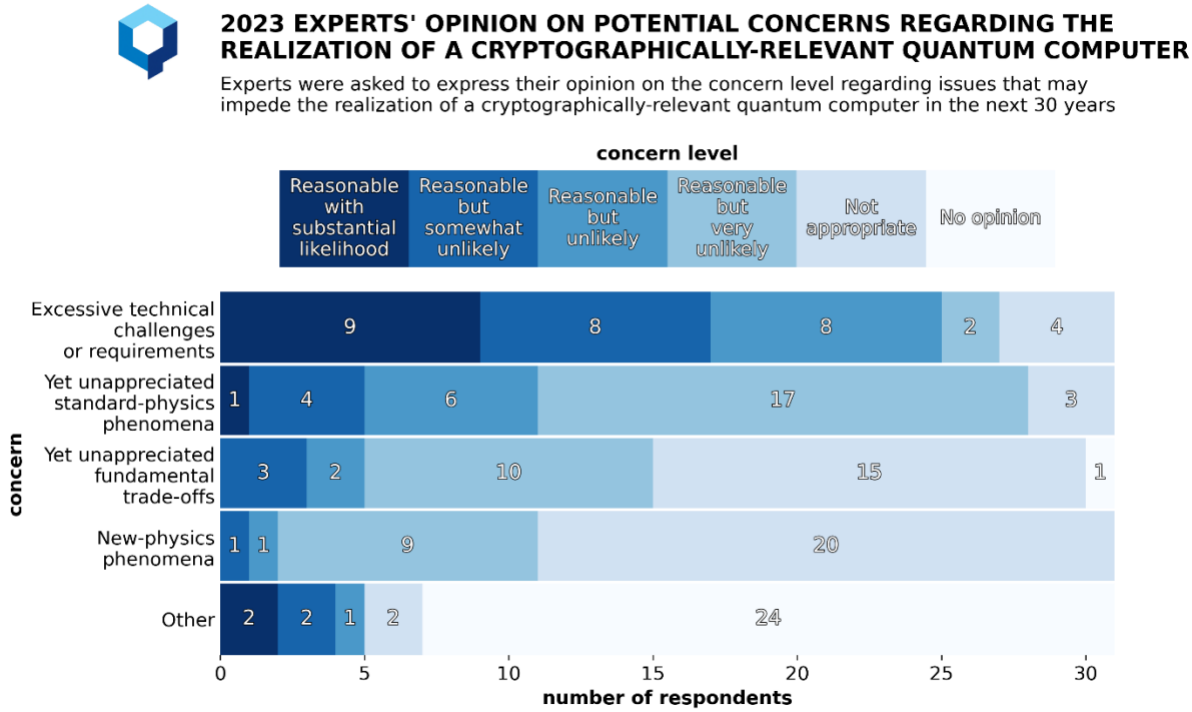

Figure 12 Experts' opinion on a number of concerns that may push the realization of a CRQC further in the future than our longest timeframe of 30 years, or impede it completely. See main text for details.

The level of concern could be classified from a highest “Reasonable concern with substantial likelihood” down to the second-lowest “Reasonable concern but very unlikely” or the lowest “Concern is not appropriate”. The results are reported in Figure 12.

Unsurprisingly, the experts consider the well-known technical challenges quantum researchers face daily as the most reasonable and likely issue that could delay the creation of a CRQC. Far second, but nonetheless judged as a reasonable concern, is the issue of known physical phenomena which may impact quantum computers more negatively that one may currently expect.

Under “Other”, some respondents indicated the possibility that the economic and societal conditions, a shift in interest, or any other dynamics within the research community will be such that not enough resources will be employed to quickly complete the path to a CRQC.

This answer by one of the respondents covers many aspects of the question and summarizes well also comments by several other experts:

> *Quite a number of quantum processor technologies are maturing in parallel, and the likelihood of an unappreciated physical phenomenon or control trade-off affecting all such technologies seems extremely low. Were fault-tolerance thresholds much lower, I would perhaps rate this possibility as significantly more likely. Discovering new physics at the high-entanglement frontier is certainly possible, as we are increasingly probing into an untested physical regime as we explore controlled entanglement of hundreds or thousands of qubits, but I have no reason to suspect such new physics nor if it did that it would be of a nature that would prohibit quantum speedups. My main concern is that the effort to build large scale quantum computers loses momentum due to external factors or dynamics within the research community.*

KEY POINTS

- The experts were asked to provide an opinion on the appropriate level of concern that engineering and scientific issues of various nature could lead to a very slow development of a CRQC, or even ultimately impede it.
- The respondents generally indicate that they do not see any real roadblock. Many perceive it simply as a matter of overcoming scientific and technical hurdles, most likely also via breakthroughs that are by their nature unpredictable but expected to happen.
- A concern some experts share is that funding for the continuous development of quantum computers may stall or diminish for several reasons, including societal and financial ones, thus slowing progress.

4.4 Most important upcoming experimental milestone toward a cryptographically-relevant quantum computer

For those tracking the quantum threat, it would be helpful to have a clear and meaningful milestone between today's current state and a CRQC that convincingly confirms that the major obstacles have been tamed. In order to better understand what such a signal could be, in the present survey we have posed the following question:

Q: What do you consider the most important upcoming experimental milestone to convincingly demonstrate the feasibility of building a cryptographically-relevant fault-tolerant quantum computer?

Most experts would like to see results regarding error correction, the logical encoding and manipulation of quantum information.

> *"The key word here is fault tolerance. Thus scaling the number of logical qubits in current demonstrations of quantum error correction beyond break even is essential."*
>
> ALEXANDRE BLAIS
> Institut quantique, Université de Sherbrooke

> *"The experimental realization of a fully controllable logical qubit prototype, that is interconnectable, and that demonstrates error suppression as the code distance increases, and that in these respects is scalable."*
>
> RESPONDENT

The conditions set by one respondent cover several aspects:

A complete and convincing demonstration of fault-tolerant computation on two logical qubits encompassing a universal set of operations, sustained over many rounds, with results for all gates and coherence times exceeding break-even versus physical qubits.

Daniel Gottesman, a professor at the Computer Science Department of the University of Maryland, hints to the fact that the key demonstration may vary based on the specific physical realization:

A clear and convincing demonstration of a full fault-tolerant protocol with error rates significantly below the unencoded physical error rates. This would apply to any system based on physical qubits. In systems involving encoding a qubit in a continuous-variable mode, it is difficult to make a fair comparison and in these systems, one big challenge so far has been in scaling them up to many modes.

Along the same line, in his response **Nicolas Menicucci** provides platform-dependent targets:

1) superconducting qubits [..] – in this case, the experimental milestones are (a) [low] error rates [..] and (b) quantum interconnection between [multiple] cryostats. [..]

2) photonic qubits [..] – a large-scale cluster state or large-scale fusion-based quantum computing must be demonstrated (> about 100 qubits). To date, we haven't seen any large-scale demonstrations of multi-photon entanglement [..].

3) bosonic qubits [..] – a Gottesman-Kitaev-Preskill state with squeezing above 10 dB. If that can be done, there's a chance for fault tolerance. Until it's done, there will always be doubts. [..]

"In short, the demonstration of horizonal scale via modularity: high universal control fidelity and entanglement fidelity across remote modular quantum processors at a rate much faster than the coherence time of the constituent qubits."

STEPHANIE SIMMONS
Photonic Inc.
& Simon Fraser University,
Co-Chair of Canada's National Quantum Strategy Advisory Council

Scalability is a key aspect and multiple respondents point to demonstrations of modular architectures that would facilitate it. One respondent writes:

A demonstration that one can run realistic quantum algorithms on a distributed quantum computer architecture, using entanglement swapping among multiple physical modules. This is important, because for many experimental systems, there are limits on how many qubits can be trapped/fabricated in a single physical module. This demonstration has to perform entanglement swapping at a high enough rate that it can run quantum algorithms that cannot be partitioned into smaller pieces (i.e., quantum algorithms that cannot be simulated efficiently [..]).

KEY POINTS

- We asked the experts about a clear and meaningful milestone, between today's current state of quantum computing development and a CRQC, that would convincingly confirm that the major obstacles towards a CRQC have been tamed.
- Most experts would like to see strong results regarding error correction and the manipulation of logical qubits, showing that errors and noise can be suppressed sufficiently well and efficiently, thus paving the road to the required scaling of the technology.

4.5 Most promising scheme for fault-tolerance

Fault-tolerance will be reached by combination of improved performance and capabilities of hardware implementation with a suitable error-correction / fault-tolerant scheme. We have asked the experts to share their opinions on the most promising among such schemes.

A straightforward answer is not possible, as by the following words by one respondent:

Quantum error correction is currently a very active research area. As time progresses, it is likely that we will see more advances, [in particular] as systems are scaled up. We may also see adaptations to various hardware architectures, hybrids of error-correction schemes, and so forth.

> *"This depends on the architecture. For most solid-state qubits (superconducting qubits or spins), the surface code is still the most practically relevant scheme since direct short-range coupling is easier than long-range coupling. For optical implementations and possibly for trapped ions with long-range interconnects, it may be feasible to switch to [..] codes with higher thresholds."*
>
> BILL COISH
> McGill University

Correspondingly, the responses of the experts were relatively nuanced, highlighting how the best fault-tolerant scheme depends on the architecture.

They pointed both to leading fault-tolerant schemes, like the surface code—and similar/associated schemes, see Appendix A.2—in superconducting implementations, and to promising fault-tolerant schemes which may improve the rate[6] at which quantum information can be reliably encoded and manipulated. Such improvements would reduce the overall number of physical qubits needed to run the same computation fault-tolerantly, but they may come at the 'cost' of using long-range interactions between physical qubits, which in turn may favour physical systems other than superconducting qubits.

This is the case for so-called quantum Low-Density Parity-Check (LDPC) codes (see Appendix A.4), which have attracted lots of interest in recent times.

Daniel Gottesman writes:

Protocols based on high-rate LDPC codes. These protocols are not yet at the stage where one could produce a really practical protocol based on these codes, but the overhead is significantly lower than for surface codes and the error rates tolerated are not too much worse, even with existing protocols. These types of protocols favor architectures that allow long-rate gates, such as photonic systems, but the possibility of long-range gates is being explored in a wide variety of different systems.

Stephanie Simmons is very optimistic about LDPC codes:

[Quantum]LDPC codes [..] are just so overwhelmingly advantageous on basically all metrics that we should all be designing for their implementation.

[6] Such a rate is the ratio between the number of encoded logical qubits and the number of underlying physical qubits.

Dave Bacon, at Google, is instead cautious, although hopeful:

> *Quantum LDPC type constructions still contain too many challenges for hardware and have not yet been shown to be robust enough to dropout / leakage type errors. But this could change in the next years as researchers work to figure out the practicality of these types of constructions.*

Other proposals for fault-tolerance see the encoding of discrete-variable quantum information (the kind of information supported by a 'standard' qubit) in so-called continuous-variable systems (like the degrees of freedom of a quantized electro-magnetic field) concatenated with discrete-variable error-correction codes. **Nicolas Menicucci** wrote:

> *While I am biased due to my expertise, I still think that bosonic codes will provide the most viable path toward fault tolerance. The ability to engineer states resilient to particular types of noise is powerful, and I expect to see significant innovations in this space in the near term.*

When considering scaling, the importance of modularity and of a distributed approach is stressed by several respondents. For example, **William John Munro**, professor at the Okinawa Institute of Science and Technology Graduate University, writes:

> *At present, it will be a distributed approach potentially using system conducting systems or ion traps. The distributed approach is important for scalability.*

Yvonne Gao, a professor at the Centre for Quantum Technologies of the National University of Singapore, speaks of a combination of error correction and modularity:

> *[M]odular approach to QEC where first layers are realised in hardware efficient manners and then concatenated to the next layer of QEC with each device as interconnected modules.*

Continuing progress in the design of fault-tolerant codes is a strict necessity, according to one respondent:

> *A viable fault-tolerance approach for cryptographically-relevant quantum computer is yet to be developed from a hardware cost perspective.*

KEY POINTS

- Error-correction codes differ in the error rate they can tame, in the rate at which logical qubits can be encoded in physical qubits, and in the practical requirements for their implementations. All the factors affect scalability.
- A relatively new family of error-correction codes, so-called quantum Low-Density Parity-Check (LDPC) codes, have attracted lots of interest in recent times, challenging the prominence of relatively traditional quantum codes, like the so-called surface code.
- Fault-tolerance will be reached by improving the performance and the capabilities of hardware implementations, combined with suitable error-correction schemes.
- Research on quantum error-correction is a very active field and could see significant breakthrough results.

4.6 Useful applications of intermediate quantum processors

Quantum computers with the capability to undermine cybersecurity might take a significant time to develop. The pace of advancements in creating these quantum machines largely hinges on the funding allocated to quantum computing research. This funding can come from research grants, venture capital, or any income generated before achieving a quantum computer with significant cryptographic impact. Although there are various means to stimulate investments and revenues, it is clear that having commercially viable applications would greatly enhance the likelihood of continued investment in quantum technology. Consequently, we sought expert opinions on the matter, by asking the following:

> *"I think it is a good question to ask: We are approaching the point in time where quantum computers will have to begin creating value by delivering solutions to practically relevant problems. This so as to ensure continued investments."*
>
> RESPONDENT

Q: Please indicate your likelihood estimates for useful commercial applications of noisy intermediate-scale quantum (NISQ) processors – or of larger/less noisy processors but anyway not yet cryptographically-relevant – going beyond proof-of-concept and/or promotional activities.

The likelihood estimates by the respondents are presented in Figure 13.

> *"We need end-user industry to get involved constructively to define desired figures of merit and utilities for quantum application to ensure we can demonstrate meaningful advantage sooner."*
>
> ELHAM KASHEFI
> UK National Quantum Computing Centre
> & University of Edinburgh
> & CNRS

While the experts express hope that there will be useful applications of early quantum computing devices – with respondents mentioning a potential useful role in simulations, chemistry, optimization, … – they also point to the many existing caveats and uncertainty, particularly with respect to commercial applications. **Daniel Gottesman** writes:

I think some quantum simulations that will provide useful information are likely within the timeframe but NISQ simulations will not give a convincing degree of confidence in their accuracy. This means that they will not be clearly better than classical approaches that make predictions based on approximations, but will still provide a different window that can help us understand these systems. This is scientifically interesting, but I am not certain how commercially viable it is. Other applications (e.g., machine learning, [Variational Quantum Algorithms]) are possible, but at this point, I view them as merely speculative still until a genuine speedup over classical approaches has been demonstrated.

Dave Bacon is concerned about the damage of hyping the performance and usefulness of early quantum computers, compared to classical computers, and suggests some best practices to avoid such a trap:

As a community we need to be honest in our assessments of current hardware. In particular, it is important that projects "red team" their claims about quantum advantage. A red team is a team that examines the claims and tries to prove them wrong, and for simulating quantum systems there is an extended group of researchers who have a lot of knowledge about this boundary which should be engaged with the red team. We are entering an era where industry has a strong incentive to make strong claims and to obscure these objections. I think this has a direct impact on NISQ being successful since a few high publicity failures could lead to quantum computing having a bad reputation for hyping results.

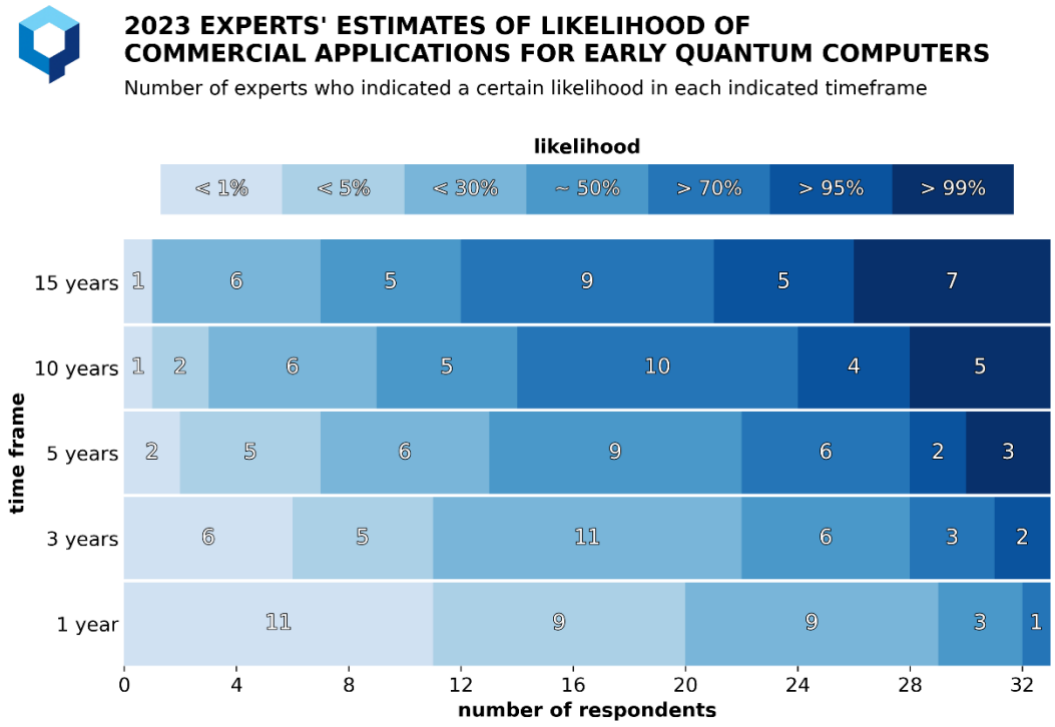

Figure 13 We asked the experts to indicate the likelihood for commercial applications of "early" quantum computers / quantum processors not powerful enough to be directly relevant from a cryptographic perspective. Not all experts expressed an opinion in this sense, but among those who did, more than half indicated a likelihood of about 50% or more within 5 years.

KEY POINTS

- The experts express hope that there will be useful applications of early quantum computing devices, significantly before a CRQC is realized.
- The respondents mention a potential useful role of such devices for tasks like physical simulations, chemistry, and optimization.
- Early applications would provide a significant role in securing further funding directed to the realization of a CRQC.
- Nonetheless, the experts point to many existing caveats and uncertainty when it comes to claims of usefulness of such devices, particularly with respect to commercial applications.
- "Hyped" claims, particularly of early usefulness, come with the risk of slowing quantum computing research long-term.

4.7 Societal and funding factors

This section contains the results for the questions meant to assess how societal and funding factors may impact the timeline of the development of a cryptographically-relevant fault-tolerant quantum computer.

4.7.1 Level of funding of quantum computing research

Substantial and sustained investments are needed to support the development a full fault-tolerant quantum computer. As world leaders in the field, involved in national and international projects and collaborations, working or consulting for industry, and at the head of start-ups, our respondents have a significant vantage point to estimate the evolution of funding. Starting in 2020, we have asked them to forecast what was likely to happen in the following two years[7].The results of the 2023 survey are presented in Figure 14 alongside the previous results. We report the percentage of respondents with a certain opinion.

> *"The general troubles that technology companies are having at the moment seems likely to reduce the amount of money available for longer-term research like quantum computing, but so far I have not seen a significant impact on the quantum industry."*
>
> DANIEL GOTTESMAN
> University of Maryland

For the first time since we started asking this question, the percentage of respondents who think funding will increase or significantly increase are an overall minority at 44%, the same percentage of people who estimate funding will stay about the same.

> *"It looks like global investment in quantum computing by governments will continue a steady rising trend. If one can hit a commercially-relevant application in the next few years, the investment in quantum computing can increase dramatically."*
>
> RESPONDENT

Such a result is in line with the trends suggested by other metrics, like the annual raised start-up investment tracked up to 2022 by McKinsey (Michael Bogobowicz et al. 2013): such a metric had seen a very rapid growth in 2020 and 2021, but increased only by about 1% in 2022, relatively steady at $2.35 billion.

The challenges the world economy and the financial markets are facing in a phase of recovery after the COVID-19 pandemic, particularly with high interest rates, play a role both in the actual dynamics of investments and in the expectations of our respondents about the level of investments. Another factor is the 'conflict' between the hype often surrounding quantum computing and the reality that developing a full-fledged quantum computer is a long-term goal, which also comes with large uncertainties. One respondent writes:

Several things are playing out at the same time: interest rates are still high, inducing investors to seek shorter-term returns; the hype cycle may turn back in the next few years; evidence against the commercial utility of NISQ devices is mounting. Overall, I think this won't sink the field yet, but will hold back further investment increase, at least for a while.

[7] Despite one slight change in wording in the question from the 2020 survey to the 2021 survey, we think the direct comparison of the 2020-2023 responses is reasonable.

It is to be noted that venture capital and institutional funding, which have both contributed to the growth of the field, are expected to exhibit different dynamics, with institutional funding much more stable or even increasing. **Winfried Hensinger**, a professor at the University of Sussex and co-founder of Universal Quantum, points to the essential role of such kind of support:

> *Government contracts will be critical to drive progress in quantum computing. Private investors alone cannot shoulder such long valley of death.*

The concerns expressed about funding have likely influenced the experts' responses regarding the likelihood estimates for a CRQC being realized within a certain future timeframe, analyzed and presented in Section 4.2. The fact that societal and economic factors were mentioned as "Other" concerns in Section 0 further supports this interpretation.

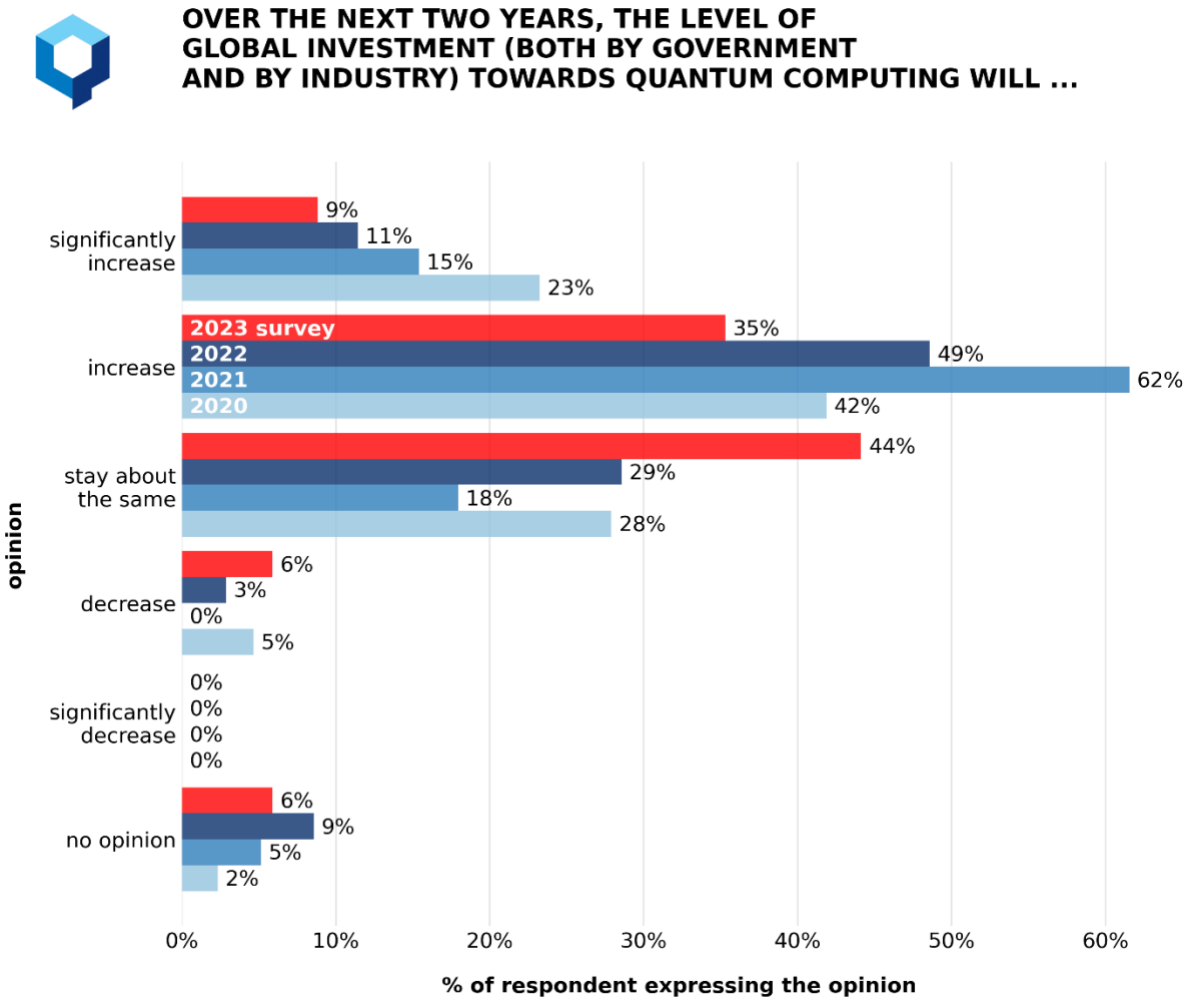

Figure 14 Expected change in the level of investment toward quantum computing in the next two years, comparing estimates by the 2020, 2021, and 2022 respondents. It appears that most experts still expect some increase in funding, but not as significant as in the recent past, and about a third of the 2022 respondents sees funding as staying the same. This is consistent with the past opinions, with levels of investment that are already high, and with the current uncertainty surrounding the global economy.

4.7.2 Global race to build a fault-tolerant quantum computer

The pursuit of a quantum computer with cryptographic significance can be likened to a race on multiple fronts. In Section 4.1, we delved into the "rivalry" among various architectures. In this section, our focus is on the contest involving both national and supranational entities, such as the European Union.

> *"Many of the innovation at the frontier will be driven by private sector in the coming decade, and a good fraction of the action will happen in North America."*
>
> RESPONDENT

Many nations actively recognize the successful creation of a quantum computer as a strategic objective (Kung and Fancy 2021). This is because such a development would revolutionize not just cryptography and much of our digital framework — the main focus of this report — but also various societal and economic sectors. For instance, consider the potential to efficiently emulate quantum systems when creating innovative materials and medicines.

This underlying rivalry significantly propels investments in the quantum computing sector. Consequently, tracking the progress and potential trajectory of this "race" offers valuable insights into the timeline of the quantum threat. Furthermore, for those responsible for addressing the quantum

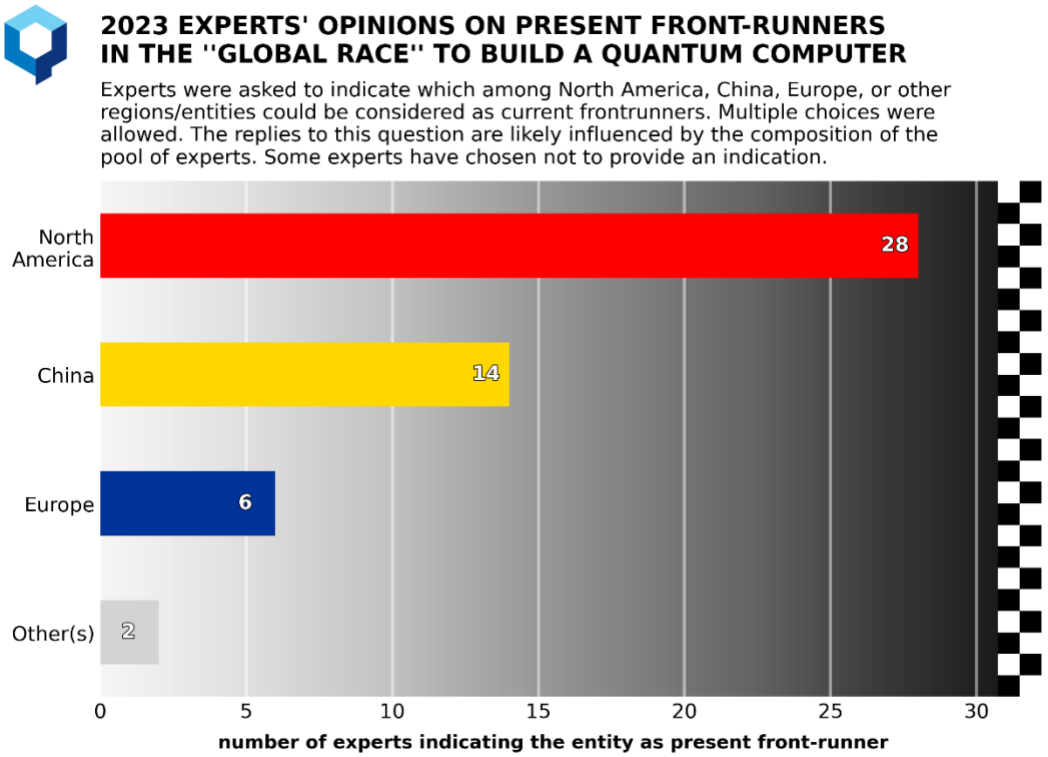

Figure 15 Number of respondents that indicated a region/entity as present front-runner in the global race to build a fault-tolerant quantum computer (multiple answers were allowed). North America appears to be in a strong position, followed by China and then Europe. The "Other(s)" answer reported here was given by a respondent who indicated uncertainty about the status of research in China.

threat, it is essential to discern its potential origin. This entails understanding which entities might first achieve a quantum computer of cryptographic significance.

We solicited expert opinions to identify which regions among China, Europe, and North America are currently leading, allowing for multiple responses and the inclusion of other regions[8].

The results are shown in Figure 15. Not all the experts provided an opinion, with one expert providing the following nuanced motivation, which highlights the importance of a qualified workforce:

> *I think that it is hard to state who is a front-runner, and therefore I have opted not to answer [..]. Recent developments have arguably been driven by US-based companies [..], and one could hence argue that in this sense North America is a front-runner. This being said, the quantum work force that produces the results we are seeing is a very international one. It is all about attracting the right competence and sufficient investments over time.* – RESPONDENT

Frank Wilhelm-Mauch directly criticizes the assumption of the question, as he thinks that "*the notion of "race" is part of the problem*".

According to those who answered with specific choices, North America appears to be the present leading world region, followed by China and Europe, in this order.

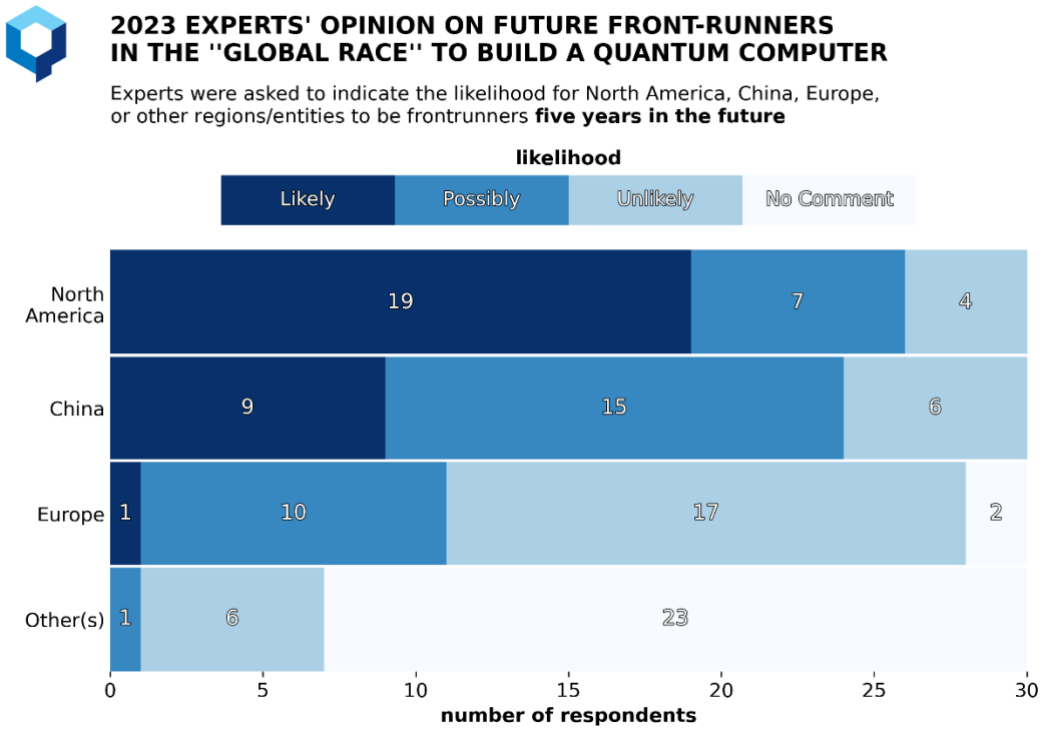

Figure 16 Number of respondents that indicated the likelihood of a given region/entity to be a front-runner in the global race to build a fault-tolerant quantum computer, five years from now. Among the "Others" mentioned: Australia and Japan.

[8] The reader may consider taking into account the geographical composition of our pool of respondents (see Section 3).

Given our interest in future trends, we also asked the experts to indicate the likelihood for each region previously considered to be a frontrunner five years from now, and whether new frontrunners may emerge. The results are presented in Figure 16. Most respondents consider it likely that North America will maintain its frontrunner position. China scores relatively highly as a likely future frontrunner and is considered to have significant potential. Europe appears to lag behind in expectations and many respondents consider it unlikely that it will have the status of frontrunner in five years.

One respondent offered a comment that may provide some rationale for the results:

> *The North American lead will not dissipate within 5 years. China is investing a lot, especially in quantum communications, but the political tensions within and around its government may backfire at some point. Europe is investing a lot, but spreading out the investments very thin.*

Australia and Japan were mentioned as potential leaders by one respondent:

> Some regions would seem to be left out of this list, notably Australia, which I think has a nontrivial – but less than even – chance of being a frontrunner in this race.

According to another respondent, Australia still leads in certain sectors, like silicon qubits, but institutions are still hesitatnt to support R&D and private capital is relatively scarce. About Japan, a different respondent commented:

> *Japan has the technical facilities, skills, and know-how to be a strong contender if they choose to join the race.*

Nonentheless, while the competition is still quite open, **William John Munro** thinks that

> *Other countries like Australia, Japan etc will take time to catch up.*

KEY POINTS

- The journey towards realizing a quantum computer is often termed the ‘quantum race’. Competition exists both at the level of nations as well as of private companies.
- The last few years have seen a surge of investments in quantum technologies, with $2.35 billion raised by quantum start-ups in 2022 alone, according to McKinsey. The experts expect that investments may not continue to grow as fast, also because of a more complex economic and financial global situation.
- We solicited expert opinions to identify which regions among China, Europe, and North America are currently leading the ‘race’, allowing for the inclusion of other regions. We further asked which regions are most likely to be leading five years from now.
- North America appears to be the present leading world region, followed by China and Europe, in this order. Most respondents consider it likely that North America will maintain its frontrunner position. China scores relatively highly as a likely future frontrunner. Europe appears to lag as future expectations go.
- Other countries like Australia and Japan are also considered to have significant potential.

4.8 Sources of unexpected speed-up

Breakthroughs might speed-up substantially and relatively unexpectedly the development of a CRQC. The field of quantum computing research is composed of several subfields, and it is of interest to understand where the largest potential for breakthroughs sits. This is also helpful in terms of monitoring progress.

To gain insight, we asked the respondents their opinion about some aspects of quantum computing research as sources of substantial and potentially unexpected progress. The results illustrated in Figure 17 indicate that while many aspects of quantum computing research could be the source of breakthroughs, the experts see hardware development – with the reduction of error rates and increased capabilities – as well as quantum error correction as the most likely ones. One respondent wrote:

I believe that the combination of error mitigation and error correction has the potential to create a breakthrough in the ability to reduce the effect of errors in quantum algorithms, making the optimal use of existing sizes of quantum computers.

While compilation ranks last, **Bill Coish** begs to differ**:**

From my perspective, hardware-aware efficient compilation schemes are where there is the most potential gain -- there are many sources of error and types of error that are highly specific to individual hardware platforms [..].

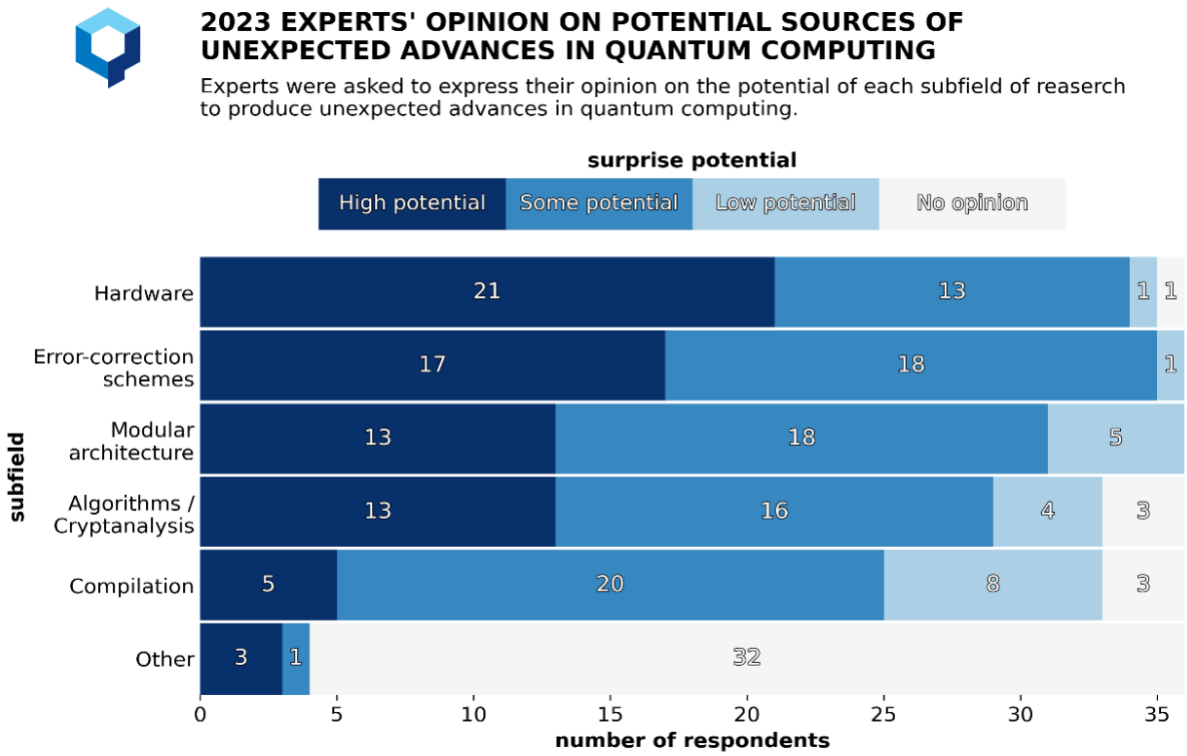

Figure 17 Number of respondents that indicated a certain "surprise potential" for several subfields of quantum computing research, which could also be seen as "layers" of a quantum stack.

KEY POINTS

- Breakthroughs might speed up substantially and relatively unexpectedly the development of a CRQC.
- The field of quantum computing research is composed of several subfields, and it is of interest to understand where the largest potential for breakthroughs sits. This is also helpful in terms of monitoring progress.
- The experts see hardware development and quantum error correction as the most likely sources of unexpected advances.

4.9 Current progress

In this section we present opinions about the status of progress in quantum computing research and development.

4.9.1 Recent developments

We asked the respondents to indicate what they considered to have been the most important advances in the field in the past year. Opinions varied but these results were mentioned repeatedly:

- progress in the demonstration of error correcting codes/steps towards fault tolerance, with more than one respondent pointing to (Acharya et al. 2023), which demonstrated break even scaling to a distance-5 surface code, and to (Sivak et al. 2023), which demonstrated real-time quantum error correction beyond break-even;
- the increase in size and/or performance of quantum processors – in terms of sheer number of physical qubits or other figures of merit;
- progress in researching quantum Low-Density Parity-Check (LDPC) codes (Breuckmann and Eberhardt 2021), which has led also to experimental implementations (Bravyi et al. 2023; Xu et al. 2023);
- advances in bosonic quantum error correction with Gottesman-Kitaev-Preskill codes (Brady et al. 2023).

> *"There had been several papers claiming that various aspects of break even point had either been reached or close to that. Reaching the break-even point is a crucial moment in the history of quantum computation."*
>
> RESPONDENT

4.9.2 Next near-term step

> *"Demonstration of memory time scaling exponentially with size of code block."*
>
> RESPONDENT

We asked our respondents to indicate a significant result on the path towards fault-tolerant quantum computation that they see as both necessary and achievable within approximately one year.

Unsurprisingly, the experts mentioned progress needed along the same lines as already considered in this report, for example, improvements in error rates, better and more convincing demonstration of quantum error correction and fault-tolerance, development of modular and hybrid architectures.

Simon Benjamin clarifies the kind of logical encoding that he would like to see and that he thinks could be reasonably achieved soon:

> *More progress on logical qubits, regardless of the code. Although we repeatedly hear that logical qubits are achieved, they aren't really — we would want to see dramatically extended lifetimes versus physical qubits, AND proper lifetime extension with code size, AND both Clifford and non-Clifford operations. Could be in the next year, if in one of the leading platforms (SC, ion trap). Similarly, more modest logical qubit progress in the less-mature systems (photonics, silicon spin, neutral atom).*

KEY POINTS

- The experts highlighted some of the most important progress in the last year. Some results were mentioned repeatedly:
  - progress in the experimental demonstration of error correcting codes/steps towards fault tolerance,
  - the increase in size and/or performance of quantum processors,
  - advances in error correction, particularly progress on quantum LDPC codes and on so-called bosonic codes.
- With respect to results on the path towards fault-tolerant quantum computation that they see as both necessary and achievable within approximately one year, the respondents pointed to:
  - improvements in error rates,
  - better and more convincing demonstration of quantum error correction and fault-tolerance,
  - development of modular and hybrid architectures.

4.10 Other notable remarks by participants

We asked the respondents to "comment freely on the present and near-future status of development of quantum computers". This section contains a selection of such comments and of other notable remarks not already quoted in prior sections of this report.

There is still a very large risk as to whether post quantum cryptography is actually safe versus quantum attacks. The vast majority of researchers in this area have little to no quantum experience, and the number of people working on novel quantum algorithms for these is much smaller than it should be, given the economic impact of the security of these systems. – DAVE BACON

A serious effort on hardware may pay off significantly at this stage. Modular architectures may be designed in closer match to error correction and error mitigation strategies and algorithms and strategies may pass from resilience against general errors to adaption to actual error models. – KLAUS MOELMER

Any prediction beyond 5 years is very difficult to make and answers will depend in part on factors such as funding and recruitment of the right people to master key challenges. The later relies on training sufficiently many researchers and engineers with the correct skills. – WINFRIED HENSINGER

In spite of the uncertainty about the future of quantum computing, I think that there's enough optimism, that investors will want to be "up-to-speed" in case the "killer-apps" materialize. – RESPONDENT

I think the community slightly undervalues actually scaling experiments: just doing a demo at fixed sized is great, but we really need to be showing that "things get better" as we get larger. No hardware is yet at that point. – DAVE BACON

As some of research becomes increasingly nationalistic, this could seriously hurt the exchange of ideas and talents, which in turn would inevitably slow down the overall progress globally. – YVONNE GAO

*At some point in the future, there could be a collective decision by the scientific community (including interested parties in industry and government) that we *do* want to build a large-scale fault-tolerant quantum computer. If/when/how that happens could impact the quantum threat timeline. – RESPONDENT*

Summary and outlook

A fully functional quantum computer is a threat for cryptosystems based on certain computational problems that are thought to be impossibly hard for present computational devices. Such problems could be easily handled by a sufficiently large and reliable quantum computer – what we call a Cryptographically-Relevant Quantum Computer (CRQC) – executing the right quantum cryptanalysis algorithms.

Achieving a CRQC will necessitate years of advancements in both science and engineering, which can only be attained through dedicated commitment and ample resources. The key challenge to overcome is the natural 'fragility' of the quantum features that make quantum computing more powerful than classical computing.

The journey towards realizing a quantum computer is often termed the 'quantum race', with competition at the level of nations as well as of private companies. This race has intensified recently, marked by the participation of big corporate entities, substantial government funding, and a surge of start-ups backed by venture capital. Nonetheless, it is more apt to liken this to a marathon than a sprint, given the prolonged research and investment required.

Cyber-risk managers may want to track developments in the experimental realization of quantum error correction to understand how quickly quantum computers are becoming a reality.

On the theory side, better error correction schemes and improvements in quantum cryptanalysis algorithms may well enable cryptanalysis with fewer quantum resources than seemingly required today, shortening the time to the concretization of the quantum threat.

That said, unexpected leaps forward are possible, owing to breakthroughs in science and/or engineering. The end goal is computations using logical qubits, a dependable way to encode and handle quantum information even when the underlying physical qubits are error-prone. We are venturing into times where increasingly credible and effective examples of such encoding and handling are achieved. Those in cybersecurity should monitor these progressions to gauge the speed at which quantum computers are materializing. In addition, one must consider the possibility of advancements in cryptanalysis algorithms, which would enable cryptanalysis with fewer quantum resources – say, fewer quantum qubits, or fewer computational steps – than the current state of the art.

In general, the expert opinions we have collected and summarized in this report – and in the series of reports it is part of – offer unique insight into the quantum threat timeline. Thirty-seven experts estimated the likelihood of the realization of a quantum computer that could break a scheme like RSA-2048 in 24 hours, and such opinions indicate a substantial likelihood within a 10-year timeframe: more than a quarter of the respondents (10/37) felt it was "about 50%" or ">70%" likely. The risk aversion/appetite of companies and institutions can vary significantly, but for critical systems, such estimated likelihoods represent a serious concern.

Importantly, the perceived imminence of the quantum threat is dynamic and can shift from survey to survey. Variables such as recent discoveries, investment fluctuations, and the economic and financial

landscape can impact both the actual threat timeline and the assessments of our experts. Our ongoing series of reports offers a lens to track these variations, but it is essential to also consider potential confounders like the changes in the composition of the pool of respondents.

The experts' estimates have been relatively consistent from survey to survey, particularly if the analysis is limited to the group of 18 respondents who have taken part in all the five surveys conducted so far. The likelihood of a CRQC is already far from negligible at 10y, but it appears to increase substantially in between the 15y and 20y marks, making a CRQC more likely than not within that kind of timeframe. Current progress in error mitigation and in error correction, the increase in the number of physical qubits available on various platforms, as well as new results in the development of efficient error-correction schemes, all fuel positive expectations for the next steps in quantum computing development.

Which physical platform will be the winner in the quantum race is not yet clear, and there will not necessarily be only one such winner. Presently, according to the experts' opinions, superconducting circuits and ion traps seem to have an edge over the competition. Other platforms continue to be developed, and some, such as integrated optics and neutral atoms, have attracted increased attention in the last couple of years, to the extent that some of the experts consider them as leading candidates. There is also the potential of combining different technologies, both to take advantage of the specific strengths each of them may have, or to create modular systems that may facilitate scaling up the number of physical and logical qubits.

The logical possibility that consequential quantum cryptanalysis is infeasible or impossible is captured in the small but non-negligible likelihood implicitly assigned in our survey to the possibility that quantumly breaking RSA-2048 will take more than 30 years. When directly queried about what could prevent the realization of a CRQC within 30 years, the respondents generally indicate that they do not see any real roadblock. Many perceive it simply as a matter of overcoming scientific and technical hurdles, most likely also via breakthroughs that are unpredictable but expected to happen, as it has occurred often in the history of technology. A concern some experts share is nonetheless that funding for the continuous development of quantum computers may stall or diminish, slowing progress down. There is the hope that quantum computing devices that are not yet cryptographically relevant will be proven to be useful enough to stimulate continued investments.

While it is up to each institution, company, and manager to decide what risk they are ready to accept, we think cyber-risk managers are naturally more concerned about the chance that the quantum threat materializes early — and potentially earlier than many could expect — rather than never. Progress in the last few years—particularly the demonstration of several aspects of quantum error correction—

"It is important to stress — not least given the roadmaps presented by industry — the importance of migrating to post-quantum secure cryptography. In particular, this is important in applications where long-term confidentiality is sought. This is because adversaries can store ciphertexts that are intercepted now for decryption sometime in the future when large-scale fault-tolerant quantum computers become available."

RESPONDENT

together with the significant momentum of the field—in terms of activities, results, and resources poured into it—should trigger caution, directed to pro-actively developing crypto-agility and resilience against quantum attacks ("Quantum Readiness Toolkit: Building a Quantum-Secure Economy" 2023; Quantum-Readiness Working Group 2023). This is particularly important for three reasons.

First, one should consider that malicious agents may adopt already now a "Harvest Now, Decrypt Later" (HNDL) approach, storing valuable encrypted data while waiting for a CRQC to become a reality. This means that the data and communication of today are already potentially at risk.

Second, there might be progress that is not public or publicized. As Stephanie Simmons, Co-Chair of Canada's National Quantum Strategy Advisory Council warns:

> *Not all progress is visible to the world's academic community anymore, and we should expect this trend to increase in the coming years. Unexpected progress will be, and is, hidden from full view.*

Third, not preparing now against the quantum threat sets the conditions where a hasty transition to quantum-safe tools may suddenly become a forced choice, with all the risks associated to it, from a breakdown of services to involuntarily creating vulnerabilities even against more traditional attacks.

The Global Risk Institute and evolutionQ Inc. have already made available a quantum risk assessment methodology for going from estimates of the threat timeline – like those provided in this series of reports – to the evaluation of the overall urgency of an institution or a company for taking action to ensure quantum safety (Mosca and Mullholland 2017).

References

Acharya, Rajeev, Igor Aleiner, Richard Allen, Trond I. Andersen, Markus Ansmann, Frank Arute, Kunal Arya, et al. 2023. "Suppressing Quantum Errors by Scaling a Surface Code Logical Qubit." *Nature* 614 (7949): 676–81. https://doi.org/10.1038/s41586-022-05434-1.

Arute, Frank, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, et al. 2019. "Quantum Supremacy Using a Programmable Superconducting Processor." *Nature* 574 (7779): 505–10. https://doi.org/10.1038/s41586-019-1666-5.

Bombin, H., and M. A. Martin-Delgado. 2006. "Topological Quantum Distillation." *Physical Review Letters* 97 (18): 180501. https://doi.org/10.1103/PhysRevLett.97.180501.

Brady, Anthony J., Alec Eickbusch, Shraddha Singh, Jing Wu, and Quntao Zhuang. 2023. "Advances in Bosonic Quantum Error Correction with Gottesman-Kitaev-Preskill Codes: Theory, Engineering and Applications." arXiv. https://doi.org/10.48550/arXiv.2308.02913.

Bravyi, Sergey, Andrew W. Cross, Jay M. Gambetta, Dmitri Maslov, Patrick Rall, and Theodore J. Yoder. 2023. "High-Threshold and Low-Overhead Fault-Tolerant Quantum Memory." arXiv. https://doi.org/10.48550/arXiv.2308.07915.

Breuckmann, Nikolas P., and Jens Niklas Eberhardt. 2021. "Quantum Low-Density Parity-Check Codes." *PRX Quantum* 2 (4): 040101. https://doi.org/10.1103/PRXQuantum.2.040101.

DiVincenzo, David P. 2000. "The Physical Implementation of Quantum Computation." *Fortschritte Der Physik* 48 (9–11): 771–83. https://doi.org/10.1002/1521-3978(200009)48:9/11<771::AID-PROP771>3.0.CO;2-E.

Feynman, Richard P. 1982. "Simulating Physics with Computers." *International Journal of Theoretical Physics* 21 (6): 467–88. https://doi.org/10.1007/BF02650179.

Fowler, Austin G., Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. 2012. "Surface Codes: Towards Practical Large-Scale Quantum Computation." *Physical Review A* 86 (3): 032324. https://doi.org/10.1103/PhysRevA.86.032324.

Gheorghiu, Vlad, and Michele Mosca. 2017. "GRI Quantum Risk Assessment Report - Part 1." Global Risk Institute. 2017. https://globalriskinstitute.org/publications/resource-estimation-framework-quantum-attacks-cryptographic-functions/.

———. 2021. "A Resource Estimation Framework For Quantum Attacks Against Cryptographic Functions: Recent Developments."

Gidney, Craig, and Martin Ekerå. 2021. "How to Factor 2048 Bit RSA Integers in 8 Hours Using 20 Million Noisy Qubits." *Quantum* 5 (April): 433. https://doi.org/10.22331/q-2021-04-15-433.

Grover, Lov K. 1996. "A Fast Quantum Mechanical Algorithm for Database Search." In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, 212–19. STOC '96. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/237814.237866.

Horsman, Clare, Austin G. Fowler, Simon Devitt, and Rodney Van Meter. 2012. "Surface Code Quantum Computing by Lattice Surgery." *New Journal of Physics* 14 (12): 123011. https://doi.org/10.1088/1367-2630/14/12/123011.

Kitaev, A. Yu. 2003. "Fault-Tolerant Quantum Computation by Anyons." *Annals of Physics* 303 (1): 2–30. https://doi.org/10.1016/S0003-4916(02)00018-0.

Kung, Johnny, and Muriam Fancy. 2021. "A Quantum Revolution: Report on Global Policies for Quantum Technology." CIFAR. April 7, 2021. https://cifar.ca/wp-content/uploads/2021/05/QuantumReport-EN-May2021.pdf.

Michael Bogobowicz, Scarlett Gao, Mateusz Masiowski, Niko Mohr, Henning Soller, Rodney Zemmel, and Matija Zesko. 2013. "Quantum Technology Monitor." 2013.

https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/quantum-technology-sees-record-investments-progress-on-talent-gap.

Mosca, Michele. 2013. *E-Proceedings of 1st ETSI Quantum-Safe Cryptography Workshop*.

Mosca, Michele, and John Mullholland. 2017. "A Methodology for Quantum Risk Assessment." Global Risk Institute. 2017. https://globalriskinstitute.org/publications/3423-2/.

Mosca, Michele, and Marco Piani. 2019. "Quantum Threat Timeline." Global Risk Institute. 2019. https://globalriskinstitute.org/publications/quantum-threat-timeline/.

———. 2021. "Quantum Threat Timeline Report 2020." Global Risk Institute. 2021. https://globalriskinstitute.org/publications/quantum-threat-timeline-report-2020/.

Nielsen, Michael A., and Isaac L. Chuang. 2000. *Quantum Computation and Quantum Information*. Cambridge University Press.

NIST. 2016. "Announcing Request for Nominations for Public-Key Post-Quantum Cryptographic Algorithms." Federal Register. December 20, 2016. https://www.federalregister.gov/documents/2016/12/20/2016-30615/announcing-request-for-nominations-for-public-key-post-quantum-cryptographic-algorithms.

———. 2023. "Post-Quantum Cryptography Standardization - Post-Quantum Cryptography | CSRC | CSRC." CSRC | NIST. 2023. https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization.

Preskill, John. 2018. "Quantum Computing in the NISQ Era and Beyond." *Quantum* 2 (August): 79. https://doi.org/10.22331/q-2018-08-06-79.

"Quantum Readiness Toolkit: Building a Quantum-Secure Economy." 2023. World Economic Forum. https://www.weforum.org/publications/quantum-readiness-toolkit-building-a-quantum-secure-economy/.

Quantum-Readiness Working Group. 2023. "Canadian National Quantum-Readiness: Best Practices and Guidelines." Canadian Forum for Digital Infrastructure Resilience (CFDIR). https://ised-isde.canada.ca/site/spectrum-management-telecommunications/sites/default/files/attachments/2023/cfdir-quantum-readiness-best-practices-v03.pdf.

Rivest, R. L., A. Shamir, and L. Adleman. 1978. "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems." *Communications of the ACM* 21 (2): 120–26. https://doi.org/10.1145/359340.359342.

Sevilla, Jaime, and C. Jess Riedel. 2020. "Forecasting Timelines of Quantum Computing." *arXiv:2009.05045 [Quant-Ph]*, September. http://arxiv.org/abs/2009.05045.

Shor, P.W. 1994. "Algorithms for Quantum Computation: Discrete Logarithms and Factoring." In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 124–34. https://doi.org/10.1109/SFCS.1994.365700.

Sivak, V. V., A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, et al. 2023. "Real-Time Quantum Error Correction beyond Break-Even." *Nature* 616 (7955): 50–55. https://doi.org/10.1038/s41586-023-05782-6.

Xu, Qian, J. Pablo Bonilla Ataides, Christopher A. Pattison, Nithin Raveendran, Dolev Bluvstein, Jonathan Wurtz, Bane Vasic, Mikhail D. Lukin, Liang Jiang, and Hengyun Zhou. 2023. "Constant-Overhead Fault-Tolerant Quantum Computation with Reconfigurable Atom Arrays." arXiv. https://doi.org/10.48550/arXiv.2308.08648.

A. Appendix

In this Appendix, we provide more detailed information about various aspects of the reports, from a complete list of the respondents, to background information about quantum computing, to aspects of our methodology.

A.1 List of respondents

A short description/bio that emphasizes the rationale for the inclusion of each respondent is provided after the table. Respondents who have participated in all the surveys from 2019 to 2023 are listed at the beginning and highlighted in grey.

#	Name	Institution	Country
1	Dorit Aharonov	Hebrew University of Jerusalem and QEDMA Quantum Computing	ISR
2	Dave Bacon	Google Quantum AI	USA
3	Simon Benjamin	University of Oxford	GBR
4	Alexandre Blais	Institut quantique, Université de Sherbrooke	CAN
5	Ignacio Cirac	Max Planck Institute of Quantum Optics	GER
6	Bill Coish	McGill University	CAN
7	David DiVincenzo	Jülich Research Center	GER
8	Runyao Duan	Baidu Quantum Computing Institute	CHN
9	Martin Ekerå	KTH Royal Institute of Technology and Swedish NCSA	SWE
10	Artur Ekert	University of Oxford	GBR/SGP
11	Daniel Gottesman	University of Maryland	USA
12	Jungsang Kim	IonQ Inc. and Duke University	USA
13	Andrea Morello	UNSW Sydney	AUS
14	Yasunobu Nakamura	University of Tokyo	JPN
15	Peter Shor	Massachusetts Institute of Technology	USA
16	Stephanie Simmons	Simon Fraser University and Photonic Inc	CAN
17	Frank Wilhelm-Mauch	Jülich Research Center Saarland University	GER
18	Shengyu Zhang	Tencent Quantum Lab	CHN
19	Sergio Boixo	Google	USA
20	Andrew Childs	University of Maryland Joint Center for Quantum Information and Computer Science	USA
21	Joe Fitzsimons	Horizon Quantum Computing	SGP
22	Jay Gambetta	IBM	USA
23	Yvonne Gao	Centre for Quantum Technologies, National University of Singapore	SGP
24	Winfried Hensinger	University of Sussex Universal Quantum	GBR

25	Elham Kashefi	UK National Quantum Computing Centre, School of Informatics, University of Edinburgh & CNRS, LIP6, Sorbonne University	GBR/FRA
26	Sir Peter Knight	Imperial College London	GBR
27	Yi-Kai Liu	US National Institute of Standards and Technology (NIST)	USA
28	Klaus Moelmer	Niels Bohr Institute, University of Copenhagen	DNK
29	William John Munro	Okinawa Institute of Science and Technology	JPN
30	Nicolas Menicucci	RMIT University	AUS
31	Kae Nemoto	Okinawa Institute of Science and Technology	JPN
32	Francesco Petruccione	Stellenbosch University	ZAF
33	John Preskill	California Institute of Technology	USA
34	Simone Severini	Amazon	USA
35	Lieven Vandersypen	QuTech, TU Delft	NLD
36	Gregor Weihs	University of Innsbruck	AUT
37	David J. Wineland	University of Oregon	USA

Dorit Aharonov
A leader in quantum algorithms and complexity, and co-inventor of the quantum fault-tolerance threshold theorem.

Dave Bacon
Leads the quantum software team at Google, facilitating the exploitation of noisy intermediate-scale quantum devices, and is an expert on the theory of quantum computation and quantum error correction.

Simon Benjamin
Simon Benjamin is an international expert in the theoretical and computational studies supporting the implementation of realistic quantum devices. He is co-founder of the company Quantum Motion and professor of quantum technologies at Oxford.

Alexandre Blais
A leader in understanding how to control the quantum states of mesoscopic devices and applying the theoretical tools of quantum optics to mesoscopic systems, he has provided key theoretical contributions to the development of the field of circuit quantum electrodynamics with superconducting qubits.

Sergio Boixo
He is the Chief Scientist for Quantum Computer Theory at Google's Quantum Artificial Intelligence Lab. He is known for his work on quantum neural networks, quantum metrology and was involved with the first ever demonstration of quantum supremacy.

Andrew Childs
Interested in the power of quantum systems to process information, he is a leader in the study and development of quantum algorithms. He is co-director of the Joint Center for Quantum Information and

Computer Science (QuICS), and director of the NSF Quantum Leap Challenge Institute for Robust Quantum Simulation.

Ignacio Cirac
One of the pioneers of the field of quantum computing and quantum information theory. He established the theory at the basis of trapped-ion quantum computation. He devised new methods to efficiently study quantum systems with classical computers, and to use controllable quantum systems (like cold atoms) as quantum simulators.

Bill Coish
A theoretician working closely with experimentalists, he is a leading expert on solid-state quantum computing, including both spin-based and superconducting implementations.

David DiVincenzo
A pioneer in the field of quantum computing and quantum information theory. He formulated the "DiVincenzo criteria" that an effective physical implementation of quantum computing should satisfy.

Runyao Duan
An expert in quantum information theory, he is the Director of the Quantum Computing Institute of Baidu. He was the Founding Director of Centre for Quantum Software and Information at University of Technology Sydney.

Martin Ekerå
A leading cryptography researcher focusing on quantum computing algorithms for cryptanalysis, and on the development of post-quantum secure classical cryptographic schemes. He is the co-author of one of the most recent and influential estimates of the resources required by a realistic and imperfect quantum computer to break the RSA public-key encryption scheme.

Artur Ekert
A pioneer in the field of quantum information who works in quantum computation and communication. He invented entanglement-based quantum key distribution and was the founding director of the Centre for Quantum Technologies of Singapore.

Jay Gambetta
He is an IBM Fellow and VP of IBM Quantum. He leads the team at IBM Thomas J Watson Research Center working to build a quantum computer.

Yvonne Gao
Leads a group to develop modular quantum devices with superconducting quantum circuits. In 2019, she was named one of the Innovators Under 35 (Asia Pacific) by MIT Tech Review for her work in developing crucial building blocks for quantum computers.

Daniel Gottesman
A pioneer of quantum error correction, and inventor of the stabilizer formalism for quantum error correction.

Winfried Hensinger
He heads the Sussex Ion Quantum Technology Group and is the director of the Sussex Centre for Quantum Technologies. He is a co-founder, Chief Scientist and Chairman of Universal Quantum, a full-stack quantum computing company.

Elham Kashefi
A leading quantum cryptography researcher, renowned for her work on blind quantum computing. She is a professor at the University of Edinburgh, a CNRS researcher at the Sorbonne University, and Chief Scientist at UK's National Quantum Computing Centre.

Jungsang Kim
An experimentalist leading the way towards a functional integration of quantum information processing systems comprising, e.g., micro-fabricated ion-trap and optical micro-electromechanical systems. He is also cofounder and chief strategy officer of IonQ Inc., a company focusing on trapped-ion quantum computing.

Sir Peter Knight
He is a pioneer in the field of quantum optics and quantum information. He has served as a fellow of the Royal Society, President of the Optical Society of America and Chief Scientific Advisor at the UK National Physical Laboratory.

Yi-Kai Liu
He is a leader in research on quantum computation, quantum algorithms and complexity, quantum state tomography and cryptography. He is the Co-Director of the Joint Center for Quantum Information and Computer Science, an Adjunct Associate Professor in the University of Maryland, and a staff scientist in the Applied and Computational Mathematics Division at the National Institutes of Standards and Technology (NIST)

Nicolas Menicucci
A leading researcher who contributed key results in the development of continuous-variable cluster states, and who further focuses on foundational quantum information and quantum theory, in particular in relation to relativity.

Klaus Moelmer
A pioneering physicist at the University of Aarhus, he has made outstanding and insightful contributions to theoretical quantum optics, quantum information science and quantum atom optics, including the development of novel computational methods to treat open systems in quantum mechanics and theoretical proposals for the quantum logic gates with trapped ions.

Andrea Morello
A leading experimentalist in the control of dynamics of spins in nanostructures. Prof Morello's group was the first in the world to achieve single-shot readout of an electron spin in silicon, and the coherent control of both the electron and the nuclear spin of a single donor.

William John Munro
A professor at the Okinawa Institute of Science and Technology Graduate University. Previously, he was

a leader in HP's development of quantum enabled technologies and headed the NTT BRL's theoretical quantum physics research group.

Yasunobu Nakamura
An international leader in the experimental realization of superconducting quantum computing and hybrid quantum systems, he contributed to the creation of the first so-called flux qubit.

Kae Nemoto
She is a professor at the National Institute of Informatics (NII) and the Graduate University for Advanced Studies. She further serves as the director of the Global Research Centre for Quantum Information Science at NII. She is a pioneering theoretical physicist recognized for her work on quantum optical implementations of quantum information processing and communication.

Francesco Petruccione
He is a professor in Quantum Computing at Stellenbosch University where he is also the interim director of the National Institute for Theoretical and Computational Sciences. He spearheaded quantum technology research in South Africa. His main is to close the gap between fundamental research, innovation, and development to solve problems and ensure sustainable development.

John Preskill
A leading scientist in the field of quantum information science and quantum computation, who introduced the notion of Noisy Intermediate-Scale Quantum devices. He is the Richard P. Feynman Professor of Theoretical Physics at the California Institute of Technology, where he is also the Director of the Institute for Quantum Information and Matter.

Simone Severini
A leading researcher in quantum information and complex systems, particularly through the application of graph theory. He is currently Professor of Physics of Information at University College London, and Director of Quantum Computing at Amazon Web Services.

Peter Shor
The inventor of the efficient quantum algorithms for factoring and discrete logarithms that generated great interest in quantum computing, and a pioneer of quantum error correction.

Stephanie Simmons
Co-leads the Silicon Quantum Technology Lab at Simon Fraser University and is an international expert on the experimental realization of spin qubits in silicon, and in interfacing them with photon qubits.

Lieven Vandersypen
Renown for realizing one of the first demonstrations of Shor's algorithm for finding prime factors. He is a pioneer in quantum computing based on semiconductor quantum dots. His current interests are to demonstrate that the fundamental process of decoherence can be reserved, and to simulate complex materials and molecules using quantum dot arrays.

Gregor Weihs
He is Professor of Photonics at the Institute for Experimental Physics at the University of Innsbruck,

where he leads the Photonics group. His research in quantum optics and quantum information focuses on semiconductor nanostructures and on the foundations of quantum physics.

Frank Wilhelm-Mauch
A leading theoretician working closely with experimentalists, he focuses on modelling and controlling superconducting circuits. He is the director of the Peter Grünberg Institute for Quantum Computer Analytics.

David J. Wineland
World-leading experimental physicist awarded the Nobel-prize winner in 2012 (shared with Serge Haroche) "for ground-breaking experimental methods that enable measuring and manipulation of individual quantum systems."

Shengyu Zhang
A global expert in quantum algorithms and complexity, including recent work on quantum noise characterization. He leads the Quantum Lab at Tencent.

A.2 Realizations of quantum computers

Physical realizations

The various physical implementations of quantum computers have advantages and disadvantages in relation to factors such as (but not limited to):

- *scalability*, that is, the possibility of building and controlling larger and larger quantum devices with more and more qubits using physical/engineering resources that grow in a manageable way;
- compatibility with—and ease of implementation of—different computational models;
- typical decoherence time (that is, for how long quantum features like superpositions remain preserved and can be exploited);
- speed and precision with which gates can be applied.

The following is a very high-level classification of some physical realizations:

- **Quantum optics**, meaning that information is stored and manipulated in states of light; this includes polarization states or photon-number states, and can be implemented also on-chip by using integrated optics.
- **Superconducting systems**, meaning that information is stored and manipulated in electric circuits that exploit the properties of superconducting materials.
- **Topological systems**, meaning that information is stored and manipulated in some topological properties—that is, properties that depend on 'global' (geometric) properties insensitive to 'local' changes—of quantum systems.
- **Ion traps**, meaning that information is stored and manipulated in properties of ions (atoms with non-vanishing total electric charge) that are confined by electro-magnetic fields.
- **Quantum spin systems**, meaning that information is stored and manipulated in the internal degree of freedom called *quantum spin*; such systems may be realized in silicon, like standard microchips are, or in less conventional systems, like diamonds with point defects known as nitrogen-vacancy (or NV, in short) centers.
- **Cold atoms gases**, where neutral atoms (rather than ions) are cooled down to close to absolute zero. While ions repel each other because of their electric charge, neutral atoms do not, and can be trapped and arranged in very regular arrays via the use of laser beams that generate so-called optical lattices; the atoms can then be controlled all the way down to the level of individual sites in the lattice.

Models of computation

Besides many possible physical realizations of quantum computers, there are also various *models* of quantum computation. While many models are known to be computationally equivalent (that is, roughly speaking, they allow one to solve the same class of problems with similar efficiency), each model offers different insights into the design of algorithms or may be more suitable for a particular physical realization. One such model is the *circuit* model—or *gate* model—where transformations are sequentially performed on single and multiple qubits (see Figure 18).

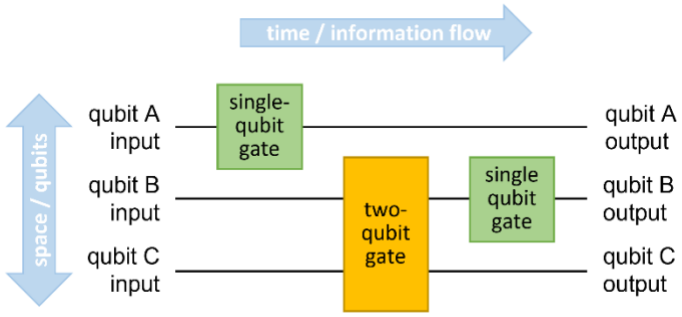

Figure 18 Illustration of the circuit/gate model for quantum computation. Each qubit corresponds to a horizontal line, so that multiple stacked lines illustrate many qubits. A qubit can be transformed individually by means of single-qubit gates, and two qubits can interact via a two-qubit gate. A given circuit transforms the initial input state of the qubits into their final output state, via the sequential action of said gates. The sequence of transformations is temporally ordered from left to right.

From the perspective of analysing the quantum threat timeline, it is useful to focus on the circuit model as there is a well-articulated path to implementing impactful cryptanalytic attacks.

In the circuit model, to perform arbitrary computations it is enough to be able to realize a finite set of *universal gates* which can be combined to generate arbitrary transformations. Such a set necessarily includes at least one gate that let multiple qubits interact, typically two at a time.

Historically, the following criteria, which are part of a larger set of desiderata, and which were listed by DiVincenzo in (DiVincenzo 2000) and hence are known as *DiVincenzo's criteria*, have been considered essential requirements for any physical implementation of a quantum computer:

1. *A scalable physical system with well characterized qubits.*
2. *The ability to initialize the state of the qubits to a simple fiducial state.*
3. *Long relevant decoherence times, much longer than the gate operation time.*
4. *A "universal" set of quantum gates.*
5. *A qubit-specific measurement capability.*

Unfortunately, the implementation of a single- or multi-qubit transformation can never be exactly the intended one, as the parameters defining a transformation are continuous, and because of the inevitable noise/decoherence. The quality of a gate implementation can be quantified by some notion of *fidelity*: the larger the fidelity, the closer the implementation of a gate is to the ideal one. A related parameter is the physical *error rate* with which gates are applied. In a sense, this parameter is the 'opposite' of fidelity. When characterizing the gate quality of experimental realizations or when studying the theory of how to correct them, most research groups use either the fidelity or the error rate.

Error correction, fault tolerance, and logical qubits

Errors and imperfections in the manipulation of (quantum) information, as well as decoherence, may be reduced by improving the physical implementation, including qubit control, but they cannot be entirely eliminated. Nonetheless, reliable storage and processing of quantum can still be achieved by employing *error correction* schemes: *logical* qubits are encoded into multiple *physical* qubits, so that errors affecting the underlying physical qubits can be detected and corrected, and logical information be protected. Error correction can ultimately lead to *fault tolerance* (Nielsen and Chuang 2000): under reasonable assumptions, one can prove that, if the error rate of the underlying physical components is low enough—below the so-called *fault-tolerance threshold*—then it is possible to implement logical encodings for information and information processing that can be made arbitrarily reliable, at the cost of using a number of physical qubits that is potentially much larger than that of the encoded logical qubits, but that still scales in a manageable way, at least theoretically.

Some more details on such codes and techniques can be found below, but they are not as relevant as keeping in mind that quantum error correction and fault-tolerance do pave the way to digital quantum computers: in principle, quantum computing devices can be made as reliable as needed, once some "quality standard" and some scalability & integration of the underlying physical qubits are achieved. We provide information on some specific error-correcting codes to 1) facilitate the understanding of the expert opinions on the topic and 2) to make it clear that developing codes that enable fault tolerance, also considering their ease of realization and tailoring them to specific physical implementation, is an on-going and very important area of research. Most relevantly, improvements in error-correcting codes and/or in their hardware implementation may speed up the quantum threat timeline.

An important issue in error correction is the kind of errors that the adopted error-correction scheme/code can detect and correct.

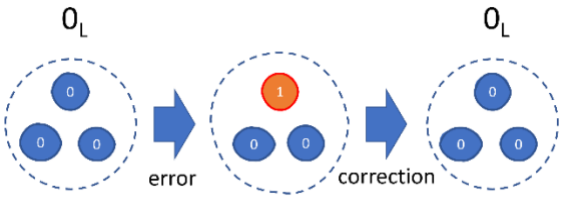

Figure 19 Example of classical information encoded logically. Several imperfect/error-prone physical bits (warped filled blue circles) are used to encode a logical 0, denoted 0_L (dashed perfectly round circle), by means of a repetition code: 0_L is encoded as 000 at the physical level. Errors can occur at the level of the physical bits, but they can be corrected, in this case by a simple majority-voting scheme, so that the logical bit is preserved. If the probability of a physical bit flipping is small enough, the probability of a logical bit being affected by an error—in this case, flipping from 0_L to 1_L—is less than the probability of a physical flip. Quantum error correction can be seen as a generalization of classical error correction to protect quantum *information; for example, a quantum code must preserve also (logical)* superpositions *of 0 and 1.*

In the case of classical bits, and excluding loss, the only possible type of error at the level of a single bit is the so-called *bit-flip*, which causes a 0 to turn into a 1, and vice versa. On the other hand, qubits can also undergo a so-called *phase-flip* error. Quantum codes can be designed and implemented that deal with just one of the two kinds of errors, but to protect quantum information both kinds need to be dealt with. Another important concept is that of *distance*, which roughly corresponds to the number of physical (qu)bits affected by an error that the error-correction scheme can handle. For example, the classical repetition code illustrated in Figure 19, using three physical bits to encode one logical bit, detects and corrects a single bit-flip error but would mishandle two bit-flips—confusing a logical 0 for a logical 1, and even introducing more physical errors upon correction. The special

properties of quantum information prevent the use of simple repetition codes, but, in general, the ability to correct against more kinds of errors and against errors affecting more qubits leads to a higher number of physical qubits needed to encode a single logical qubit.

Examples of error correcting codes

Surface codes, which are an instance of so-called topological quantum error correcting codes (Kitaev 2003), are currently among the leading candidates for large-scale quantum error correction.

The surface code (Fowler et al. 2012) allows for the detection and correction of errors on a two-dimensional array of nearest-neighbour coupled physical qubits via repeatedly measuring two types of so-called stabilizers generators. A single logical qubit is encoded into a square array of physical qubits. A classical error detection algorithm must be run at regular intervals (surface code cycle) to track the propagation of physical qubit errors and, ultimately, to prevent logical errors. Every surface code cycle involves some number of one- and two-qubit physical quantum gates, physical qubit measurements, and classical processing to detect and correct errors (i.e., decoding). Surface codes can provide logical qubits with lower overall error rates, at a price of increasing the number of physical qubits per logical qubit and the cost of decoding.

The *color code* (Bombin and Martin-Delgado 2006), is a generalization of surface codes, produced by tiling a surface with three-colorable faces and associating a distinct variety of stabilizer generator with each color (usually red, green, and blue). The surface code is a color code with only two colors (two types of stabilizers). These color codes combine the topological error-protection of the surface code with transversal implementations of certain gates (so-called Clifford gates), allowing for increased ease in logical computation, at a price of less efficient decoding algorithms.

Lattice surgery is a technique to merge and split surface codes to implement fault-tolerant interactions between qubits encoded in separate surface codes (Horsman et al. 2012).

Low-Density Parity Check (LDPC) codes have widespread use in the handling of classical information, as they have an essentially optimal scaling in terms of rate of encoding—the ratio between reliable logical bits and underlying faulty bits. Significant effort has recently been put into researching good *quantum LDPC codes*, which are characterized by the constraint that the number of underlying physical qubits involved in each error check and the number of checks each qubit is involved in are bounded by a constant (Breuckmann and Eberhardt 2021). One challenge with quantum LDPC codes is that the qubits used in the encoding and in the error correction, despite being “few”, may be far apart.

A.3 Questions

Regarding the wording of the core questions, in general we wanted to minimize the chances that the respondents could interpret them very differently. For example, questions like "when will we have useful quantum computers?" or "is it likely that a quantum computer will break cryptography in 10 years?" would have been far too vague. Some could have assumed that a useful quantum computer could have just a few dozen physical qubits that can demonstrate some proof-of-concept speed-up over currently known classical methods. Others could have assumed that a useful quantum computer will require thousands of logical qubits (and thus perhaps millions of physical qubits) and should be performing something of immediate commercial value. Even sticking to cryptographic applications, it is important to pose questions in the right way: a quantum computer breaking RSA-2048 in 10 years may be unlikely, but is it 49%, 10%, or 1% unlikely? Some of the above considerations and goals are in—perhaps, unavoidable—tension for some of the questions.

Given the scope of our survey, and the above general principles and considerations, we proceeded as follows:

- We kept the questions largely focused on the issue of the implementation of fault-tolerant quantum computers that would be able to run quantum algorithms posing an actual threat to cryptosystems.
- We sought a range of relevant perspectives. Already in 2019, we invited a select number of respondents with authoritative and profound insights. They provided a great variety of expertise on the most recent developments and the next steps needed towards the realization of fault-tolerant quantum computers. The same philosophy guided the selection of respondents in the subsequent surveys, including this one.
- Considering the quality of the pool of respondents, all very busy professionals and researchers, we kept the questions limited in number, so that the estimated time to complete the questionnaire was less than 30 minutes. In some cases, to secure responses to at least the major key question revolving around the quantum threat timeline, we gave the option to provide input about only such a key question.
  NOTE: Given the latter flexibility, not all respondents have provided answers to all questions, some of which were optional to begin with.
- Given the inherent uncertainty in the progress towards realizing a quantum computer, we asked the respondents to indicate in a relatively coarse-grained fashion how likely something was to happen.
- We did keep several of the questions at the basis of previous reports the same or very similar, so to be able to detect a change in opinions.
- On the other hand, we modified to some extent the set of questions from survey to survey, due to:
  - recent developments in the field (such as the efforts shifting more and more towards quantum error correction and the realization of logical qubits) and in the economic, political, and social scenario;
  - the respondents' feedback from previous surveys;
  - the desire to seek opinions about other relevant aspects of the quantum threat timeline.
- For the non-free-form multiple-choice answers, we gave the possibility to leave more nuanced comments. This mitigated to some extent the issue of the experts potentially responding to the same questions under a different set of assumptions and allowed us to collect insightful opinions.

Preliminary questions involved identification of the respondent and gauging their familiarity with different subfields of quantum computing research as well as implementations.

Here is a list of the main questions, grouped by questionnaire section.

Questions about "Implementations of quantum computing"

***Q**: Please indicate the potential of the following physical implementations for realizing a digital quantum computer with ~100 logical qubits in the next 15 years.*

Physical implementations listed: Superconducting Systems, Trapped Ions, Quantum Optics (including integrated photonics), Quantum spin systems in Silicon, Quantum spin systems not in Silicon, Topological Systems, Cold Atoms, Other

Options for answer: "Not promising", "Some potential", "Very promising", "Lead candidate", "No opinion"

Questions about "Timeframe estimates"

***Q (key question)**: Please indicate how likely you estimate it is that a quantum computer able to factorize a 2048-bit number in less than 24 hours will be built within the next 5 years, 10 years, 15 years, 20 years, and 30 years.*

Possible classification for each period of time:

1. Extremely unlikely (< 1% chance)
2. Very unlikely (< 5% chance)
3. Unlikely (< 30 % chance)
4. Neither likely nor unlikely (about 50% chance)
5. Likely (> 70 % chance)
6. Very likely (> 95% chance)
7. Extremely likely (> 99% chance)

***Q:** Various reasons for why a cryptographically-relevant quantum computer make take 30 years or longer to be built (if ever) have been articulated. Please indicate your opinion on the issues listed below, which may be among the reasons for an exceptionally long timeline.*

Concerns listed:

- Yet unappreciated fundamental trade-offs in controlling quantum features for cryptographically-relevant computational advantage (something akin to the uncertainty principle)
- Yet unappreciated standard-physics phenomena that may disrupt quantum computation (e.g., some unappreciated unavoidable source of correlated noise)
- New physics phenomena (e.g., random collapse of the wavefunction)
- Excessive technical challenges / requirements (e.g., the required scaling is practically impossible) not attributable to any of the above
- Other

Possible levels of concern:

- Concern is reasonable and has substantial likelihood (>30%)
- Concern is reasonable but somewhat unlikely (15%< likelihood <30%)
- Concern is reasonable but unlikely (5% < likelihood <15%)
- Concern is reasonable but very unlikely (likelihood <5%)
- Concern is not appropriate (likelihood <1% or the concern is unreasonable)
- No opinion

***Q**: What do you consider the most promising scheme for fault-tolerance?*

***Q**: What do you consider the most important upcoming experimental milestone to convincingly demonstrate the feasibility of building a cryptographically-relevant fault-tolerant quantum computer?*

***Q**: Please indicate your likelihood estimates for useful commercial applications of noisy intermediate-scale quantum (NISQ) processors – or of larger/less noisy processors but anyway not yet cryptographically-relevant – going beyond proof-or-concept and/or promotional activities, within the next 1 year, 3 years, 5 years, 10 years, and 15 years.*

Possible classification for each period of time the same as for the key question.

Questions on "Non-research factors that may impact the quantum threat timeline"

***Q**: You think that, over the next two years, the level of global investment (both by government and by industry) towards quantum computing will ...*

Options: Significantly Increase, Increase, Stay about the same, Decrease, Significantly Decrease, and Prefer not to answer

***Q**: Which of the following is currently the front-runner in the "global race" to build a scalable fault-tolerant quantum computer?*

Options [multiple selection was possible]: China, Europe, North America, Other(s)

***Q**: How likely are the following to be front-runners in the "global race" to build a scalable fault-tolerant quantum computer in five years?*

Each of "China", "Europe", "North America", "Other(s)" could be assigned one evaluation among "Likely", "Possibly", "Unlikely", "No Comment"

Questions on "Current progress in the development of a cryptographically-relevant quantum computer"

***Q**: What has been the most significant recent (since the second half of 2022) achievement in the progress towards building a fault-tolerant quantum digital computer?*

***Q**: What do you consider to be the next essential step towards building a fault-tolerant quantum digital computer? (something that could reasonably be achieved by approximately Summer 2024)*

***Q**: Please comment freely on the present and near-future status of development of quantum computers.*

A.4 Responses and analysis

In this section of the Appendix we provide some details on our methodology in handling and analyzing the responses.

Quantum factoring responses and analysis

We asked the respondents to provide an informative but rough estimate of the likelihood of the availability of a quantum computer able to factorize a 2048-bit number in less than 24 hours within a certain number of years. We provide here the raw aggregate counts of the responses.

LIKELIHOOD ESTIMATE	Within 5 years	Within 10 years	Within 15 years	Within 20 years	Within 30 years
Extremely unlikely (< 1% chance)	24	8	0	0	0
Very unlikely (< 5% chance)	6	12	7	0	0
Unlikely (< 30% chance)	4	7	10	6	0
Neither likely not unlikely (~ 50% chance)	2	4	10	11	87
Likely (> 70% chance)	1	4	4	9	10
Very likely (> 95% chance)	0	2	6	6	10
Extremely likely (> 99% chance)	0	0	0	5	9

We may associate each of the seven possible likelihood estimates to a sentiment between 1 and 7. One can then proceed to compute a (numerical) mean sentiment for each timeframe, averaged over the sentiment distribution of the experts. Note that this number carries both the uncertainty of the original estimates and the arbitrariness of the sentiment value assigned, but also note that we could have directly asked the experts to indicate how optimistic they were about the realization of a cryptographically relevant quantum computer in a given timeframe, on a scale from 1 to 7, where 1 is "Extremely unlikely (< 1% chance)", 2 is "Very unlikely (< 5% chance)", etc. It is reasonable to assume the answers would have been the same.

To derive from the responses the cumulative probability distributions as shown in Section 4.2, we assigned the following cumulative probabilities to each response, which are the largest and smallest ones compatible with the ranges among which the respondents could choose:

LIKELIHOOD ESTIMATE	OPTIMISTIC ASSIGNMENT	PESSIMISTIC ASSIGNMENT
Extremely likely (> 99% chance)	100%	99%
Very likely (> 95% chance)	99%	95%
Likely (> 70 % chance)	95%	70%
Neither likely nor unlikely (about 50% chance)	70%	30%
Unlikely (< 30 % chance)	30%	5%
Very unlikely (< 5% chance)	5%	1%
Extremely unlikely (< 1% chance)	1%	0%

The period option "More than 30 years, if ever" was implicit (not listed), and is trivially associated with a cumulative probability of 100%.

The resulting cumulative probabilities of the experts have simply been averaged for both the optimistic assignment and the pessimistic assignment.

General considerations on the reliability of the experts' estimates

We list here some considerations about factors that may influence the general reliability of the responses and/or lead to apparent changes in opinion trends:

- First and foremost, a general warning and an invitation to caution:
  - While the experts' likelihood estimates provide insight into the quantum threat timeline, the results of our surveys must always be interpreted cautiously.
  - The experts who take part in our surveys are uniquely qualified to estimate the quantum threat timeline, but that does not imply that any of them can correctly indicate what is going to happen and when.
  - Both in this survey and in the previous ones, several experts themselves have explicitly admitted the difficulty of making reliable forecasts.
- Considering averages does not provide necessarily the *best* possible estimates.
- When the pool of respondents changes from survey to survey, it may affect substantially the averages / the consensus.
- Statistically speaking, the number of respondents in our surveys is relatively small. Moreover, the time frame considered as well as the likelihood intervals constitute few, relatively coarse-grained bins. These factors may combine so that resulting estimates fluctuate noticeably form survey to survey, just because of few respondents answering slightly differently than they had done in the past. For example, if a respondent feels that a likelihood is around 25-35%, they might reasonably select "<30%" or "approximately 50%", and "switch" choice from one survey to the next, relatively randomly.
- The previous point is relevant even further when we adopt the approach of estimating likelihood ranges by interpreting optimistically or pessimistically the experts' likelihood estimates; the reasons is that some of the likelihood ranges associated with some answers are larger than others.

- Especially from the perspective of someone working in quantum computing research and taking a survey like ours, the "time when a cryptographically relevant quantum computer will become available" is not a random value whose probability distribution is fixed. Our respondents are hard at work to make such a device become a reality, and the progress they achieve year after year is such that they are gaining a better understanding of the hurdles towards building it and of what needs to be done for circumventing them. This better understanding might increase confidence in the eventual realization of a quantum computer, but might also allow them to better estimate how long it might take to overcome certain challenges. This corresponds to updating the above-mentioned distribution, for example making it more peaked some time in the future and, without contradiction, lower in the shorter term.
- Societal factors, including real or perceived issues related to the economy, may affect both the actual progress and perceptions/expectations about progress.