Global Risk Institute **FULL REPORT**

2021 Quantum Threat Timeline Report

GLOBAL RISK INSTITUTE

Authors: Dr. Michele Mosca
Co-Founder & CEO
evolutionQ Inc.

Dr. Marco Piani
Senior Research Analyst
evolutionQ Inc.

January 2022

Scope

This report focuses on estimates for the timeline of the threat posed to cybersecurity by quantum computers. It reflects the views of nearly fifty experts in the field of quantum computing research. The report follows versions compiled in 2019 (Mosca & Piani, 2019) and 2020 (Mosca & Piani, 2021); it provides the most recent opinions offered by these experts and examines the evolution of their views over the past three years due, for example, to scientific or technological developments or to changes in investment levels.

Contents

1 Introduction / background

Here and in the Appendix, we provide the background necessary to understand how quantum computers pose a threat to cybersecurity and how building such computers is an incredible scientific and technological challenge. This material is similar to what present in the previous reports (Mosca & Piani, 2019; Mosca & Piani, 2021). It is included to make the current report self-contained. Nonetheless, we provide more details on error-correction concepts, which become more and more relevant as the field moves towards the experimental implementation of error correction.

1.1 Quantum computing

Quantum mechanics is our best description of the inner workings of nature. It allows us to explain the behaviour of matter and energy at small physical scales, including the behaviour of fundamental particles like electrons, or of atoms and molecules. Quantum phenomena are 'fragile'. E.g., the uncontrolled interaction of a quantum system with its environment tends to 'wash out' its quantum features, in a process referred to as *decoherence*. Such fragility is of the utmost importance when we consider the following:

> *Quantum computing requires preserving and controlling quantum behaviour at a level and with a precision that has no precedence in human history.*

Information ultimately needs a physical substrate to be stored. A standard *bit* corresponds to binary information, either "False" (0) or "True" (1), and is stored in physical systems like a lightbulb or a switch which may be "off" or "on". Standard—also known as *classical*—computers process such kind of binary information. Is it possible to leverage quantum behavior to store and process information in a different way?

Quantum computing (Nielsen & Chuang, 2002) was born from taking this possibility seriously, and from the idea proposed by physicist and Nobel laureate Richard Feynman of a quantum computer that could allow us to study problems in physics that appear to be nearly impossible to handle with classical computers (Feynman, 1981).

The basic unit of quantum information manipulated by quantum computers is the quantum bit, or *qubit*. Unlike a standard bit, a qubit can store not only the two values 0 and 1, but also a *superposition*—technically, a linear combination—of them: the two values may be thought as "coexisting" and being processed at the same time.

Not only quantum computers will allow us to simulate quantum systems as proposed by Feynman but, by exploiting quantum features like superposition, and through cleverly designed algorithms, they will be able to tackle several mathematical, optimization, and search problems much faster than conventional computers (Nielsen & Chuang, 2002).

1.2 Quantum threat to cybersecurity

Widely used public-key cryptographic schemes rely on mathematical problems that are thought to be intractable by classical computers, the best-known example probably being the Rivest–Shamir–Adleman (RSA) cryptosystem (Rivest, Shamir, & Adleman, 1978). RSA is based on the difficulty of finding the prime factors of large numbers.

Such schemes may be broken by quantum computers. For instance, RSA can be attacked by implementing Shor's algorithm (Shor, 1997). Furthermore, the ability of a quantum computer to search through a solution space with 2^n values (i.e., all the possible combinations of values that n bits can assume) in roughly $2^{n/2}$ steps (Grover, 1996) would also weaken symmetric-key cryptography.

The threat posed by quantum computers could lead to a catastrophic failure of cyber-systems, both through direct attacks and by disrupting trust. The quantum threat can be mitigated by adopting new cryptographic tools which are designed to be resistant to quantum attacks. These so-called *quantum-safe cryptographic tools* can be conventional or quantum in nature. The first kind amounts to adopting cryptographic protocols based on problems that are hard or, at least, strongly believed to be hard also for quantum computers. The second kind of quantum-safe tools are based on quantum phenomena themselves, as in the case of quantum key distribution (Nielsen & Chuang, 2002). However, transitioning to quantum-safe cryptography is both arduous and delicate (Mosca M. , 2013): it requires the development and deployment of hardware and software solutions, the establishment of standards, the migration of legacy systems, and more.

With the necessity to devote enough time to an orderly and safe transition to a "post-quantum" world, the urgency for any specific organization to complete the transition to quantum-safe cryptography for a particular cyber-system relies on three simple parameters[1]:

- $T_{\text{SHELF-LIFE}}$ **(shelf-life time)**: the number of years the information must be protected by the cyber-system;
- $T_{\text{MIGRATION}}$ **(migration time)**: the number of years needed to properly and safely migrate the system 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.

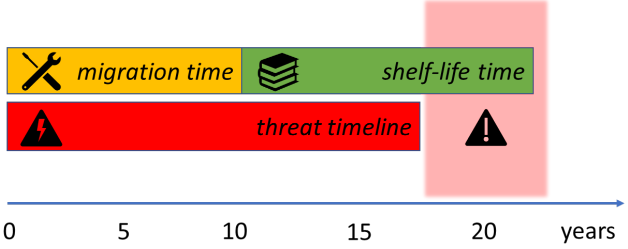

Figure 1 The timeline for the development of quantum computers that may pose a threat to cybersecurity should be compared with the time needed to migrate the cyber-system to post-quantum security combined with the shelf-life time of the data to be protected. See main text for details.

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).

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

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

Rushing the process of migration might itself create security issues which could be exploited even by attackers using only standard computers.

For example, problems might arise from gaps and omissions, from design flaws, or from implementation errors. Interoperability and backward compatibility may also suffer.

While the security shelf-life $\boldsymbol{T}_{\text{SHELF-LIFE}}$ is generally a business decision or dictated by regulations, assessing the threat timeline $\boldsymbol{T}_{\text{THREAT}}$ is not a straightforward task. There are numerous scientific and engineering challenges to overcome before building a quantum computer capable of breaking existing cryptographic schemes. While these challenges imply that the deployment of cryptographically-relevant quantum computers is likely to happen only many years in the future, it also means that unexpected breakthroughs may suddenly accelerate such a deployment.

Investments into the development of quantum computers and, in general, quantum technologies also play a major role in the speed of progress and may reduce the maximum available migration time. It is then worth considering that such investments have grown enormously in recent times (see also Section 1.5), coming from all kinds of sources: governments/funding agencies, (large) pre-existing companies, and private investors supporting newly established start-ups.

1.3 Realization of quantum computers

Quantum information can be encoded and processed in many different physical systems that behave quantumly. The latter include, for example, quantum spins, or the polarization of fundamental particles of light—so-called photons.

Besides the issue of the specific physical realization, there are also various *models* of computation. One such model, particularly useful when discussing the quantum threat, is the *circuit* model—or *gate* model—where transformations are sequentially performed on single and multiple qubits. More details can be found in the Appendix.

Regardless of the specific implementation, a common issue to contend with is that of the previously mentioned decoherence. One specific reason is that a quantum system designed to be controlled by an external user is particularly prone to interact with its environment. Such an interaction leads to the loss of the very quantum features used to encode and process quantum information[2].

It is vital to ensure adequate preparation of the physical system, to maintain control of it, and to measure it reliably—a step necessary to, e.g., extract the answer that one has run the quantum computation to find—while isolating it from the surrounding environment. Given the miniature scale of the systems at play and the numerous potential sources of decoherence, this is a daunting task.

1.3.1 Physical realizations

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

[2] Classical information encoded and processed in classical computers can be ‘corrupted’ too, and steps are taken to protect it, but its ‘classical nature’ makes it much more robust.

- *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.

1.3.2 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 & Chuang, 2002): 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 in the Appendix, 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. In the Appendix 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.

1.4 Quantum computing before achieving fault tolerance

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

Despite their limitations, NISQ devices already constitute a very significant achievement in terms of our ability to control quantum systems. Substantial effort is being poured into finding ways in which such devices may be useful well before full-fledged quantum computers become available, also to justify and strengthen investments in the area and generate returns on such investments sooner rather than later. Related research is also directed to conclusively proving at least in principle that progress in quantum computation research has already widened the range of feasible computations.

“Quantum supremacy” (Preskill, 2018) [3] may be generally described as the ability for a quantum device to perform some computation that would be practically impossible for classical computers, independently of the usefulness of such computation. Criteria for firmly establishing whether a device has achieved quantum supremacy are somewhat ‘fuzzy’. The reason is that it is difficult to establish that no classical means—including even the most powerful existing classical supercomputers, and the best possible classical algorithms—would allow one to perform the same computation in a ‘reasonable’ time. Even if one was content with just ’known’—rather than abstractly ‘possible but still unknown’—algorithms, quantum supremacy can be considered as a moving target, because classical computers and known classical algorithms improve over time.

Google argued to have achieved quantum supremacy in (F. Arute et al., 2019), and the 2020 version of this report included a collection of opinions by the experts about the significance of such a result (Mosca & Piani, 2021).

The respondents we surveyed this year pointed to new and improved demonstrations of quantum supremacy as among the most significant results of the past year (see Section 5.5).

1.5 The flourishing quantum landscape

Quantum technologies—in particular, quantum computing—have received growing attention from major private companies, universities, and research centres, as evident also by the affiliations of our

[3] This terminology is somewhat controversial because it recalls, e.g., racial supremacy. Nonetheless 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 in some strictly technical sense.

pool of respondents. This interest has been supported and boosted by several national and transnational initiatives, like the National Quantum Initiative in the United States (Raymer & Monroe, 2019) and the Quantum Flagship Initiative in the European Union (Max, Kovacs, Zoller, Mlynek, & Calarco, 2019), with investments in the field of quantum technologies seen as strategic. In addition, many start-ups that specialize in various aspects of quantum computing research have been established, often supported by venture-capital investments. While a detailed description of such a flourishing quantum landscape is beyond the scope of this report, it is important to stress that investments in the area have been growing stronger year after year. The map of Figure 2 provides a necessarily partial illustration of how numerous "quantum companies" are and of their location/distribution worldwide. While most names are likely unfamiliar to someone outside the field, the reader may recognize the names of some major "traditional companies" that have nonetheless been investing towards the development of quantum computers.

The commercial interest in quantum computers is certainly positive and has significantly sped up their development. Nonetheless, already in the 2019 report several respondents had indicated the risk created by a combination of "hype" and high expectations for the field, which could lead to reduced funding and investments some time in the future, if those high expectations are not met due to a slow or slower-than-anticipated progress. Such a scenario, called by some "quantum winter", could trigger a negative feedback loop involving slow progress and decreasing funding, resulting in a substantial slowdown in the development of a cryptographically-relevant quantum computer. To better understand the likelihood of this scenario, and similarly to what done in 2020, we have asked this year's

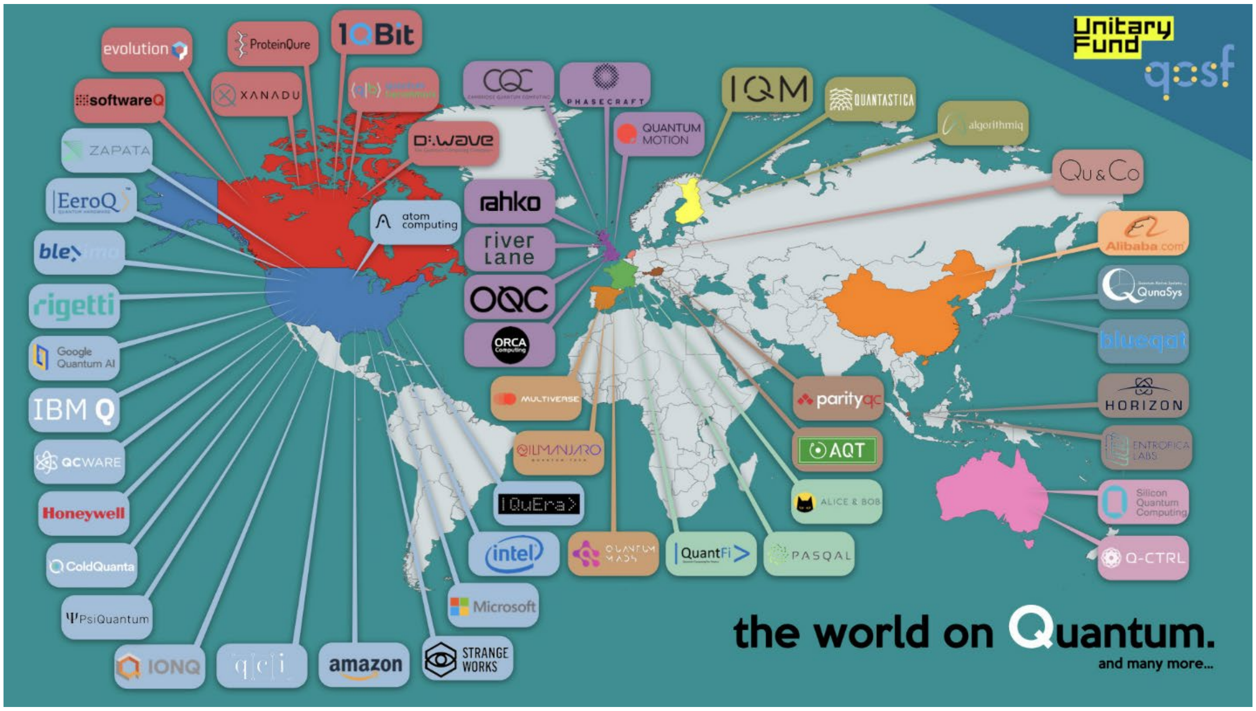

Figure 2 A partial map of some of the companies/start-ups focused exclusively on quantum technologies or at least with significant stakes in quantum technologies, grouped by country. Note the similarities with the map of the geographic location of our respondents, presented in Figure 4. [Map by Stefano Mangini]

respondents to express an opinion on whether they see the funding in the field increase, decrease, or stay stable in the next two years (Section 5.4.1).

2 Scope of this report

This document reports the results of a survey conducted by evolutionQ Inc., with the participation of 47 internationally leading experts on quantum computing. Following similar surveys conducted in 2019 and 2020, 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 some of the key questions via email.

We stress that we aim not only to provide a snapshot of the experts' opinions, but also 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. Examples of such external factors are the level of funding and societal changes, including the ones triggered by the ongoing COVID-19 pandemic.

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

The threat that quantum computers pose to RSA-2048 had already been considered previously. Within its more-than-200 pages, the National Academies of Sciences (NAS) report (National Academies of Sciences, Engineering, and Medicine, 2019) articulated an opinion on when quantum computers would threaten RSA-2048:

> *[…] [I]t is highly unexpected that a quantum computer that can compromise RSA 2048 or comparable discrete logarithm-based public key cryptosystems will be built within the next decade.*

While this insight may be helpful to someone who is responsible for managing the quantum threat to cybersecurity, it does not provide a threat timeline by estimating/reporting when we may see the realization of a cryptographically-relevant quantum computer; furthermore, much has happened since the NAS report was compiled.

Our series of reports differs from other reports, like the NAS report, in the following ways:

- *We seek a fine-grained picture of what leading experts think with respect to the timing and likelihood of the quantum threat.* While our survey obviously cannot provide definitive answers, we aim to depict a much better picture of what experts think about this question. Chances of 1%, 10% and 49% are meaningfully different shades of "unexpected". For example, do all experts agree, or is there a wide variance in opinions even amongst experts? What about 5 years, 15 years, 20 years? What does this timeline depend mostly on?
- *We want to provide estimates that take into account recent advancements.* There are many fast-moving parts in the field, and much can change in even just one year. Indeed, since the beginning, the development in the field has been characterized by steady progress combined with breakthrough results—even unexpected ones. Risk managers need to know how estimates of the quantum threat timeline are affected by such changes. For example, just in the year after the NAS report, in the public literature the overall cost estimates of breaking RSA-2048 has gone down by

about four orders of magnitude (Gidney & Ekerå, 2021; Gheorghiu & Mosca, 2019), and since then other players have joined the quest to build quantum computers.

- *We aim to track how the opinions of experts evolve over time.* This year's report is the third one in this series, with the potential to continue running such reports as long as the community finds them helpful[4].
- *The scope of our survey/report is much tighter and more focused than that of other reports, as it revolves about the quantum threat timeline alone.*
- *We ask specific questions individually to several leading researchers* and compile relevant statistics.
- *We ask the respondents to indicate what they judge as the most important milestones to pass*, or the necessary steps in the creation of a quantum computer.
- *We give the respondents the chance to provide free-reign comments on the status and expected evolution of the field, in doing so gaining substantial insight on what to expect and look out for in the future*.

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 & Riedel, 2020) the authors try to forecast future progress in the domain of quantum computing extrapolating past progress in the field, by looking at relevant metrics—roughly speaking, at how many effective logical qubits are available for computation. Sevilla & Riedel focus on superconducting implementations, and, similarly to what we do, on the task of breaking RSA-2048. Their estimates for when (super-conducting) 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 & Piani, Quantum Threat Timeline, 2019)". They also write that a cryptographically relevant quantum computer 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 implementations 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. For example, IBM announced in 2020 the expected progress of their family of superconducting quantum chips, going in parallel with software advances (Gambetta, 2020). Such a roadmap includes, for example, a 1,000-qubit machine in 2023, and machines with millions of qubits realizing fault-tolerant quantum computation after that (see Figure 3).

[4] Feedback from the quantum computing research community about the previous reports suggests that this is the case. We consider this is an indication that the report fills a relevant gap. It also corroborates the idea that the information / estimates the report provides have significant value for risk managers.

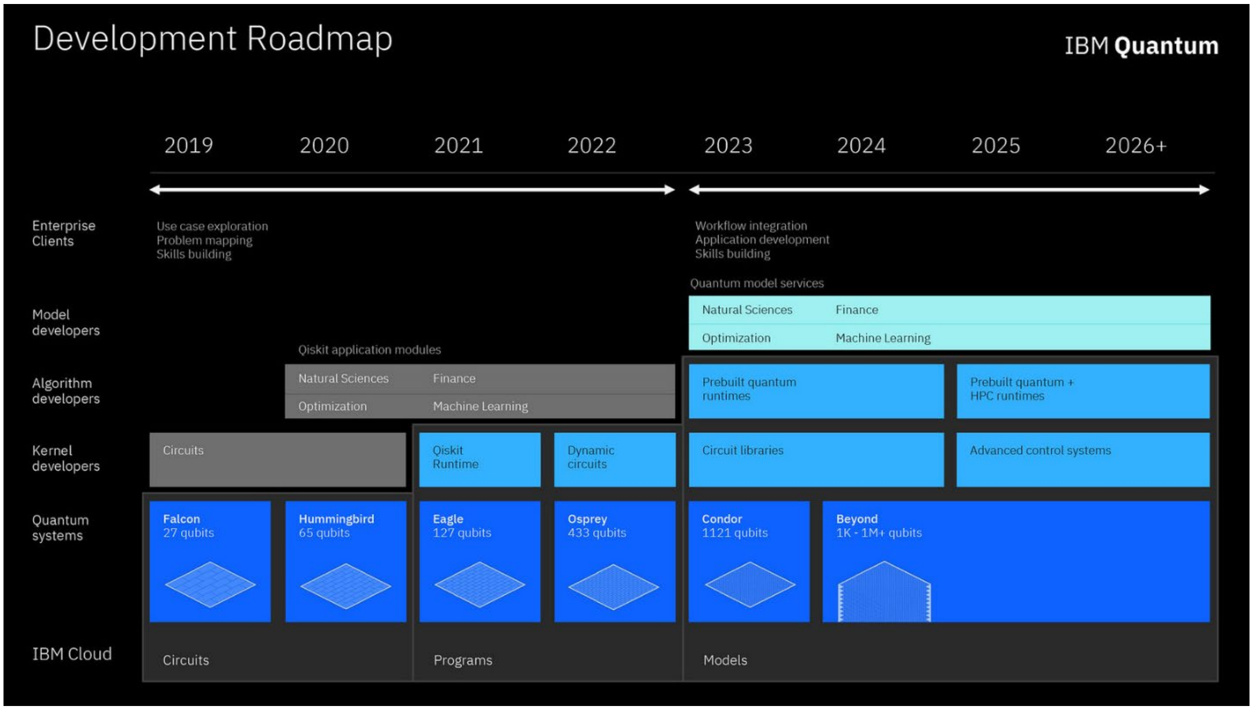

Development Roadmap

IBM Quantum

2019 2020 2021 2022 2023 2024 2025 2026+

Enterprise Clients

Use case exploration
Problem mapping
Skills building

Workflow integration
Application development
Skills building

Model developers

Quantum model services

Natural Sciences | Finance | Optimization | Machine Learning

Algorithm developers

Qiskit application modules

Natural Sciences | Finance | Optimization | Machine Learning

Prebuilt quantum runtimes

Prebuilt quantum + HPC runtimes

Kernel developers

Circuits

Qiskit Runtime

Dynamic circuits

Circuit libraries

Advanced control systems

Quantum systems

Falcon 27 qubits

Hummingbird 65 qubits

Eagle 127 qubits

Osprey 433 qubits

Condor 1121 qubits

Beyond 1K - 1M+ qubits

IBM Cloud

Circuits

Programs

Models

Figure 3 Development Roadmap of IBM Quantum, by IBM Research (unmodified; some rights reserved under the CC BY-ND 2.0 license). The inclusion of this roadmap in this report does not imply any endorsement and is only meant to provide additional information about the pace of development in the field of quantum computing.

3 Survey design and methodology

There was a range of non-trivial considerations in the phrasing of the questions, which we further discuss in the Appendix.

It was most important to understand the perspectives of the diverse range of expertise of the people who would be asked to complete the survey, and how the target audience of the report itself would interpret the questions and possible answers.

Given that this is a series of reports, we also wanted to strike a balance between keeping a fixed subset of questions—to be able to track the shift of opinions in time—and improving/replacing/adding questions to take into account both changes in the field and past feedback from the respondents.

Some of the questions were optional, as in past surveys. This applies, for example, to providing less technical and more free-form opinions on the state of the field.

To facilitate frank answers, we made the point—shared in advance with the respondents—of analyzing the estimates in an aggregate and anonymized fashion. For free-form answers/input, we gave respondents the option to avoid being quoted in this report, or, if quoted, to be quoted while still preserving anonymity.

Overall, we are confident that the above setup has allowed the respondents to provide answers to the best of their knowledge and expertise.

3.1 Questions

In what we consider an improvement with respect to the questionnaires of the previous years, the 2021 questionnaire was split into explicit sections about:

- the potential of various platforms / physical implementations for the realization of a scalable fault-tolerant quantum computer;
- estimates about when to expect the realization of a cryptographically-relevant fault-tolerant quantum computer, or of important steps along the way;
- societal and funding factors that may impact the timeline of the development of a cryptographically-relevant fault-tolerant quantum computer;
- recent and near-future significant progress towards building a cryptographically-relevant quantum computer.

A complete listing of the most relevant survey questions can be found in the Appendix.

The key question of the survey was:

Please indicate how likely you estimate 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,

with the following possible classification for each period:

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)”.

The choice of timeframe increments and of likelihood values is the result of trying to strike a balance between seeking informative answers while acknowledging the inherent uncertainty of this kind of forecast/estimate. We do not necessarily imagine that the notion of chance/likelihood has been statistically interpreted in a rigorous way by the respondents, but we wanted to make sure that all respondents assigned a similar meaning to words like, e.g., “likely” and “very likely”. In the analysis of the responses, we have taken the stance of considering the answers as simply indicating a “sentiment” or, alternatively, of interpreting them as actual probability estimates. We note that some respondents have explicitly indicated that the constraint for the solution to be found by the quantum computer in less than 24 hours is significative and has impacted their estimate.

Several factors inform the opinion of the experts, either consciously or unconsciously. They include, for example:

- potential recent breakthroughs in the science and in the technology being developed;
- changing levels of investment by countries, institutions, and companies;
- relatively unexpected events that affect society and the economy (e.g., the COVID-19 pandemic).

Correspondingly, one may expect a shift in opinions over time. We note that tracking changes in the opinions may be considered as important as taking snapshots of such opinions, because a shift in opinions may itself point to speed-ups or slowdowns in the development of quantum computers. For this reason, the above key question has been kept the same since the 2019 survey.

The other question that we considered as very important regards the likelihood of the realization of a fault-tolerant qubit, and was stated as follows:

> *Please indicate how likely you estimate that a single fully controllable fault-tolerant (logical) qubit within a scheme / architecture viable for scaling will be demonstrated within the next 1 year, 3 years, 5 years, and 10 years,*

with the range of potential answers the same as in the above factorization question.

We remark that this year’s version of the fault-tolerant-qubit question differed slightly but significantly from the wording used in the 2020 survey. Indeed, this year we have stressed the importance of *scaling*. The point is that the “plain” demonstration of a logical qubit—that is, of the reliable encoding and manipulation of an “abstract” unit of quantum information in multiple physical qubits—does not suffice to ensure a viable path to a full-fledged digital quantum computer. This is because the single logical qubit may be realized in a way that is impossible to scale to a sufficiently large number of logical qubits. In Section 5.3 and in the Appendix we report comments from the respondents that indicate how this requirement can be interpreted differently, depending also on the architecture used.

4 Participants

For the sake of consistency and of tracking trends, we aimed at securing the participation of as many as possible of the respondents to the 2019 and 2020 surveys. We were pleased that this was possible for the vast majority. Notably, 21 or the 22 participants in the 2019 survey took part also in the 2020 survey and in the present one.

Our respondents have been selected out of an initial shortlist of more than one hundred leading experts. From this list, we contacted several who, like the original respondents of 2019, were intended to provide a balanced—e.g., with respect to implementation types—and insightful range of opinions on the state of development of the field. Those who accepted were asked to complete the online questionnaire in about two weeks, but the response time has varied substantially.

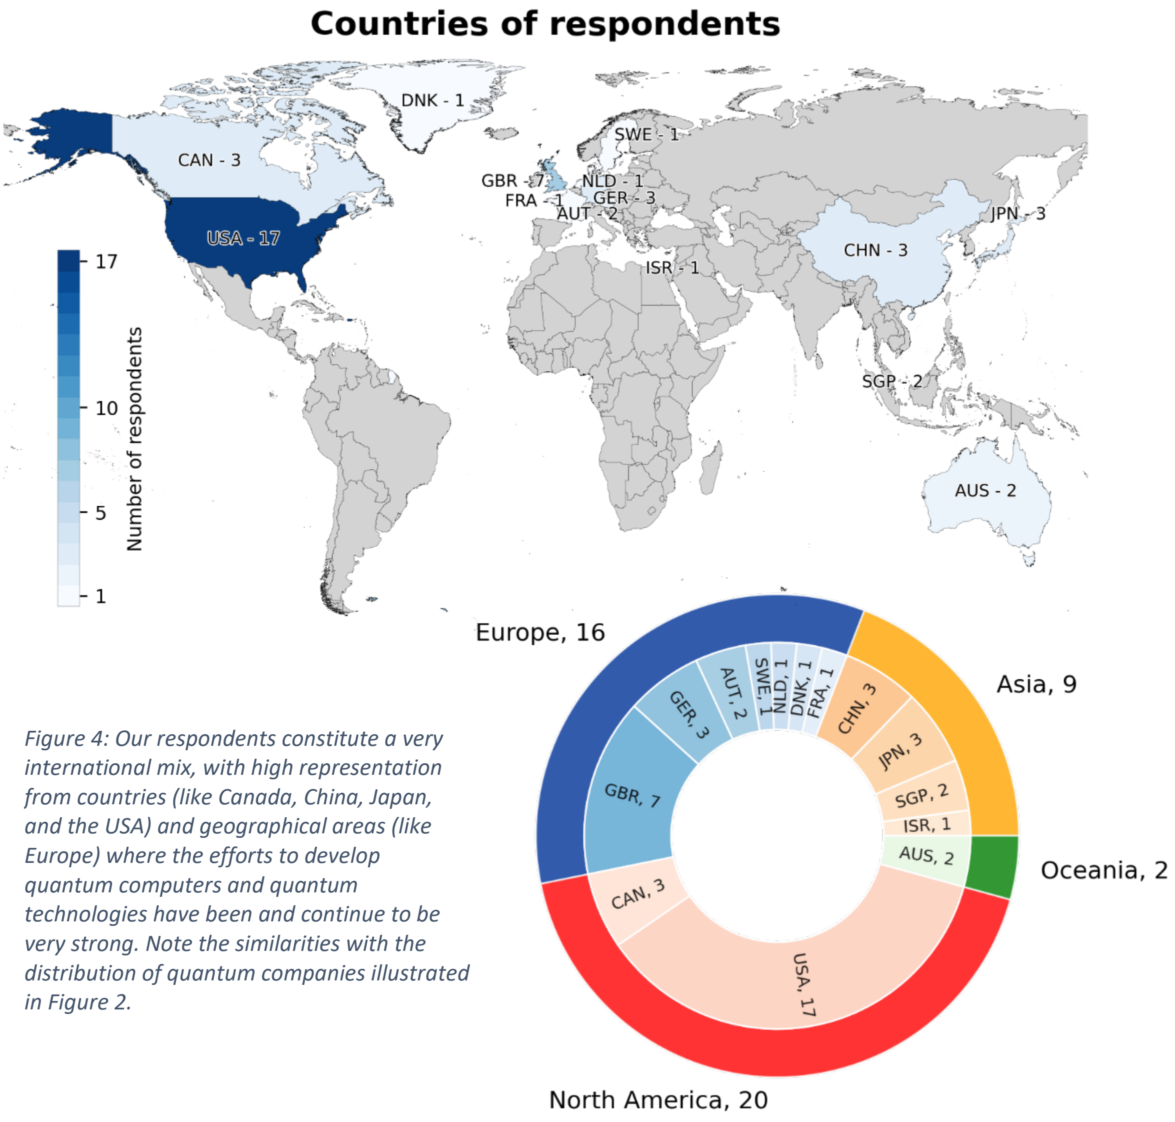

Figure 4: Our respondents constitute a very international mix, with high representation from countries (like Canada, China, Japan, and the USA) and geographical areas (like Europe) where the efforts to develop quantum computers and quantum technologies have been and continue to be very strong. Note the similarities with the distribution of quantum companies illustrated in Figure 2.

Some candidate respondents we contacted did not reply to our invitation. Others reported that they were unable to complete the questionnaire for various reasons, ranging from personal circumstances, to being too busy, to business strategy.

Overall, we were able to collect responses from 47 experts (see the Appendix for a complete list). Here we summarize graphically the composition of the group in terms of:

- country where they work (Figure 4),
- kind of activity they lead (Figure 5),
- kind of organization they belong to (Figure 6).

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

In summary, the pool of respondents comprised a diverse set of expertise and nationality, and a mix of university and private-sector researchers, representative of the diversity of the quantum computing community among its top players. We observe that the number of academics who also play some role in companies has grown in the years from survey to survey, reflecting and proving how the attention and

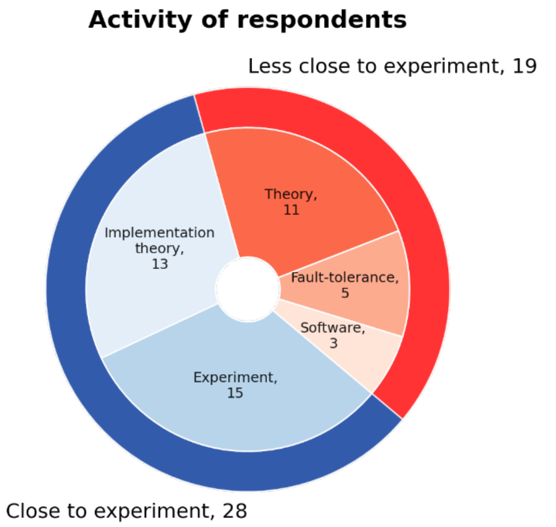

Figure 5: Our respondents cover a wide range of research activities. While the major division is between non-experimental research and experimental one, research that is not directly experimental can be very different. E.g., implementation theory focuses on guiding, supporting, and, in general, facilitating experimemental effort. Respondents are classified under simply "theory" if their more abstract abstract activity is not specificically related to experiments or implementations, or to fault-tolerance, or to software development.

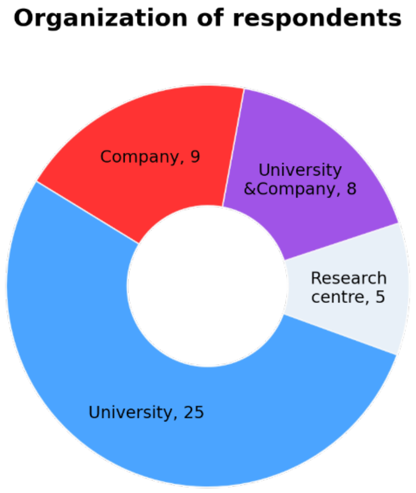

Figure 6 Most of the respondents work at universities, but some work at companies or research centres. Some researchers/academics may have some role in—or at least collaborate closely with—external companies. Compared to the previous years, a larger fraction of our respondents falls in the latter category, also because some past academic respondents have joined or founded companies.

the effort towards the commercialization of quantum technologies and of quantum computing is getting stronger and stronger.

In presenting the results of our survey in this report we have adopted a high-level/coarse-grained classification of research activities, which does not reflect the high specialization of the actual activities (Figure 5). We note that we refer to the kind of theoretical work that contributes to experiments, or that is in general concerned with implementations, as to *implementation theory.* We have also grouped together experimentalists and theorists who work on implementation theory, under the umbrella of experts who are “close/closer to experiment”. Conceivably, the latter group has a very informed&informative vantage point when it comes to judging the ongoing progress towards building a quantum computer.

5 Survey results

This section mirrors the survey structure indicated in Section 3.1. It provides an aggregate analysis of the key responses about:

- the potential of various physical implementations/platforms for quantum computing (Section 5.1),
- the quantum threat timeline (Section 5.2),
- the timeline for the implementation of a fully controllable fault-tolerant qubit in a scalable architecture, with experts' comments about fault-tolerant schemes (Section 5.3),
- the expected change in funding in support of quantum computing research (Section 5.4.1),
- the status and potential development of the so-called "quantum race" (Section 5.4.2),
- the estimated delay for the progress of the field due to the COVID-19 pandemic (Section 5.4.3).

It also provides:

- a selection of opinions about:
  - key recent developments in the field of quantum computing research, as highlighted by the respondents,
  - near-future (that is, approximately, by mid-2022) 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-2022;
- a collection of other notable remarks made by the respondents.

Where we deem appropriate, we analyze shifts in the responses as compared to responses from the last two 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. Similarly to what done in 2019 and 2020, we sometimes consider separately the responses of those participants who are deemed to be close/closer to experiments.

Not all the 47 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 have been modified/tweaked in their wording. 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. Nonetheless, where we notice a trend that could potentially be significant, we point it out, and, where feasible and/or appropriate, we provide a rationale that may explain it.

5.1 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.

Compared to 2019, in 2020 we slightly modified the categories the experts had to provide an opinion on (e.g., we introduced the category of "cold atoms"). This year, we have rephrased and condensed two of the questions posed in the previous years, with the goal to obtain what we hope are slightly more informative estimates. Specifically, we asked about the perceived potential for the realization of a significant number of logical qubits (~100) within what could be considered an intermediate time window with respect to the timeframes "probed" in our survey—15 years, while the timeframe that is the furthest in the future is 30 years.

The responses indicate a significant consensus that the present leading platforms are superconducting systems and trapped ions (Figure 7). This is consistent with the opinions collected in previous two surveys. We notice that superconducting qubits might have gained on trapped ions in our opinion poll; besides actual scientific and technological achievements, this could be due also to announcements of "aggressive" roadmaps by major companies (see, e.g., Figure 3 in Section 2).

It also appears that quantum optical implementations are strengthening their perceived potential. This might be due in part to the announcements of progress and future plans by companies that are working on such a type of implementation, one example being PsiQuantum, a quantum company which only recently came out of "stealth mode", revealing significant investments and an ambitious roadmap.

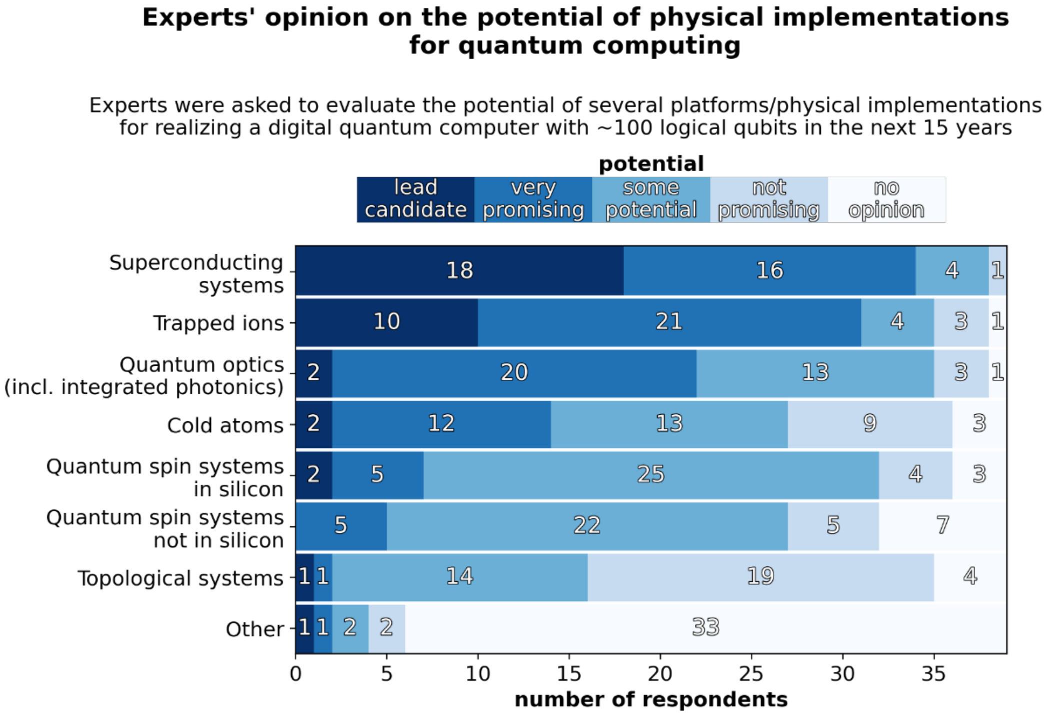

Figure 7: Similarly to previous years, superconducting-system implementations, followed by ion-trap implementations, are perceived as presently having some edge over other physical realizations.

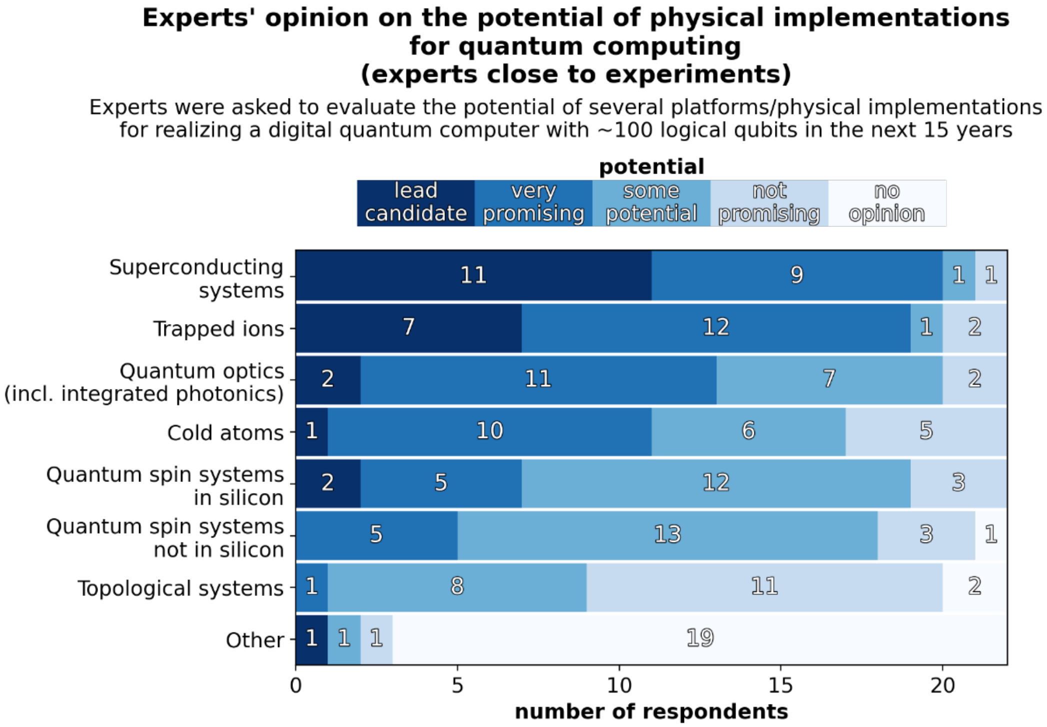

Figure 8 The opinions of the experts closer to experiments is not too dissimilar to that of the whole cohort of experts; nonetheless, there are small but potentially significant differences.

We note that among the "Other" systems indicated by the respondents are, for example, electron-charge qubits, which **Simon Benjamin** characterizes as

> *[..] being at interface of 'not promising' and 'some potential' since [such a platform] would require very substantial innovations[,]*

and hybrid solutions.

Figure 8 shows the responses of only those experts whom we consider as being closer to experiments. It shows a slightly stronger lead of superconducting systems and trapped ions; interestingly, it points also to a stronger positioning of cold atoms and of spin systems over quantum-optical implementations.

Appendix A.5 provides a selection of inputs from the experts that offer further insight into the potential of various platforms, and that both informed and corroborated the above interpretation of the multiple-choice results.

5.2 Quantum factoring

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 factor a 2048-bit number—that is, able to break RSA-2048—in less than 24 hours. See Section 3.1 for the exact formulation of the question.

Estimates on the practical requirements to achieve such a feat, also considering the imperfections of physical implementations, were presented for example in (Gheorghiu & Mosca, 2019)[5] and in (Gidney & Ekerå, 2021)[6].

The key outcome of our annual survey is presented in Figure 9, which provides the aggregate distribution of the responses of the experts and shows the estimated increase of the likelihood of the quantum threat as one moves from the relatively short-term future to the relatively long-term one. In Figure 10 we provide a graphical representation of the individual estimates. Figure 11 and Figure 12 are the corresponding representations of the responses for only those experts that are deemed closer to experiments. Several respondents articulated the difficulty inherent in making such kind of prediction, to the extent that one expert chose not to indicate any likelihood. Hence, only 46 opinions rather than 47—the latter being the total number of respondents who contributed to our survey—appear in all the figures and tables in this section.

It is possible to appreciate the ample range of opinions, particularly when looking at the plots of individual responses (Figure 10 and Figure 12). Some experts appear to be relatively optimistic and some others relatively pessimistic about the rate of development of quantum computers. The graphs of individual responses also illustrate that the pace of progress may be differently estimated, as some respondents assign an initial low probability, which subsequently, in later timeframes, grows faster than in the responses of other experts. Another observation is that for some respondents the likelihood estimate "saturates" earlier than 30 years in the future and/or at a likelihood lower than the highest possible assignment. This may be interpreted as an expression of uncertainty about the future, including for example the chance that some unexpected non-trivial technological challenge—perhaps even a fundamental showstopper—may emerge; such an eventuality could send us back to the drawing board on some key aspects of building a large-scale quantum computer.

[5] The Global Risk Institute has published regular updates of the estimates of (Gheorghiu & Mosca, 2019); the updates consider recent developments and complement from a more technical perspective the present opinion-based series of reports (Gheorghiu & Mosca, 2021).

[6] One of the authors of the latter paper is part of our pool of respondents.

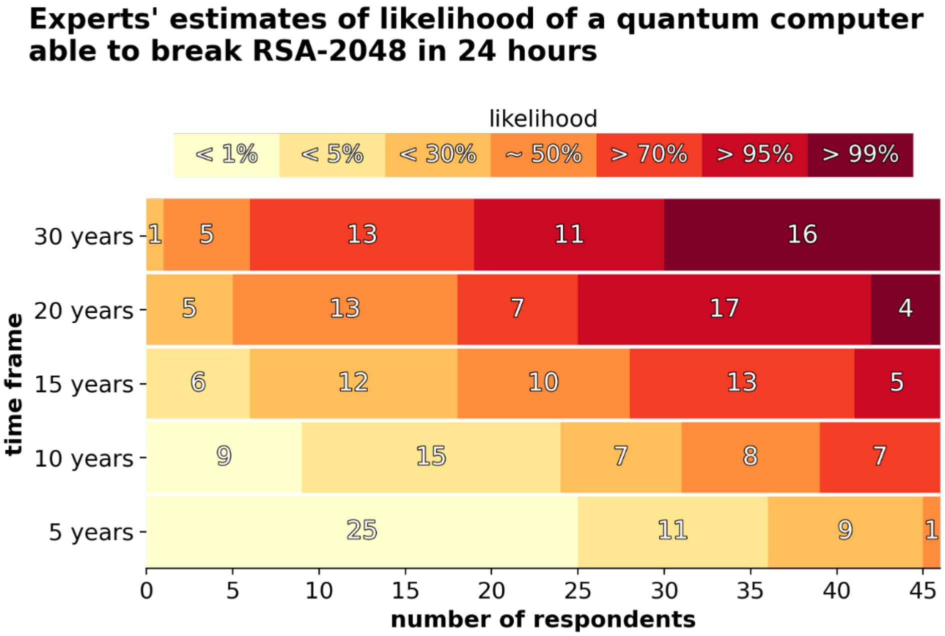

Figure 10 This figure illustrates the central information collected through our survey. The experts were asked to indicate their estimate for the likelihood of a quantum computer that is cryptographically relevant—in the specified sense of being able to break RSA-2048 in 24 hours—for various time frames, from a short term of 5 years all the way to 30 years.

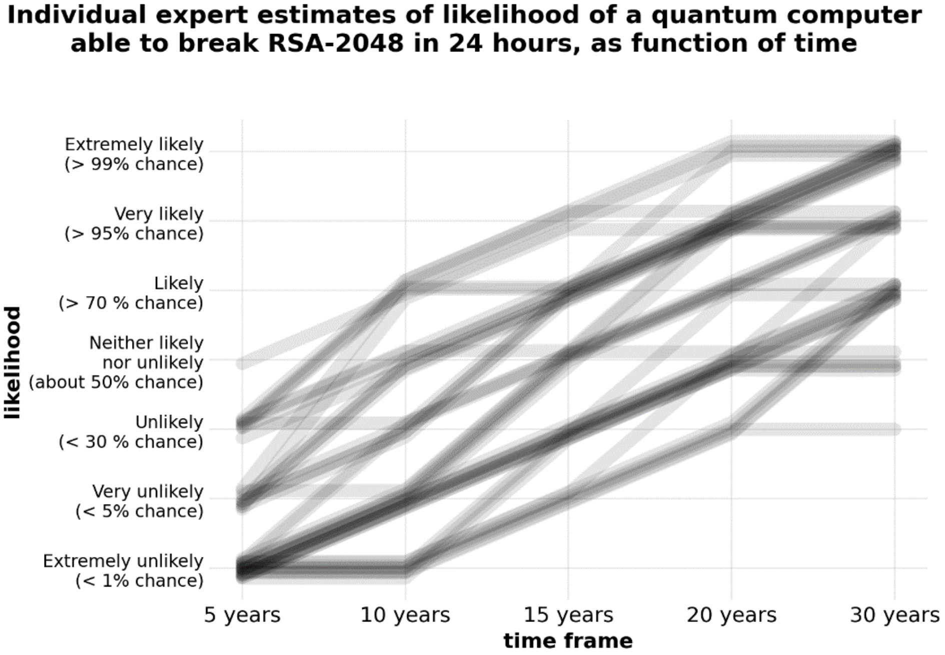

Figure 10 Opinions of individual *experts about the likelihood of the existence of a quantum computer able to factorize a 2048-bit number—that is, able to break RSA-2048—in at most 24 hours. With this figure we want to convey the diversity of opinions of the experts, both in terms of the likelihood within each time frame and of how the likelihood is estimated to evolve in time. Experts whose estimates grow in a similar way with the timeframe considered correspond to "reinforcing trajectories" in this representation, giving rise to a "joint trajectory".*

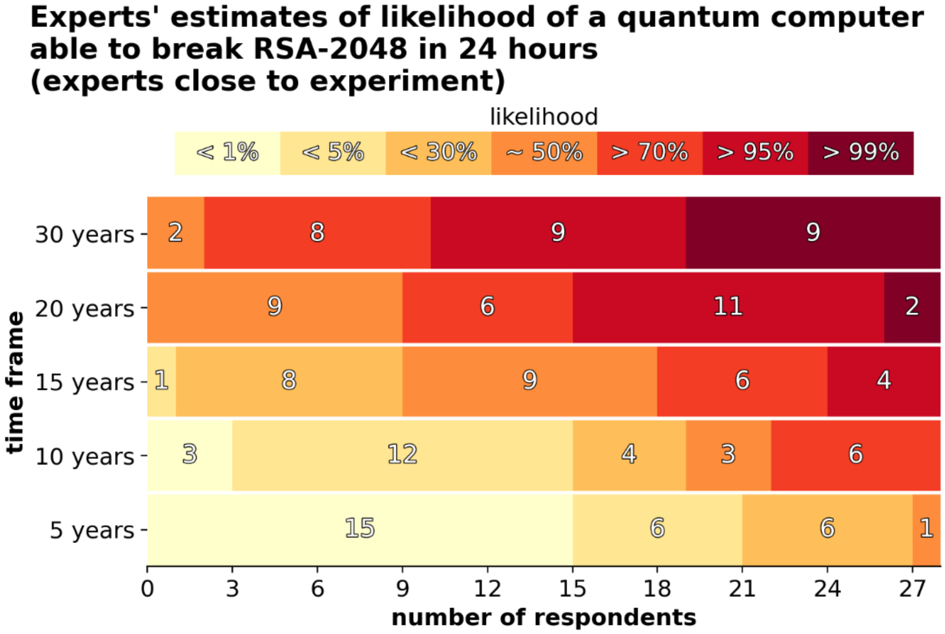

Figure 12 Estimates for the likelihood of a quantum computer that is cryptographically relevant—in the specific sense of being able to break RSA-2048 in 24 hours—for various time frames, limited to the 28 experts deemed to be closer to experiments. Such a subset of experts appear to provide estimates that do not differ substantially from those of all respondents (see Figure 9).

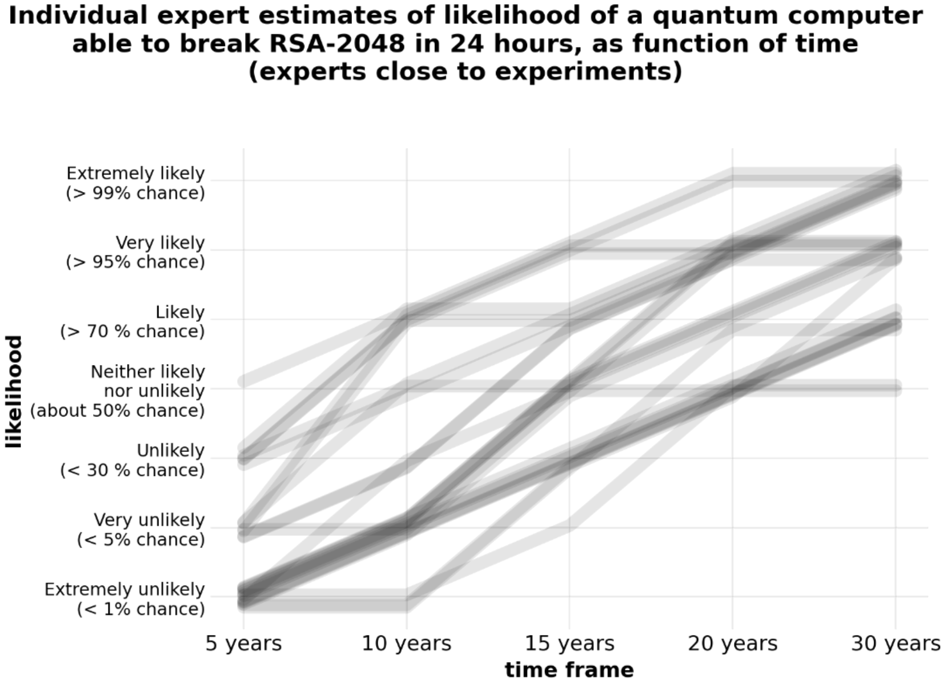

Figure 12 Opinions of only the individual experts close to experiments about the likelihood of having a quantum computer able to factorize a 2048-bit number—that is, able to break RSA-2048—in, at most, 24 hours. Each line represents the opinion of one expert. See caption of Figure 10 and main text for more details.

Despite the great variability of the responses, some valuable patterns emerge (see Figure 9):

NEXT 5 YEARS: Most experts (25/46) judged that the threat to current public-key cryptosystems in the next 5 years is “<1% likely”. About a quarter of them (11/46) judged it relatively unlikely (“<5% likely”). The rest selected “<30%” (9/46) or “about 50%” (1/46) likely, suggesting there is a non-negligible chance of an impactful surprise within what would certainly be considered a very short-term future.

NEXT 10 YEARS: Still more than half of the respondents (24/46) judged the event was “<1%” or “<5%” likely, but already 15/46 felt it was “about 50%” or “>70%” likely, suggesting there is a significant chance that the quantum threat becomes concrete in this timeframe.

NEXT 15 YEARS: More than half (28/46) of the respondents indicated “about 50%” likely or more likely, among whom 13 indicated a “>70%” likelihood, and 5 an even higher “>95%” likelihood. This time frame appears as a tipping point, as the number of respondents estimating a likelihood of “about 50%”, or larger, become the majority.

NEXT 20 YEARS: Roughly 90% (41/46) of respondents indicated “about 50%” or more likely, with 21/46 pointing to “>95%” or “>99%” likely. This indicates a significant bias toward viewing the realization of the quantum threat as substantially more likely than not within this timeframe.

NEXT 30 YEARS: Forty experts out of 46 indicated that the quantum threat has a likelihood of 70% or more this far into the future, with 16/44 experts indicating a likelihood greater than 99%. Thus, there appears to be a relatively low expectation of any fundamental show-stoppers or other reasons that a cryptographically-relevant quantum computer would not be realized in the long run.

We can directly represent via a heatmap the percentage of respondents that gave a specific likelihood estimate for a certain time frame (see Figure 13). Notice that it is akin to a coarse grained and blurred version of the “trajectory plot” of Figure 10. What gets eliminated is the individual “trajectory” information, that is, the information on how the likelihood estimated by each single respondent changes moving from one time frame to the next time frame.

The heatmap representation shows and emphasizes both the variance of opinions at every time frame, and the shift of the estimates in time towards larger likelihoods. It indicates that the experts tend to agree that the quantum threat is (very) unlikely to become concrete in the short 5-year term but that it will instead likely materialize within the long 30-year term. What the experts “disagree” about is how quickly the likelihood of the quantum threat grows in time, to move from (very) unlikely in the short term to (very) likely in the long term.

To gain more insight into how the experts’ estimates shift from one timeframe to the next one, we can adopt at least two ways to further summarize the experts’ estimates.

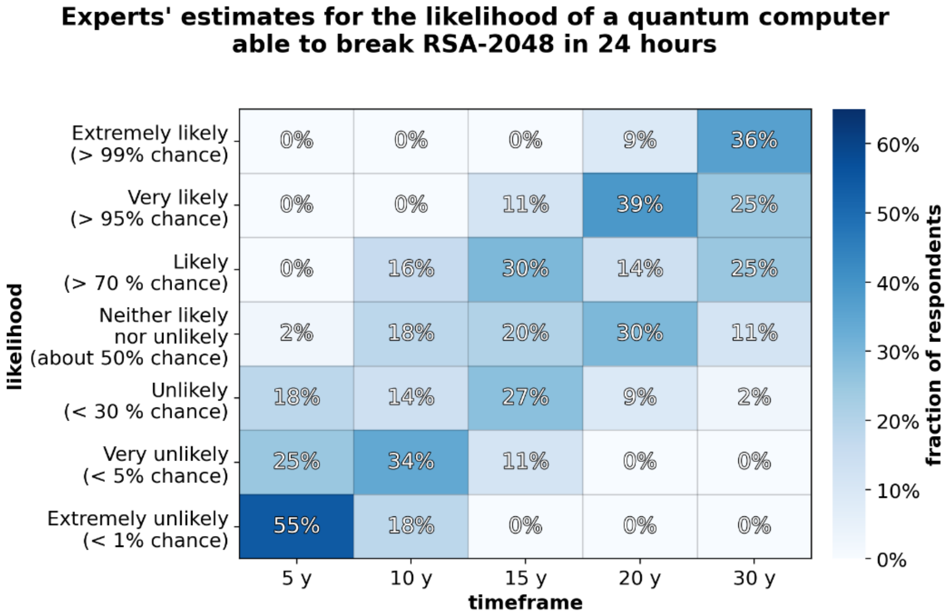

Figure 14 This heatmap represents the fraction of experts who assigned one of the specific listed likelihoods (left axis) to the existence of a quantum computer able to break RSA-2048 in less than 24 hours, within a certain time frame in the future (horizontal axis).

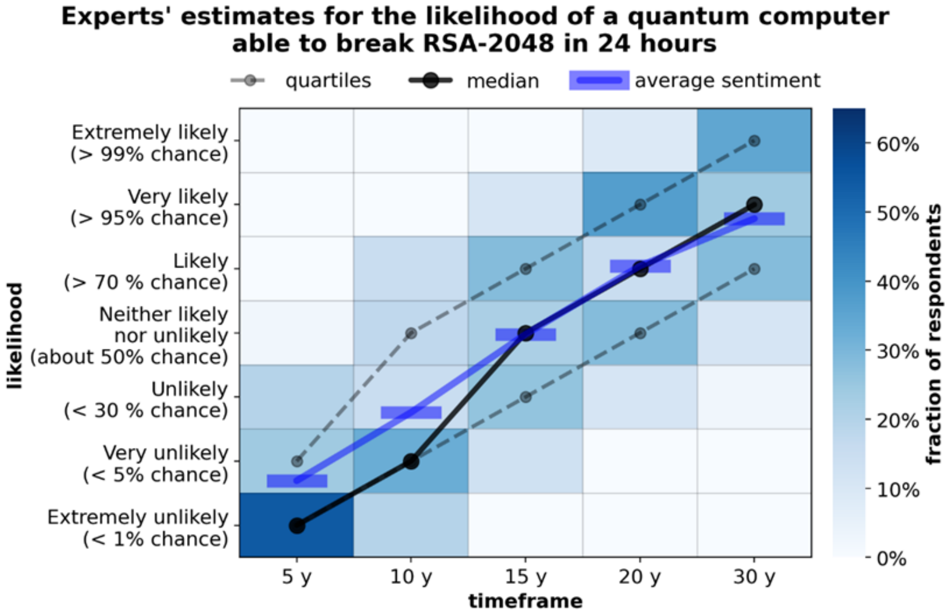

Figure 14 For each timeframe we can calculate the average sentiment *of the respondents, as discussed in the main text, and indicated here by the horizontal bar within each timeframe column. E.g., in the 5-year timeframe, the average sentiment is "equivalent" to a likelihood in between "extremely unlikely" and "very unlikely". The continuous blue line connects the average sentiment we calculated for each time frame. We also plot a graphical representation of the naïve* median *sentiment for each time frame (continuous black line with circle markers) and quartiles (dashed partially transparent black lines with circle markers).*

In the first approach, the expert likelihood estimates are considered a measure of how optimistic or how pessimistic each respondent is about the realization of a cryptographically-relevant quantum computer within each timeframe—i.e., (a measure of) their *sentiment* in that regard.

The result of this approach is shown in Figure 14, along with the heatmap introduced in Figure 13 as backdrop to indicate the variance in the opinions.

In a second approach, 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[7]. 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" or, alternatively, "optimistic" interpretation of the answers' ranges. This approach allows us to calculate an average cumulative probability distribution, both for the optimistic and pessimistic interpretation. Had each respondent selected a precise estimate within the respective ranges, then the point estimate for the likelihood would sit in the range between the optimistic-interpretation and pessimistic-interpretation curves. In turn, the latter

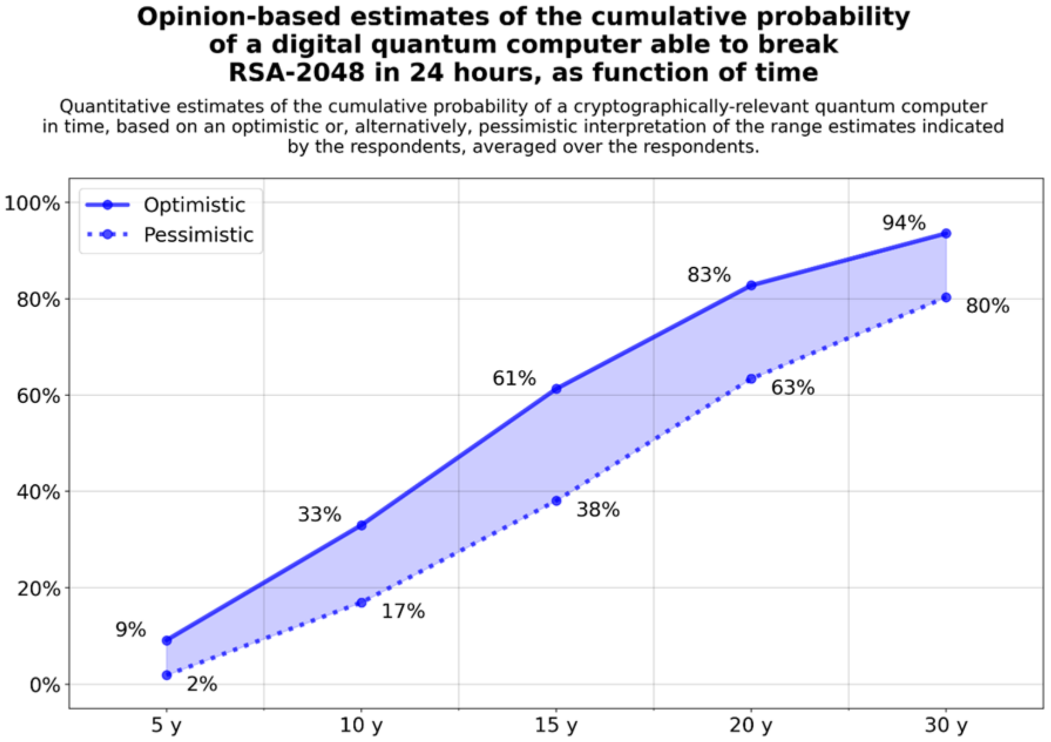

Figure 15 One way of reducing the set of likelihood estimates provided by the experts to some aggregate likelihood is that of interpreting optimistically or, alternatively, pessimistically, the answers of each respondent, and averaging over the respondents. This approach provides a reasonable range for what could have been the average of the point estimates of the experts, had they been asked to provide one single probability.

[7] Note that each expert could have preferred to provide a range rather than a point estimate, if given the opportunity.

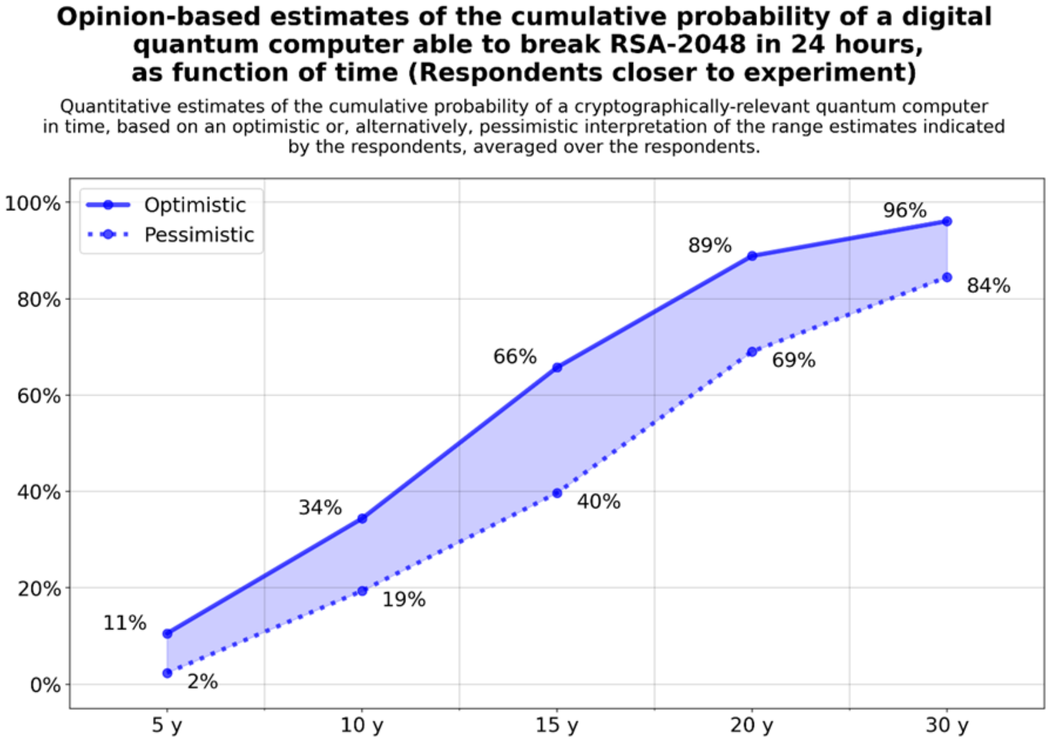

Figure 16 Same graph of cumulative probability of a cryptographically significant quantum computer as in Figure 15, but limited to the opinions expressed by respondents who are closer to experiment. The estimates are in line with those of the larger group of experts. There is a slightly faster growth of the likelihood at 15y and 20y.

two curves provide what we may consider a reasonable notion of uncertainty about the average likelihood assigned by the experts. The result is presented in Figure 15. More details on the method are given in the Appendix.

Figure 15 is as a coarse-grained summary of the experts' opinions based on the analysis we described, and it should be interpreted cautiously. On the other hand, we think it provides precious insight into the quantum threat timeline. For example, even in a 'pessimistic' interpretation of responses as the lowest compatible probability for a given likelihood option, the probability associated by the above-described analysis to the disruptive quantum threat is already ~17% in the next 10 years, growing steadily in the years that follow: the estimated probability is more than ~38% by the 15-year mark, and more than ~63% by the 20-year mark. In Figure 16 we have also plotted the cumulative probability distribution obtained by averaging the probabilities associated with the opinions of only the respondents close to experiments; it does not differ significantly from the all-respondent one.

The above two approaches—based on the average sentiment and the average likelihood—are meant to facilitate reasonable summary interpretations of the opinions we have collected, at the cost of losing some of the details presented in, e.g., Figure 9. One advantage is that, once the results of the survey are coarse-grained in such a fashion, a straightforward comparison with the results of the 2019 and 2020 surveys, which is important for understanding how the opinions of the experts may have changed from survey to survey, becomes feasible.

5.2.1 Comparison with previous years

The starting point of our comparison with the results of the surveys of the previous years is a simple joint rendition of the results for all the three years, presented in Figure 17. A direct inspection reveals a general trend towards larger likelihoods being assigned earlier, with the most evident increases of the likelihood estimates starting at the 15-year mark.

The kind of analysis discussed in the previous section for the 2021 data allows us to make such a trend more directly evident, by comparing the change of the average sentiment expressed by the respondents (Figure 18, previously illustrated for 2021 alone in Figure 14), or in the cumulative probability

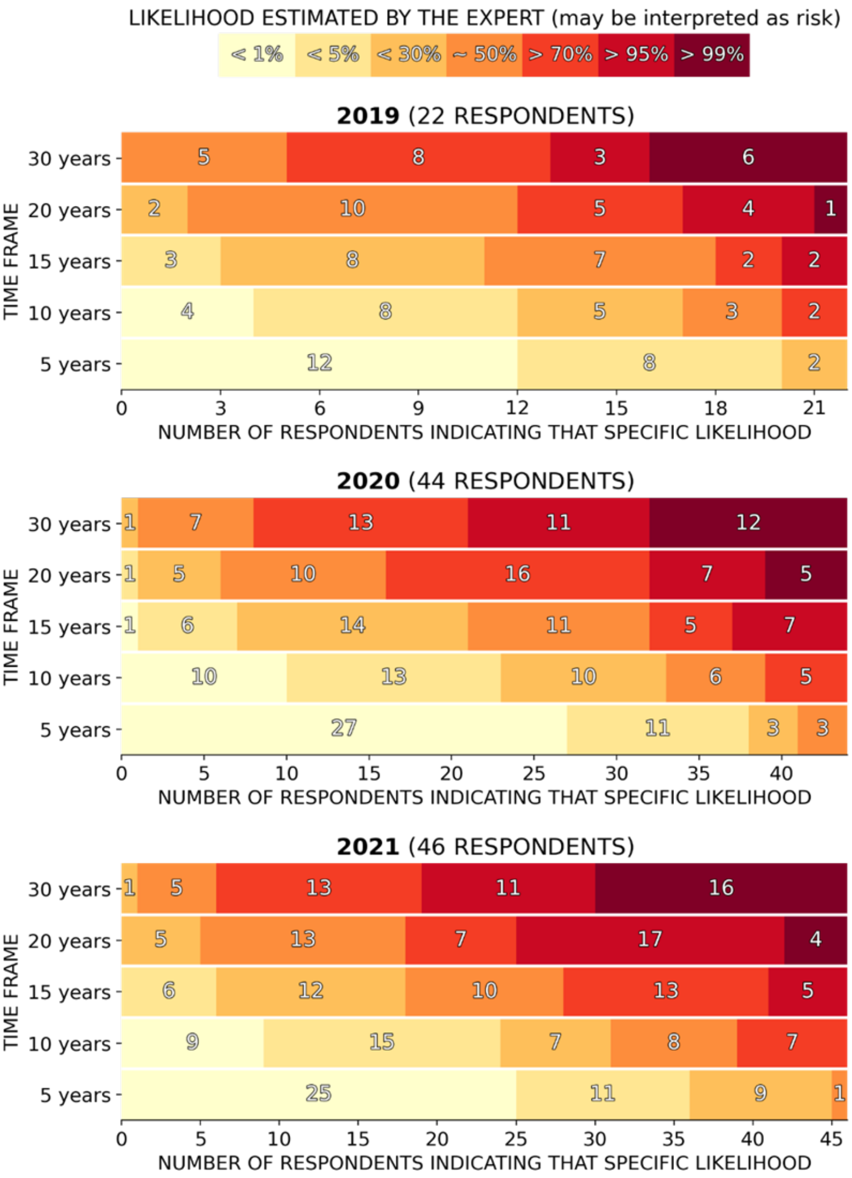

Figure 17 Comparison of the distribution of the opinions of experts about the likelihood of a quantum computer able to break RSA-2048 in 24 hours across the three years of the survey. A trend to estimate a higher chance within a shorter time frame may be appreciated upon inspection. See main text and other graphs in this section that highlight such a trend.

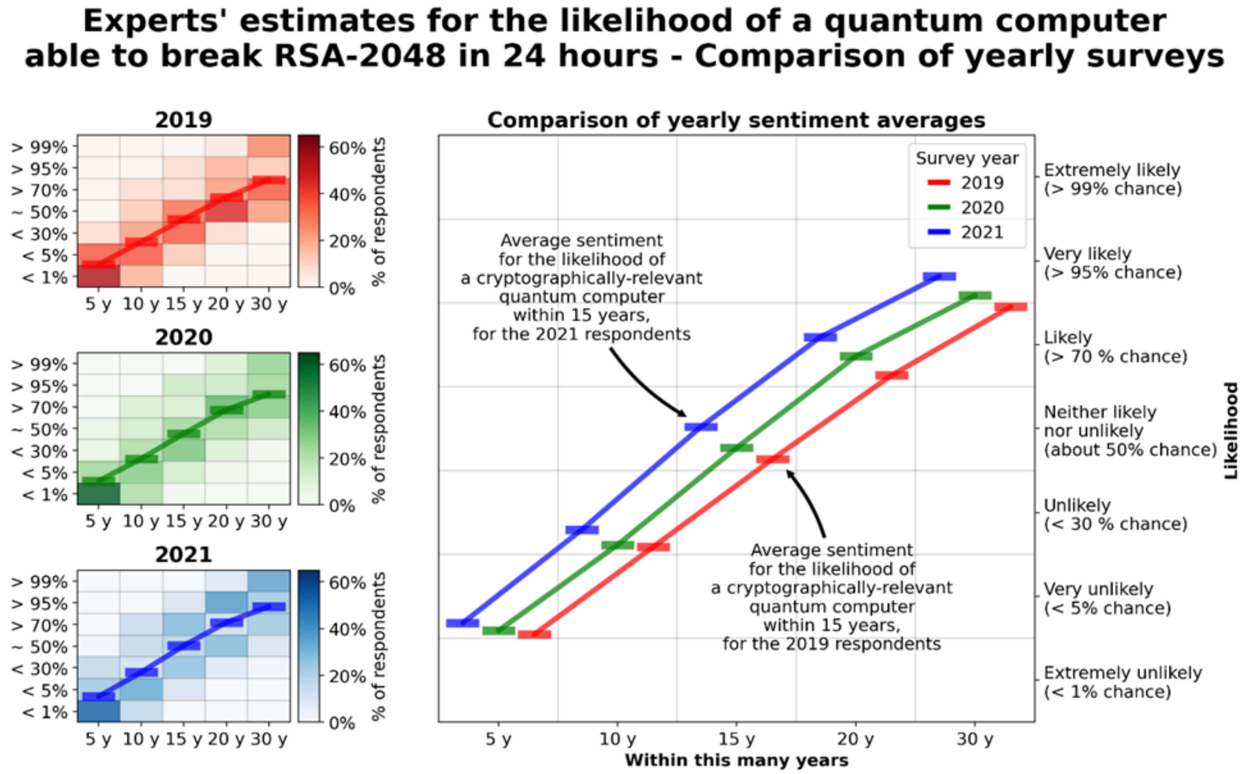

Experts' estimates for the likelihood of a quantum computer able to break RSA-2048 in 24 hours - Comparison of yearly surveys

2019

2020

2021

Comparison of yearly sentiment averages

Figure 18 Evolution of the likelihood estimates by the experts in the surveys about the quantum threat timeline conducted so far. In the three graphs on the left: heatmaps / distributions of responses for the 2019, 2020, and 2021 surveys, with indication of the increase of the average sentiment as one considers longer and longer timeframes (see Figure 14 for details for 2021, and the main text of Section 5.2 for a an explanation of what is being plotted). Large graph on the right: comparison of the trajectory of the average sentiment for the three years, showing a more optimistic sentiment year after year for every timeframe considered in our surveys. The upward shifts in 2021 are larger than the shifts from 2019 to 2020; the larger increase is at the 15-year mark.

distribution based on the “optimistic” or, alternatively, “pessimistic” interpretation of the experts’ likelihood estimates (Figure 19, previously illustrated for 2021 alone in Figure 15).

Both these approaches reveal what could be considered an increasing optimism about the realization of a fault-tolerant quantum computer that is cryptographically relevant. Upward shifts for both the average sentiment and for the range of the average estimated likelihood are consistent across the board for all timeframes. We notice that the upward shifts from 2020 to 2021 are more pronounced than from 2019 to 2020.

We have already indicated how Figure 15, representing the increasing likelihood in time according to the 2021 respondents, provides insight into the quantum threat timeline but should be interpreted cautiously. Caution is also advisable when comparing the change of the likelihood estimates from survey to survey, as, for example, the composition of the pool of respondents has changed. To mitigate this confounding factor, in Figure 20 and Figure 21, we present the survey-to-survey comparison only for the subset of 21 respondents who so far have taken part in all three surveys.

Focusing on such a “stable subset” of experts changes what appears to be the consensus for the short-to-medium term but preserves the increasing trend for the sentiment/likelihood estimates for the

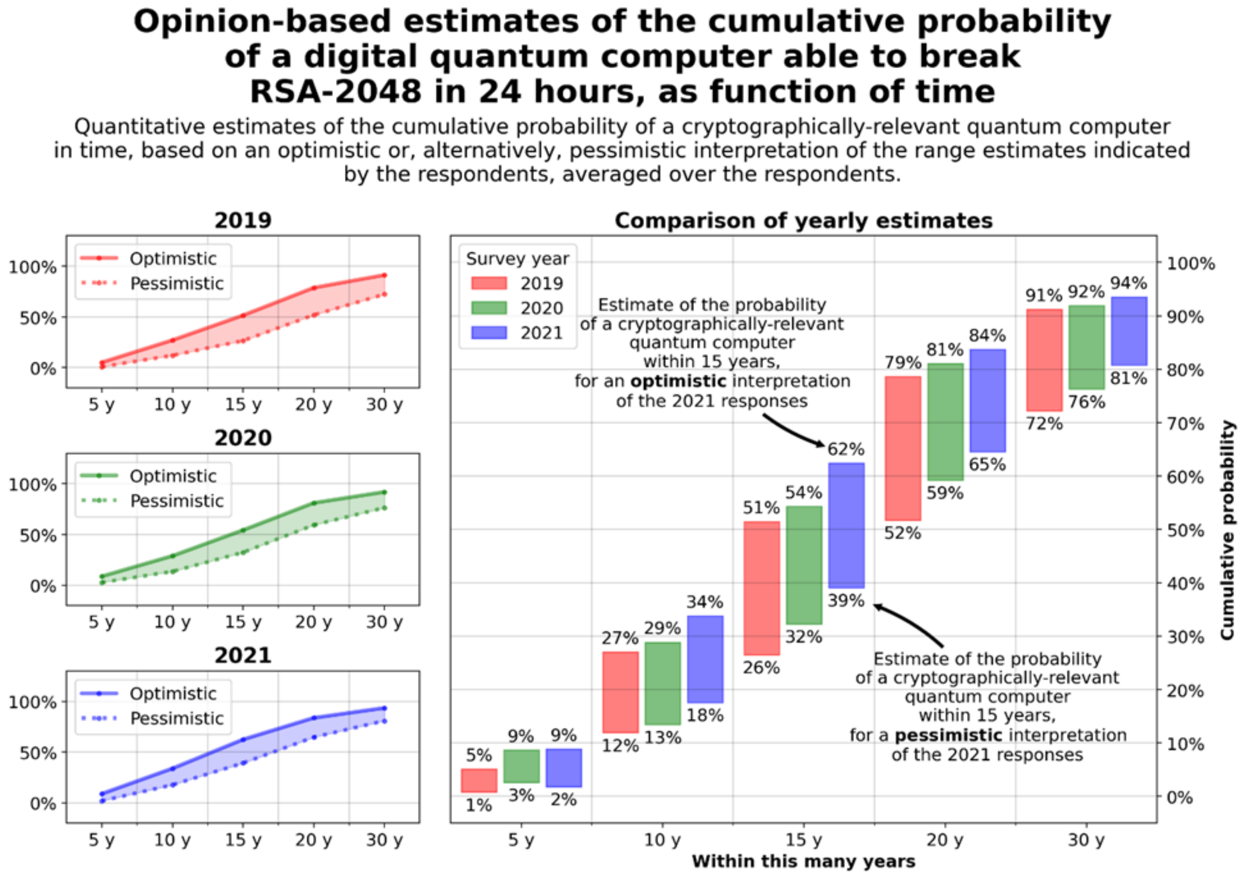

Opinion-based estimates of the cumulative probability of a digital quantum computer able to break RSA-2048 in 24 hours, as function of time

Quantitative estimates of the cumulative probability of a cryptographically-relevant quantum computer in time, based on an optimistic or, alternatively, pessimistic interpretation of the range estimates indicated by the respondents, averaged over the respondents.

2019

2020

2021

Comparison of yearly estimates

Estimate of the probability of a cryptographically-relevant quantum computer within 15 years, for an **optimistic** interpretation of the 2021 responses

Estimate of the probability of a cryptographically-relevant quantum computer within 15 years, for a **pessimistic** interpretation of the 2021 responses

Figure 19 Evolution of the likelihood estimates by the experts in surveys about the quantum threat timeline conducted so far. In the three graphs on the left: probability estimates based on the optimistic or, alternatively, pessimistic interpretation of the responses for the 2019, 2020, and 2021 surveys (see Figure 15 for details for 2021, and the main text of Section 5.2 for a an explanation of what is being plotted). Large graph on the right: side by side and timeframe by timeframe comparison of such estimates. Both the lower and the upper end of the average likelihood estimate have been rising survey after survey, for each timeframe considered, the only exception being the lower end of the 5-year estimate. The increases at the 15-year and 20-year marks appear stronger in the most recent survey.

medium-to-long term. More specifically, the average sentiment is pretty much constant (and relatively low) survey-to-survey for the 5-year timeframe but getting *lower*—albeit not by much—at the 10-year mark. Similarly, for the stable subset of experts the cumulative probability for a cryptographically relevant fault-tolerant quantum computer appears to be shifting down at the 10-year mark from survey to survey. On the other hand, both average sentiment and average likelihood estimates stay high and growing at the 15-year mark and beyond for this subset of respondents; in some cases, they even reach values that are higher than for the whole 2021 respondent cohort.

In Appendix A.5 we provide some considerations about potential factors that one may want to keep in mind when determining the reliability of the experts' estimates, and which may explain some of the just noticed variation between the results for the whole pool of respondents and for the "stable subset".

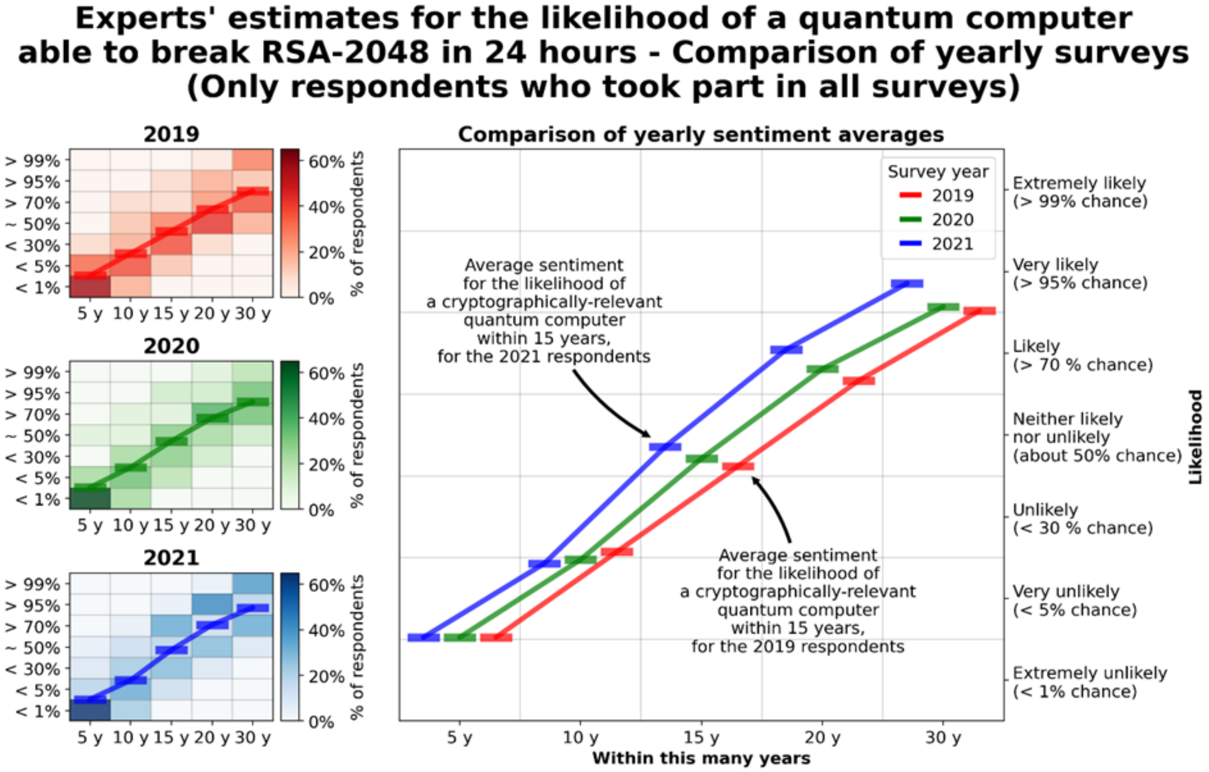

Figure 20 Evolution of the likelihood estimates by the experts in surveys about the quantum threat timeline conducted so far, limited to only those respondents who have taken part in the whole series of surveys (*21 respondents). See caption of Figure 18 for details on what is represented. Contrary to what shown there, the average sentiment does not grow for every considered timeframe; it is instead seen reducing at the 10y mark, before growing with respect to the previous surveys at the 15y timeframe and beyond.*

5.2.2 Comments by respondents

Some of the experts provided comments that add insight to the summary quantitative results just presented, including addressing the point of how much the time constraint we considered in the question ("[..] a quantum computer able to factorize a 2048-bit number *in less than 24 hours* [..]") influenced their estimate. More specifically, with respect to the latter point, and having been prompted to do so by a follow-up question in our questionnaire, some experts commented on how their answer would have changed if the constraint had rather been "*in less than one month*". While some respondents write that in such a case their answers would have indeed indicated higher likelihoods for earlier time frames, other respondents point out that, all considered, the difference in the required "speed" has limited impact.

Klaus Moelmer states that his answers would have shifted

The 10- and 15-years prospects would become more likely.

He is echoed by at least three respondents who chose to stay anonymous, one of whom wrote:

Increasing the target computing time to one month [would] change my estimates - I believe within 10 years it will become likely, and within 15 it will become very likely.

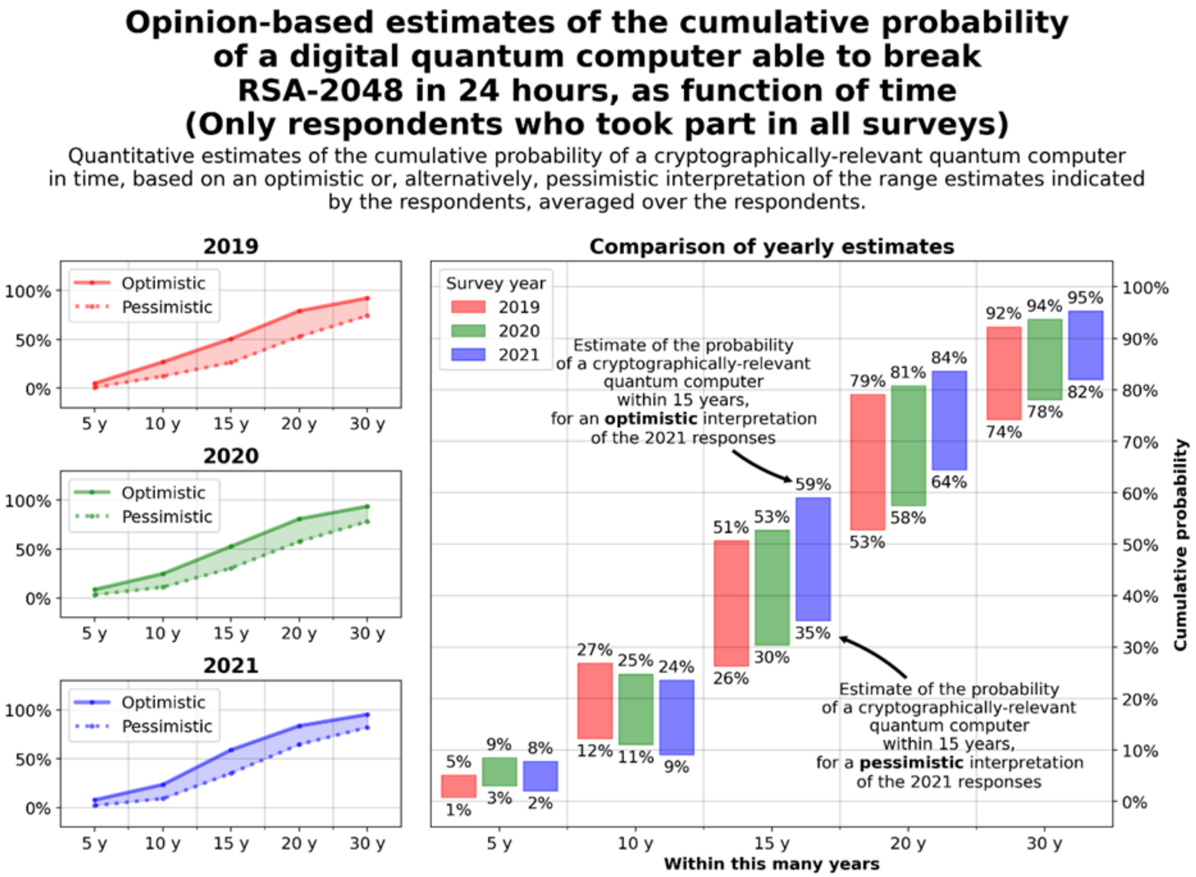

Figure 21 Evolution of the likelihood estimates by the experts in the surveys about the quantum threat timeline conducted so far, limited to only those respondents who have taken part in the whole series of surveys *(21 respondents). See caption of Figure 23 for details on what is represented. Contrary to what shown there for the overall average estimates of all our respondents, for the stable set of respondents the average likelihood-estimate interval does* not *shift upward for every considered timeframe; it is instead seen shifting down at the 10y mark, before nonetheless growing compared to the previous surveys at the 15y timeframe and beyond.*

Kae Nemoto is uncertain about the impact of a change of roughly a factor ten in the time requirement, but sees such a constraint as relevant in general:

> *One month is only one order of magnitude difference, and I am not sure if it is enough. [I]f the time constraint is loosened by two orders of magnitude, the estimation for [the] long term could be rather significantly affected.*

Ashley Montanaro explains why he instead does not think the time allowed for the factorization task is such a key parameter, at least for the considered change in the constraint:

> *My estimate [..] would not change significantly if the computing time allowed increased or decreased somewhat - in my view the biggest challenge lies in reaching the scale of quantum computing hardware required to obtain a relatively large number of high-quality logical qubits; once this is reached, I anticipate that faster and larger hardware will be produced on quite a short timescale.*

Stephanie Simmons agrees and stresses that ultimately it is a question about the solution of the fundamental technological issue of scalability; she thinks that finding such a solution would transform the problem of "speed" into a problem of resources at disposal:

> *The distinction between one month and one day does not really influence my answer to this question. The basic question being asked is "when will quantum scalability be solved" – because when it is, a ["one day" vs "thirty days"] question will likely become a resources/financial question[,] not a technology question, and the resources will absolutely be available at that point.*

Daniel Gottesman also agrees that the difference between one day and one month is not that significant, and explicitly points to parallelization as a rationale:

> *If it can be done in one month, I doubt it will be much longer before it can be done in 24 hours. By then, all the scaling problems will have been largely dealt with, so in the worst case, all that is needed is a bit more parallelization.*

Joe Fitzsimons supports this view by explicitly pointing to the large gap that exists between present quantum processors and the discussed future ones:

> *The difference between a quantum computer that can factor RSA-2048 in one day versus in month is very small compared to the gap between today's quantum processors and ones capable of factoring a 2048-bit number.*

A respondent explains the difficulty of the kind of estimate the experts were asked to make:

> *[I]t is very hard to predict the future: [my] answer represents a current best guess, based in part on current roadmaps from industry leaders such as Google and IBM. We still have a considerable road to travel, and many obstacles to overcome. There are several different avenues to achieving the stated objective. This makes it hard to make predictions at this point in time.*

The respondent chooses to be cautious in two separate senses. First, about the estimate itself:

> *I think it very unlikely that the stated objective (i.e., a quantum computer able to factor a 2048-bit number in less than 24 hours) will be reached before the end of the decade, i.e., in approximately ten year's time. In 20 years' time, I expect that the objective will either have been reached, or that we will have identified and understood one or more key obstacle preventing it from being reached. This explains why I will not go above a 50% chance in the above estimates.*

Second, the expert warns about the concept of risk, and of its handling, in a cryptographic context, as even a "small" likelihood estimate could mean an unacceptable risk:

> *[I]n the context of cryptography, already a 1% risk of currently widely deployed asymmetric cryptography being broken is unacceptably high, requiring some form of mitigating action.*

Considering this, the expert points out that[8]

[8] We find this kind of comment very relevant both for risk managers and for the future design of our survey.

> *[O]ne could discuss what scale is best to use here, with which ticks, and what labels should be associated with each tick, and so forth.*

Within the recognized uncertainty, **Dave Bacon** is optimistic, specifically about progress that we cannot quite anticipate:

> *[I]n 10 years I suspect that we will have both algorithmic and physical breakthroughs which make [building a cryptographically-relevant quantum computer] less challenging, resource-wise, than we currently believe.*

One respondent supports the idea that there is some inherent unpredictability related to potential (near-)future breakthroughs, also providing a fitting historical precedent:

> *With noise levels as they are now, it is extremely hard to extrapolate forward with much confidence. Great leaps in new technology are usually the result of paradigm shifts and fundamental breakthroughs (e.g., transistors) rather than steady improvement in the same direction (e.g., improving vacuum tubes). These breakthroughs usually involve fundamental shifts in perspective, which makes them notoriously hard to predict. Tasks that look nearly impossible now may, with new insight or new experimental breakthroughs, suddenly become possible – with little or no warning.*

Andrea Morello appears to be already quite satisfied with the current rate of progress, which he sees as likely to continue; he instead chooses to be cautious about potential yet undiscovered "show-stoppers" related to the behaviour of large quantum systems:

> *I have set an upper bound of 95% to the likelihood of success (no matter how far in the future) to account for possible fundamental surprises in our understanding of large quantum systems. Other than that, progress will continue at great pace.*

Alexandre Blais explains that what drove his estimates is mostly the consideration of how many underlying physical qubits will be needed, which in turn can be cast as a non-trivial question about the efficiency that can be achieved by quantum error correction. He further suggests that such an efficiency could be boosted by a stronger interplay between the development of quantum error correction (QEC) methods and of hardware:

> *My answers [..] are more governed by the number of physical qubits that are necessary rather than the overall computing time. The gate times are unlikely to change dramatically so the change in computing time from 24 hours to one month is, at least to me, related to a change in how efficient error correction can be made. This is a difficult question. To move that needle, it is my impression that we need more QEC research that is "hardware aware" and more quantum hardware research that is more "QEC aware".*

Bill Coish briefly discusses explicit estimates for the required number of physical qubits, how different platforms may be better or worse suited for the task (see also Section 5.1), and how the control of a very large number of qubits remains a challenge for any implementation in the short-to-medium term, hence keeping the likelihood of a cryptographically-relevant quantum computer relatively low:

> *Scaling to >10 million physical qubits with a cycle time <1 microsecond will almost certainly be necessary. I don't see these system sizes as reasonable for transmon-style[9] superconducting qubits (due to the size of the resonator) and the cycle time is very difficult to achieve for ion-trap qubits. Spins in semiconductor quantum dots could achieve the required cycle time and are small enough to imagine up to ~100 million qubits in a very small volume, but controlling so many qubits is still a huge engineering task that will likely require more than a decade to solve in any implementation.*

Another respondent agrees that breaking RSA-2048 is and will remain a huge challenge, and stresses the importance of intermediate applications in order to secure the necessary level of long-term investment:

> *Building a quantum computer that can factor such a large integer remains a very very challenging problem and will only happen if other applications are developed sooner to continue the technical investment.*

John Martinis suggests that recent progress in China and considerable focused efforts there could lead to a shorter quantum threat timeline:

> *Companies like Google are talking about 1M qubits in 10 years. You need about 30M to do factoring, so it will take a bit longer. However, you have [to consider] that China has now caught up to Google, so it could be faster since they are way more motivated to do this quickly than anyone else.*

[9] In this report, we have considered very coarse-grained classes of physical implementations; different superconducting implementations of qubits exist, including the mentioned one.

5.3 Logical qubits and fault-tolerant schemes

Arguably, the next major milestone towards building a fault-tolerant quantum computer is the realization of a controllable logical qubit (see also Section 5.6). This would mean having achieved the ability to prepare, store, and manipulate the quantum information of a single “ideal” qubit for some large—potentially arbitrary large, at least in principle—sequence of operations. This feat would be attained by encoding the logical qubit in a sufficiently large number of physical qubits through quantum error correction, with a physical error rate low enough to reach fault-tolerance.

Significant demonstrations of various aspects of error correction—e.g., reduction of at least some types of error rates via encoding, the repeated read-out of physical errors, or running the correction in real time—have already been achieved experimentally, albeit not necessarily in the same experiment, in at least some architectures, like superconducting circuits (see, e.g., (Chen et al, 2021), (Ristè et al, 2020), (Andersen et al, 2020)) and ion traps (see, e.g., (Egan, 2021), (Ryan-Anderson et al, 2021)).

The importance of these recent results is stressed by our respondents in Section 5.4.2, but significant work remains to be done, like realizing *all* the relevant aspects of error correction in the same system/experiment, and, very importantly, doing so in a way that is amenable to scaling—like asked in our question.

Nonetheless, the progress so far has been sufficient to induce the majority of experts who provided an estimate (44 in total) to suggest that the realization of a scalable logical qubit is quite close. Thirteen out of the 44 respondents indicated that it will be demonstrated with “about 50%” probability or higher within one year. Even more strikingly, most respondents (31/44) suggested that this will happen with “about 50%” or more probability within three years (see Figure 22 and Figure 23).

A selection of comments is provided in Appendix A.5 It is important to remark that some experts have expressed the perspective that the realization of an *individual* logical qubit that is scalable is not necessarily a well defined or sensible milestone. Some of the reasons provided go from the opinion that focusing on the realization of an individual logical qubit—with, say, the idea/intuition that it might be possible to combine many instances of it afterwards—does not capture well how quantum computing implementations are intended to work, to the opinion that claims of scalability are relatively vacuous until that scaling is realized.

Some experts have interpreted the question in some specific way that made the question the most meaningful to them and have indicated such an interpretation explicitly. On one hand, this may be considered a warning about:

- a coarse-grained analysis of the responses, as not all answers may be based on the same assumptions;
- the close scrutiny future claims of results towards fault-tolerance may be subject to.

On the other hand, the clarifying comments of the experts reveal:

- the complexity of establishing sensible ways of evaluating (claims of) demonstrations of error correction and fault-tolerance;
- the challenge of realizing a fault-tolerant quantum computer.

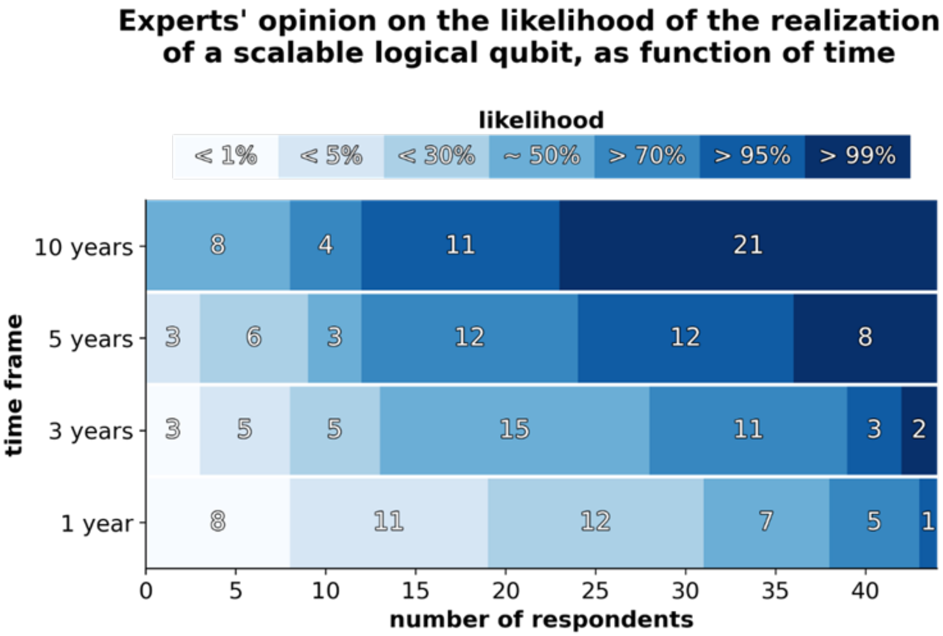

Figure 23 Quantum information is fragile, and its manipulation imperfect. Nonetheless, the experts appear to be of the general opinion that we will soon see the realization of logical qubits which make use of error correction to counteract such issues. Most importantly this appears to be likely even considering the requirement of scalability *of the encoding scheme, that is the possibility of realizing and handling a growing number of such logical qubits through a manageable increase in resources and complexity of operations.*

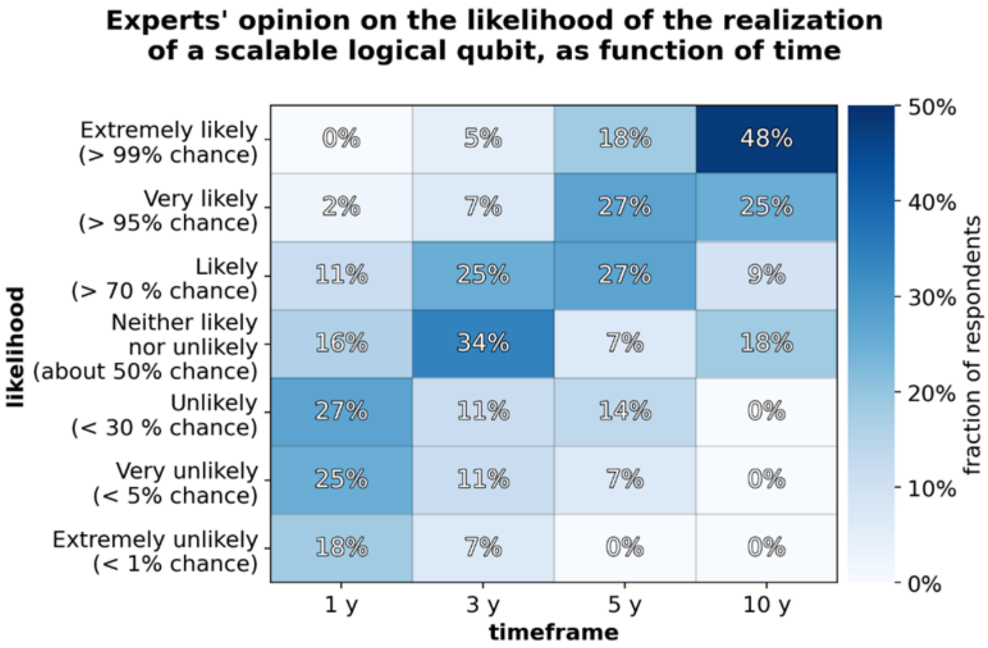

Figure 23 Heatmap representation of the distribution of answers presented in Figure 22. It is possible to appreciate the fast increase of the estimated likelihood over time frames that are relatively close in the future.

5.3.1 Most promising fault-tolerant schemes

The notion of redundancy makes the possibility of error-correction for classical information relatively intuitive, but quantum mechanical properties prevent the use of "naïve" redundancy, because quantum information cannot in general be copied (Nielsen & Chuang, 2002). In this sense, the conceptual possibility of quantum fault-tolerance is remarkable per se, independently of its actual realization. Nonetheless, when it comes to the latter, it is important that the implementation requirements be feasible, for example that the threshold for error rates be within experimental reach, or that the complexity of the control of the underlying physical qubits be compatible with the large number of physical qubits needed to encode logical qubits. In turn, these key aspects depend on the interplay between the properties of the physical platform used for the implementation, and on the details of the fault-tolerant scheme, that is, on the "recipe" to combine physical qubits into logical qubits.

We have asked the experts to share their opinion on the most promising fault-tolerant schemes. Many respondents point to the surface code—and similar/associated schemes, see Appendix—in superconducting implementations as leading proposal. Nonetheless several other respondents indicate promising alternatives, which may improve the rate at which quantum information can be reliably encoded and manipulated, intended as the ratio between the number of encoded logical qubits and underlying physical qubits. 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 favor physical systems other than superconducting qubits. Other proposals that are attracting substantial interest 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. Finally, there is interest in tailoring error correction to the specific kind of noise that affects a certain implementation.

In general, all the above can roughly be seen as attempts at making the best possible use of the freedom in the encoding of quantum information and of the specific properties of the physical systems used to encode it, including the specific noise, with the goal of attaining a robust and efficient encoding.

Some quotes from the experts can be found in Appendix A.5 .

5.4 Societal and funding factors

In this section we report the results of the questions directed to assess societal and funding factors that may impact the timeline of the development of a cryptographically-relevant fault-tolerant quantum computer

5.4.1 Level of funding of quantum computing research

The present level of investments in quantum technologies, and specifically in quantum computing, is at a historical high. This is very relevant, because 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 / consulting for industry, and at the head of start-ups, our respondents have a significant vantage point to estimate the evolution of funding. In 2020, we asked them to forecast what was likely to happen in the following two years, and we have repeated the question this year[10].

The results of the survey are presented in Figure 24. A large majority of the respondents expect investments towards quantum computing to increase or even significantly increase. Compared to last year, the percentage of respondents who think the increase will be substantial has gone down, but this appears to be more than compensated by the facts that:

- we have already seen substantial increases in investments in the past years, and
- the percentage of those who foresee at least an increase has grown, with the percentages of those who forecast stable or decreasing funding going down, even to zero, for the latter option: compared to last year, no expert indicates any expectation of a decrease.

The perception about the dynamics of investments in the field varies, as evident also from the quotes presented in the Appendix. Some respondents think that, while there might be a growth, we may be close to peak, while others see the growth as unabated and solidifying. Nonetheless there seems to be some consensus that a moment will come relatively soon when there will be a consolidation of the start-up landscape. How that process will take place and will be perceived appears to be important, as a "collapse" of one or more start-ups may lead to a slow-down of investments in the whole field. Reasons provided to expect continued or growing funding comprise:

- momentum following large recent investments in the field, with various entities interested in having a stake in a promising and growing field;
- continuous government spending, which may increase depending on the location (several respondents mentioned China), or seen as at least staying stable if that location had already seen large investments in the recent past;
- related to the previous point: world competition in a technological sector considered as strategic, fuelled also by political and economic tensions;
- investors are not expecting yet profitable results;

[10] Compared to 2020, in 2021 we adopted a slightly different wording, adding the qualifier "global", to make sure that the respondents considered the level of worldwide funding, rather than specific local realities. We think the nonetheless the direct comparison of the 2020 and 2021 responses is reasonable.

- hype.

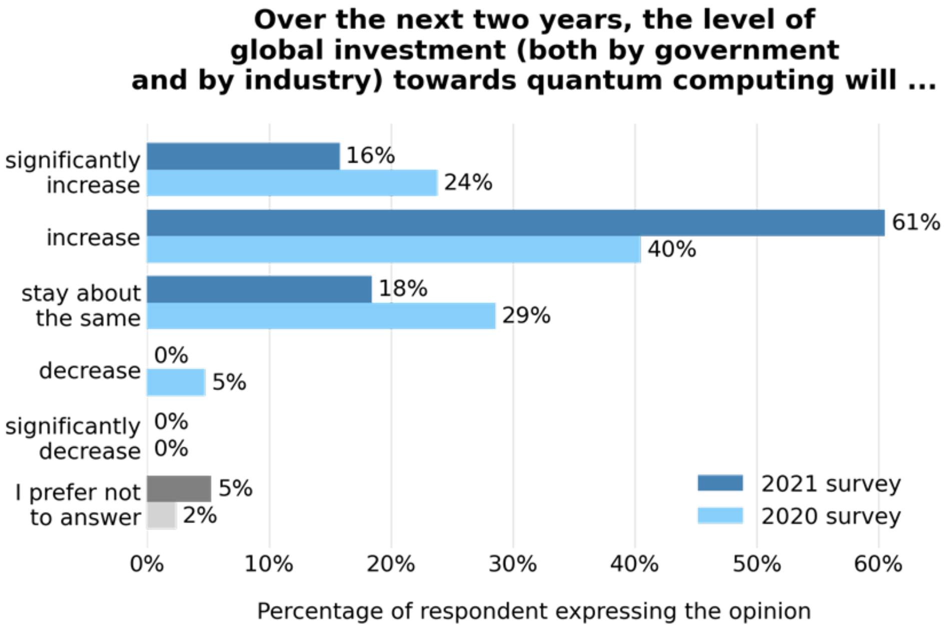

Figure 24 Expected change in the level of investment toward quantum computing in the next two years. While a smaller percentage of the 2021 respondents forecast a significant increase in funding compared to the 2020 respondents, the fraction of those who expect some increase is even larger than last year, and no respondent has forecast any kind of decrease.

We note that hype is nonetheless considered also as dangerous, because it may create excessive expectations in terms of capabilities/results, which may be impossible to satisfy.

In general, some respondents point out that continuous progress is needed to keep seeing investments pour into the field. Importantly, the large investments in the area and the rapid growth of the quantum landscape/ecosystem have made quantum information processing researchers a relatively scarce resource; focus on training is necessary to have qualified personnel that can support such a growth.

5.4.2 Global race to build a fault-tolerant quantum computer

The development of a cryptographically-relevant quantum computer can be seen as a race, at multiple levels. In Section 5.1 we have already discussed a competition between architectures. Here we are interested in the competition at the level of both national and supranational (like the European Union) entities.

The successful development of a quantum computer is explicitly considered a strategic goal by many countries. The reason is that it would be game-changing in many ways, not only for cryptography and for much of the digital infrastructure—the sense most relevant to this report—but also for other societal and economic activities, starting, e.g., from the ability to simulate quantum systems in the design of new advanced materials and drugs.

The resulting competition is a major driver of the investments in the quantum computing area, and understanding how the “race” is going and how it may develop provides insight into the quantum threat timeline itself. Moreover, for risk managers tasked with handling the quantum threat it is important to understand *where* the threat may come from, which means understanding which players could have earliest access to a cryptographically relevant quantum computer.

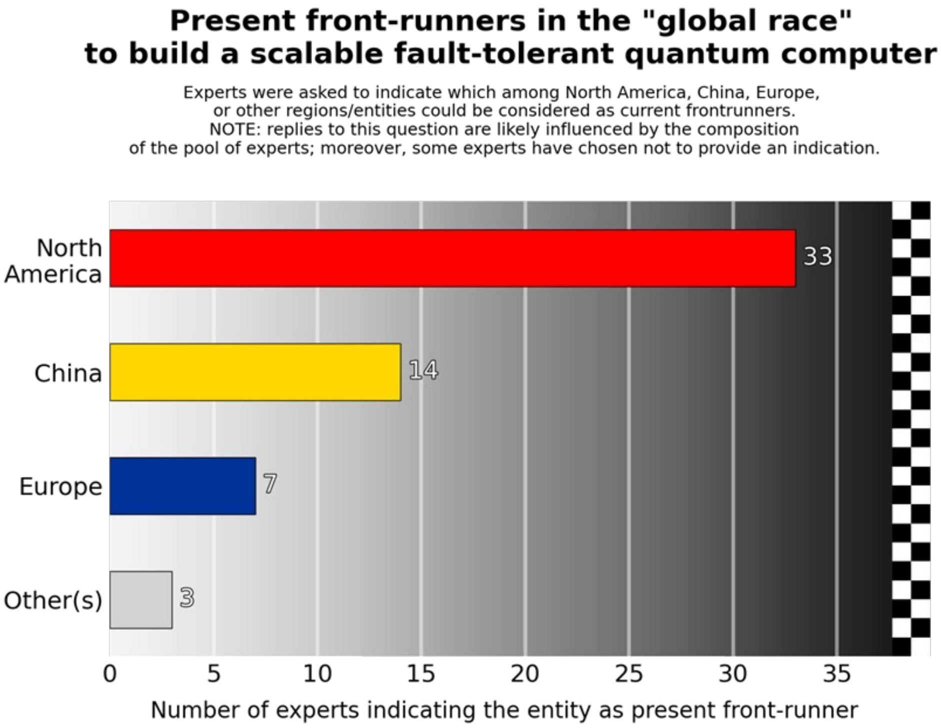

Figure 25 Number of respondents that indicated a region/entity as present front-runner in the global race to build a fault-tolerant quantum computer (multiple answers were allowed). North America appears to be in a strong position, followed by China and then Europe. "Other(s)" options included collaboration among the regions or stressed the notion of "global companies" leading the efforts.

We point out that the reader may take into account the geographical composition of our pool of respondents by referencing Section 4.

We have asked the experts to indicate which geographic areas among China, Europe and North America are current frontrunners, with the option to provide multiple answers and/or alternative names. The results are shown in Figure 25. Not all the experts provided an opinion but, according to those who did, North America appears to be the present leading world region, followed by China and Europe, in this order. Among the "Other(s)" options indicated are the suggestions that what is driving the development are "global" companies, and that the development of a quantum computer will be the result of the interaction / collaboration between geographic regions.

Given our interest in future trends, we also asked the experts to indicate the likelihood, for each region previously considered, to be a frontrurnner five years from now, and whether new frontrunners may emerge. The results are presented in Figure 26. Most respondents consider it likely that North America will maintain its frontrunner position. On the other hand, China scores relatively highly as a likely frontrunner and is considered to have significant potential. Europe appears to lag behind in expectations, and it is worth remarking that many respondents (15/38) consider it unlikely that it will

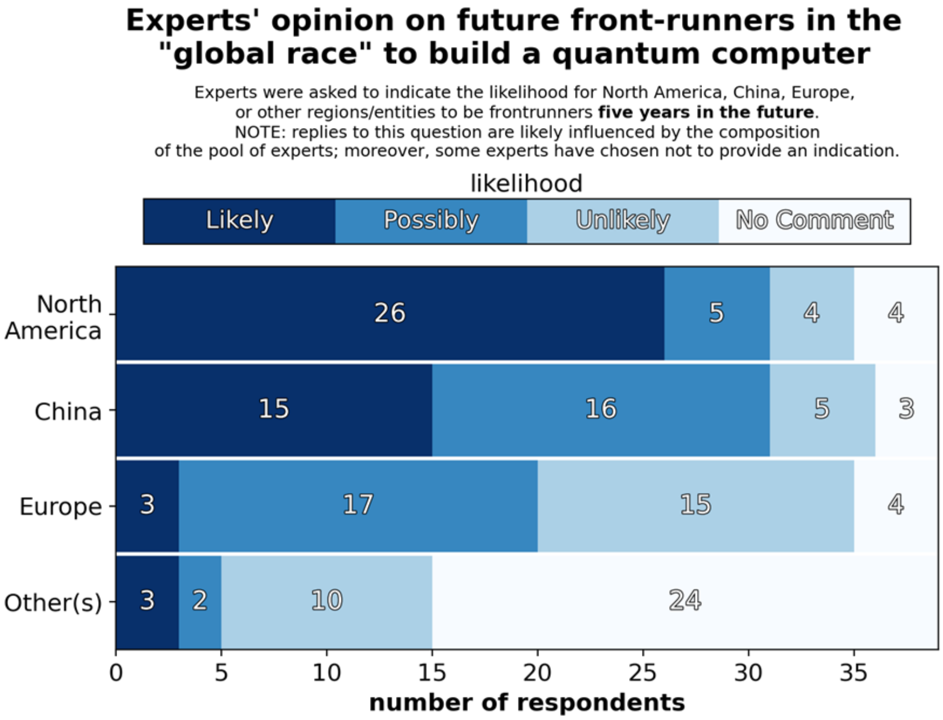

Figure 26 Number of respondents that indicated the likelihood of a given region/entity to be a front-runner in the global race to build a fault-tolerant quantum computer five years from now.

have the status of frontrunner in five years. Australia and Japan were named as "other" countries that are potential future frontrunners.

Some experts provided relevant comments, which offer some rationale for the results of the survey. In particular, some respondents point to the issues of availability of talent, of resources (particularly financial/funding), and of focus/coordination as determinant in influencing the quantum race. See Appendix A.5

5.4.3 How COVID-19 is affecting quantum computing research

In the 2020 report we asked the respondents to comment on how the ongoing COVID-19 pandemic was affecting quantum computing research, and they expressed various degrees of concern (Mosca & Piani, 2021).

An important difference was highlighted overall between experimental and theoretical work, with the first kind being affected more due to several key factors, going from restrictions on accessing equipment on campus to disrupted supply chains.

This year we asked the respondent to be quantitative, by posing the following question:

Q: Assuming that the COVID-19 pandemic has had / will ultimately have impacted negatively the progress being made towards achieving the construction of a scalable fault-tolerant quantum computer, how large do you estimate such a slowdown to be?

Figure 27 summarizes the opinions of those experts who provided one answer among the choices provided; the percentages refer to such a subset of respondents. In Figure 28 we report the distribution of answers restricted to the subset of experts who are deemed to be closer to experiments.

The respondents agree that the pandemic has had a substantial impact, with 22% indicating a large delay of between one and two years, and only 14% indicating a delay of less than three months. As one could have expected, the percentage of those who foresee a large delay is higher among the respondents closer to experiment, who have had to face the disruption of work at laboratories and of supply chains (Figure 28).

Shengyu Zhang comments that present estimates about the final slowdown are necessarily speculative, as the pandemic is still ongoing and

> *[t]he answer [..] also depends on the future control of the pandemic (policy, vaccine, drug[s]...).*

Tracy Northup stresses the different impact on theoretical rather than experimental work:

> *I think work in my own academic research group has been set back maybe by about 6 months so far. I've heard estimates that were higher or lower, depending on whether these were experimental or theory groups, and for experiments, whether they were at a stage at which they could be run remotely or whether a lot of hands-on construction was being done.*

She also points out how the impact goes well beyond delaying current work, because the pandemic has been affecting and is continuing to affect scientific exchange, collaboration, and training:

> *I chose 1-2 years both because we are not out of the woods yet, sadly, but also because you have to factor in longer-term effects, like the missed opportunities for in-person scientific exchange at conferences and research visits that didn't happen, or the delays in hiring and training new team members.*

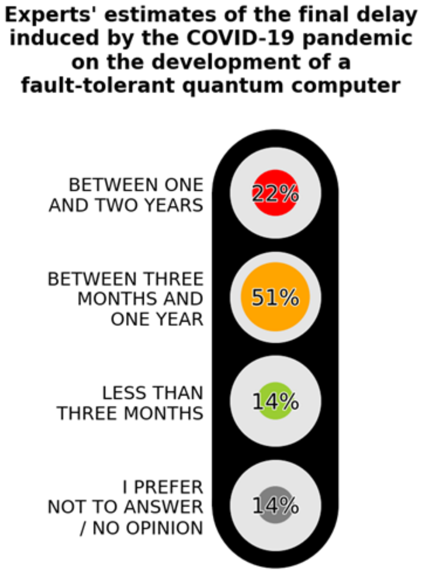

Figure 27 The pandemic has caused a slowdown of activities of all kinds, including quantum computing research. We asked the experts to provide an estimate of the pandemic-induced delay for the development of a fault-tolerant quantum computer.

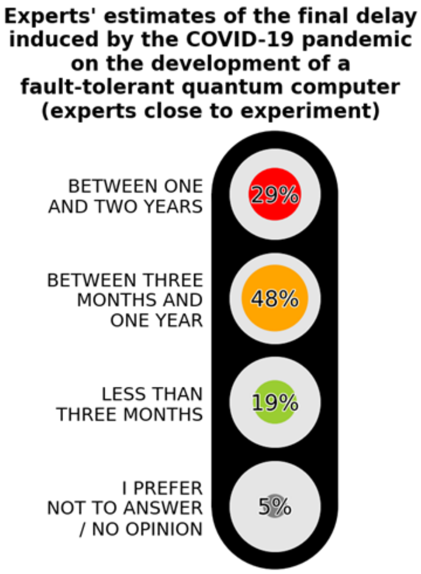

Figure 28 Slowdown estimates for the subset of respondents whom we consider closer to experiments. The percentage of those who estimate a relatively large slowdown is higher than in the overall set of respondents.

We emphasize that the pandemic-induced slowdown may have affected (at least) the short-term estimates for the likelihood of a cryptographically relevant quantum computer (Section 5.2) and of the implementation of a controllable fault-tolerant qubit (Section 5.3).

We refer the reader to the previous installment in this series of reports (Mosca & Piani, 2021) for more comments by the experts about the ways the pandemic has been affecting quantum computing research.

5.5 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. Many mentioned experiments demonstrating various aspects of quantum error correction (see also Section 5.3), or new experiments demonstrating so-called quantum supremacy, with particular emphasis on the fact that they were performed in China. Some highlighted also promising progress in the theory of error-correction and fault-tolerance, potentially able to significantly reduce the number of physical qubits needed to reliably encode and process quantum information.

Here are some representative quotes.

Dave Bacon is among those who point to

> *[t]he first quantum error correcting experiments. In particular[,] the experiments by Monroe and Brown in trapped ions [(Egan, 2021)], Honeywell's experiments (repeated rounds [of error correction (Ryan-Anderson et al, 2021)]), and Google's larger bit and phase flip work [(Chen et al, 2021)].*

Ashley Montanaro is one among several to highlight the most recent realization of quantum supremacy (Wu et al, 2021):

> *A very impressive achievement is the group of Jian-Wei Pan showing that a quantum computer can outperform the world's best supercomputers, reproducing and improving on Google's 2019 result.*

He also points to important steady progress:

> *I'd highlight ongoing improvements in the quality and maintainability of the hardware platforms used by the major quantum hardware companies, combined with their development of long-term roadmaps for achieving fault-tolerant quantum computing.*

He is not the only one providing as an answer a combination of pushing the boundaries of quantum computation research while improving quality and reliability. One respondent wrote:

> *Most obvious answer: China's 56-qubit superconducting device.*
> *Less obvious answer: detailed published designs for reaching 99.99% fidelity for modified superconducting designs.*

Another respondent highlighted other improvements that may not immediately attract attention but may prove very important in the medium to long term:

> *We have [..] seen progress [..] in how signals are routed in and out of dilution refrigerators, and on how to better handle correlated errors and error burst caused by cosmic rays [..]. Such development may likely also prove important when scaling up systems in the future.*

One respondent pointed to the realization of quantum supremacy in a platform other than superconducting circuits:

> *Gaussian boson sampling [(Zhong et al, 2021)] displaying quantum supremacy/advantage – this is the answer from optics to Google's quantum supremacy result in superconducting circuits.*

Simon Benjamin added to the list of achievements seen in the last year the significant progress made in several platforms alternative to the ones considered as the leading ones so far:

> *[S]everal less mature technologies (versus superconducting and ion trap) have achieved high fidelity operations suggesting they are only ~5 years behind.*

Another respondent agrees:

> *The reported high-fidelity gates with Rydberg atoms bring it into the field as a highly scalable and strong candidate for both surface code and Shor/Steane code implementations. A wealth of ideas for multi-qubit gates and natural interfacing by light brings promises for this system far into the future.*

Tracy Northup adds further perspective about this kind of progress:

> *It's not yet clear whether [Rydberg atoms] will be well-suited for digital quantum computing, and [such kind of system is] not yet as advanced as superconducting qubits and trapped ions (for example, in implementing error correction), but the rate of improvement and the degree of control over hundreds of qubits is impressive.*

Both **Daniel Gottesman** and **Stephanie Simmons** point to progress in theoretical quantum error correction, particularly the discovery and rapid development of low-density parity check codes of greater distance than previously known.

Alexandre Blais and **Joe Fitzsimons**, as well as other respondents, mention research results demonstrating how to leverage bias in noise to reduce the overhead in fault-tolerant schemes.

5.6 Next near-term step

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.

Many pointed to progress in the experimental demonstration of error correction, including, among the most demanding desiderata, the realization of error-corrected operations among at least two-qubits.

Here are some opinions that are echoed also by other respondents who are not directly quoted.

Dave Bacon:

> *[A c]ouple of good steps: first [Quantum Error Correction] experiments with gain as you scale the distance*[11]*, first magic state distillation*[12] *that improves fidelity, first encoded two qubit gates.*

Ashley Montanaro:

> *It's plausible that an impressive small-scale demonstration of error suppression via the surface code, or some other error-correction procedure, could be achieved by summer 2022. This is the next key milestone towards fault-tolerant quantum computing.*

Simon Benjamin emphasizes the issue of scalability in his answer:

> *A demonstration like [the exponential suppression of errors of (Chen et al, 2021)], but with the full surface code (or other e.g., colour code) and with two logical qubits; importantly, not using any support elements that are non-scalable.*

Along similar lines, **Andrea Morello** would like to see *"quantum operations between two adjacent logical qubits"* and **Tracy Northup** the "*fault-tolerant error correction of multiple logical qubits*".

One respondent points to scalability as a central issue, and would like to see the

> *[d]evelopment and engineering of quantum computer hardware that scales, not necessarily focusing on a fault-tolerant logical qubit.*

Another respondent points to further strengthened demonstrations of quantum supremacy / of a quantum advantage, and the realization of high-fidelity (physical) multi-qubit modules:

> *1) [R]unning a supremacy algorithm on a significantly larger processor, e.g., 300 qubits.*
>
> *2) Demonstration in a scientific paper of a 4-qubit module where all gate operations are at better than 99.99% fidelity.*

One respondent wishes to see the continuous variable approach prove itself with respect to the progress towards fault tolerance:

[11] See Introduction, in particular Section 1.3.2.

[12] "Magic states" are quantum states that have the role of resource in certain error-correction scheme. "Distillation" refers to a procedure where imperfect resources can be transformed into improved resources (in this case, in terms of "fidelity") at the cost of reducing the number of resources.

> *I believe that the next critical step is demonstrating a GKP state ("grid state") in optical hardware. I think this is challenging but achievable in the next year.*

Alexandre Blais judges that the demonstration of a single logical qubit remains the essential next step. In that respect, **Daniel Gottesman** would be pleased with

> *[d]emonstrating a not-necessarily universal fault-tolerant logical qubit with error rates convincingly below the unencoded error rates.*

One respondent believes that

> *[..] the most significant progress would be in understanding how to put together several systems of less than 1000 qubits each so that they can function efficiently as a single quantum computer.*

Kae Nemoto casts a similar goal in terms of an

> *[..] interface to quantum mechanically connect two qubits on different chips in a scalable manner.*

Frank Wilhelm-Mauch points to the importance of *"understanding the role of global errors triggered by ionizing radiation"*, where he refers to the issue that in reality the errors experiences by physical qubits may not be independent, one such case being that of errors induced by cosmic rays.

Stephanie Simmons provides an opinion looking at the issue from several angles:

> *[A] demonstration of a non-fault-tolerant (NISQ) quantum algorithm of appreciable commercial value would be the most important leap forward; it would spur so much investment that the overall goal of quantum fault-tolerance would come much closer.*
>
> *Further progress on reducing overheads by the continued development of better fault-tolerant codes would be a very significant and achievable objective in that timeframe.*
>
> *From a hardware perspective, each major platform has substantial milestones that could be met within the next calendar year which could unlock more resources. By way of example, the arrival of neutral atoms as a controllable quantum platform at the scale of hundreds of qubits will be a major event if firms [developing such king of platform] are successful in their stated 2022 goals.*

5.7 Next milestone

We strive to identify relevant milestones that can be considered as highly significative in the development of a fully-fledged and cryptographically relevant digital quantum computer. This is the kind of milestone that, for example, the realization of a fault-tolerant scalable qubit may constitute and whose timeframe we may ask the respondents to assess (see Section 5.3).

We asked for the input of the respondents to identify what kind of milestone may be an intermediate one between the realization of a fault-tolerant scalable qubit (see Section 5.3) and the realization of a cryptographically-relevant quantum computer (see Section 5.2). This input will inform the next installments in this series of reports. It also has value in itself *now* when it comes to understanding how the experts think the development of quantum computers will unravel.

5.8 Other notable remarks by participants

We asked the respondents to tell us about “the status of [their] own research” and to “comment freely on the present and near-future status of development of quantum computers”. We report here a selection of their replies and comments. We attribute quotes for those respondents that have given us permission to do so.

Some themes that appear repeatedly are:

- the progress and the excitement that permeates the field, which some see as sparking more discussions and interactions between experts in various subfields, or between experimentalists and theorists[13];
- the remarkable and productive—but sometimes problematic, especially in terms of attitudes towards illustrating achieved or expected results—interaction between academia and the private sector;
- also related to the previous point: the dangers of hype and of high (and potentially, too short-term) expectations from funders, government and the public;
- the hope that quantum computing will make a difference in several areas of science, thanks to the ability to simulate quantum systems.

Dave Bacon

> *There is still tremendous good work going on, despite many grumblings about hype. Quantum machine learning is undergoing a cycle of discovery [where] we are seeing a [wave of] more rigorous results [compared to previous efforts]. We are entering the first experiments for quantum error correction, when Shor's second great discovery will finally be realized. In industry, those who are able to focus on results, will pull out into the lead. I suspect those who have to focus on business, will fall behind. The healthy skepticism, but conservatism, of academia, will continue to butt heads with the optimism and overpromising press releases coming from companies. This is a good thing, not a bad thing, though in the end both sides would be better served by being sent to their corners for a time out sometimes.*

Ashley Montanaro

> *This continues to be an extremely exciting time for the development of quantum computers. Major experimental advances continue to be made, hand-in-hand with new algorithms and other theoretical improvements. Multiple groups have demonstrated quantum algorithms outperforming classical supercomputers. Yet there are still significant challenges faced between the point we are at today and the achievement of fault-tolerant quantum computing, and it is crucial that expectations for the likely performance of quantum hardware over the coming few years are kept realistic.*

One respondent

[13] We note that this reported increased interaction goes somewhat against a view of quantum computing research as a competitive race, and favors seeing it as a collaborative effort that benefits science and society.

> *Solutions to many difficult technical problems will require hard work and patience; overall however, I am optimistic about the future.*

Simon Benjamin

> *The field is over-hyped, but perhaps that is a necessary phase in the development of any radical technology. There is a 'prisoner's dilemma' scenario of different (especially, commercial) players needing to one-up each other in terms of promising rapid accomplishments, to their mutual harm. Clear standards and definitions are needed.*

One respondent

> *"It is absolutely remarkable to see how the quantum eco-system is evolving—the level of excitement, the depth of the questions and the interconnections between communities and research groups.*
>
> *[A]t this point, it seems that the community is on an amazingly promising and exciting track towards not only getting closer to realizing a quantum computer, but also making very important progress with much larger scientific impact on physics, computer science and mathematics.*

Daniel Gottesman: *It remains an exciting time in the development of quantum computation*.

Alexandre Blais: *It is a very exciting time to be doing research in quantum computing.*

Stephanie Simmons: *The pace of progress is astonishing and accelerating. The next five years will be a very wild ride.*

One respondent

> *We are approaching the point in time where quantum computers will have to begin creating value by delivering solutions to practically relevant problem. This [to] secure a future stream of investments into the field.*
>
> *Companies [that] have committed to public roadmaps whereby their progress can be measured on more or less a year-by-year basis, will furthermore be evaluated with respect to how well they will be able to meet their own milestones.*
>
> *In short, we live in interesting times. The decade to follow will likely be* very [Note: emphasis added in editing] *interesting.*

Kae Nemoto

> There will be two trends in the development of quantum computers. One is of course for fault-tolerant quantum computers and the other is for [small- to middle-] scale quantum computers. As these two share the core technology, up to now there is not much different, but they will probably grow in quite different directions.

One respondent:

> *I'm hopeful that simulations with analog and digital quantum platforms will provide a variety of insights into highly correlated quantum many-body quantum systems over the next 5-10 years.*

One respondent

> *As a field, we need to continue to be careful to not overhype applications of quantum computers, in particular the potential for quantum computers today to outperform classical solutions. Need to help people understand the real "power" of today's machines vs. full promise of quantum. Need to be clear that we are still very much in development phase and the industrial-scale applications will come when we get to 1M qubits or more.*

One respondent

> *There are important breakthroughs that will be necessary to develop a fault-tolerant quantum computer. It appears that none of the systems under study today gather all the required attributes.*

Tracy Northup

> *It's an exciting time to be working in this area! I continue to be surprised by how quickly the field is evolving on so many fronts, and what's very rewarding is how much discussion and collaboration there is across the board: between experimentalists and theorists, between experimentalists working on different platforms, ...*

Summary and outlook

A fully working quantum computer can be seen as the 'holy grail' of quantum technologies, but also as a major threat for cybersecurity. Specifically, it is a threat for cryptosystems that are based on the difficulty of solving certain mathematical problems with present computational devices. Those problems would be relatively easily solvable by a quantum computer large and reliable enough to run the appropriate quantum algorithms.

Building a quantum computer requires scientific and engineering advances that will take several years to be developed and implemented as well as focused effort and resources. The key challenge to overcome is the natural "fragility" of the quantum features that we think make quantum computing more powerful than classical computing.

The quest for a quantum computer has been often described as a 'quantum race' (Hsu, 2019), with competition at the level of nations as well as of private companies. This competition has substantially heated up in recent years, with the entry of new major private players, large grants from governments, and the birth and growth of many start-ups fuelled by venture capital. It has also been described as a marathon, rather than a sprint race, because of the relatively long-term research and investments that will be needed.

Nonetheless, there could be sudden accelerations, which may come in the form of scientific or engineering breakthroughs. We expect improvements both in hardware implementations and from new schemes for error correction and fault tolerance, that is, from schemes intended to overcome the fragility of quantum features and allow quantum calculations to be done using so-called logical qubits, reliably encoding and processing quantum information even when dealing with underlying physical qubits prone to errors. The convincing demonstration of such logical encoding and processing, in ways that indicate a feasible path to realize a full-fledged quantum computer, is the next big milestone targeted by quantum computing research, also according to the surveyed experts. Cyber-risk managers may want to track developments in that direction to understand how quickly quantum computers are becoming a reality.

In general, the expert opinions we have collected and summarized in this report offer unique insight into the quantum threat timeline. We have more than doubled the number of respondents since the first report in 2019, also tracking changes in opinions. Forty-six experts estimated the likelihood of the realization of a quantum computer that could break a scheme like RSA-2048. While most of the experts (25/46) judged that the development of such a quantum computer within the next 5 years is very unlikely ("<1%"), several (21/46) indicated the likelihood as non-negligible. We find it remarkable that only 24/46 judged the likelihood as small as "<1%" or "<5%" within 10 years. Within the latter timeframe, the rest of the respondents indicated already a significant likelihood, to the extent that 8/46 judged it about as much likely as unlikely ("about 50%") and 7/46 considered it even likely (">70%"). The risk aversion/appetite of companies and institutions can vary significantly, but we expect that for critical systems such estimated likelihoods represent a serious concern.

The likelihood the experts assign to the quantum threat may change from yearly survey to yearly survey, because several factors—from recent results in the field, to changes in investment levels—influence both the actual threat timeline and the opinion the experts have on it. Our series of reports allows one to track such an evolution. Comparing this year's opinions to the results of the surveys we conducted in 2019 and 2020, the experts appear to be more confident about the quantum threat becoming concrete in the medium-to-long term.

At the technological and scientific level, there are several competing potential physical implementations for quantum computing. It is not yet clear which will be the winner, nor that there will be necessarily only one winner. Presently, according to the experts' opinions, superconducting circuits and ion traps seem to have an edge over the competition, but other platforms continue to be developed, and some, like integrated optics, have attracted renewed attention in the recent times. In general, surprises could come from any implementation, and several respondents point to 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—or eventually be necessary for—scaling up the number of physical and logical qubits.

Whenever one deals with opinions rather than hard facts, it is appropriate to consider how reliable or partisan such opinions might be. Our respondents are generally devoting their careers to quantum information science and quantum computing. One could therefore wonder whether they are necessarily biased toward believing in the possibility of realizing a fault-tolerant computer, and/or believing it could happen sooner rather than later. Alternatively, there may be an instinct by some scientists to "under-promise" and continue to "over-deliver". Nevertheless, we are confident that our respondents have tried to provide the best possible realistic estimates. Quantum computing corresponds to changing the paradigm of computation itself. Working in a field that pushes the conceptual and practical limits of what humans and human-made tools are capable of requires some optimism, but it also requires a deep critical capacity that is necessary to identify and overcome roadblocks. The experts we surveyed are leading scientists also because they excel at such critical thinking.

The logical possibility that consequential quantum cryptanalysis is, for some reason, 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. 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 / earlier than could be expected, rather than never.

While building a cryptographically-relevant quantum computer is a formidable task, it is important for people managing cyber-risk to understand that there is nothing close to a scientifically convincing or established argument for why the efforts currently underway are likely to fail, especially in the medium-to-long term. Progress in the last year, including the demonstration of several aspects of quantum error correction and further realizations of so-called "quantum supremacy", as well as the significant momentum of the field—in terms of activities, results, and resources—should probably trigger caution, directed to developing crypto-agility and resilience against quantum attacks. A respondent wrote:

> *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.*

In similar spirit, John Martinis, a pioneer of superconducting implementations and leading the first demonstration of the quantum advantage of a programmable quantum processor over classical devices, suggests a corresponding prudent timeframe for action, based on the rate of progress he is seeing:

> *[T]he takeaway message is that quantum safe encryption needs to be developed and deployed in the next 5 years to be reasonably safe. Right now would be better.*

The Global Risk Institute and evolutionQ Inc. have already made available a quantum risk assessment methodology for taking estimates of the threat timeline and evaluating the overall urgency of taking action (Mosca & Mulholland, A Methodology for Quantum Risk Assessment, 2017).

The Global Risk Institute and evolutionQ Inc. will provide an update of this survey in approximately one year. This will allow us to further track the evolving opinion of experts and any changes in the expected timeline for the quantum threat to cybersecurity.

References

Andersen et al. (2020). Repeated quantum error detection in a surface code. *Nature Physics, 16*, 875.

Bombin, H., & Martin-Delgado, M. A. (2006). Topological quantum distillation. *Phys. Rev. Lett., 97*, 180501.

Chen et al. (2021). Exponential suppression of bit or phase errors with cyclic error correction. *Nature, 595*, 383.

DiVincenzo, D. P. (2000). The Physical Implementation of Quantum Computation. *Fortschritte der Physik, 48*, 9.

Egan, L. e. (2021). Fault-tolerant control of an error-corrected qubit. *Nature, 598*, 281. doi:https://doi.org/10.1038/s41586-021-03928-y

F. Arute et al. (2019). Quantum supremacy using a programmable superconducting processor. *Nature, 574*, 505.

Feynman, R. P. (1981). Simulating physics with computers. *International Journal of Theoretical Physics, 21*, 6/7.

Fowler, A. G., Mariantoni, M., Martinis, J. M., & Cleland, A. N. (2012). Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A, 86*, 032324.

Frolov, S. (2021). Quantum computing's reproducibility crisis: Majorana fermions. *Nature*, 350-352.

Gambetta, J. (2020, 9 15). *IBM Research Blog.* Retrieved from https://www.ibm.com/blogs/research/2020/09/ibm-quantum-roadmap/

Gheorghiu, V., & Mosca, M. (2019). Benchmarking the quantum cryptanalysis of symmetric, public-key and hash-based cryptographic schemes. *arXiv:1902.02332*.

Gheorghiu, V., & Mosca, M. (2021). *A Resource Estimation Framework For Quantum Attacks Against Cryptographic Functions: Recent Developments.* Global Risk Insitute.

Gidney, C., & Ekerå, M. (2021). How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum, 5*, 433. doi:https://doi.org/10.22331/q-2021-04-15-433

Grover, L. K. (1996). *Proceedings, 28th Annual ACM Symposium on the Theory of Computing (STOC)*, (p. 212).

Horsman, C., Fowler, A. G., Devitt, S., & Van Meter, R. (2012). Surface code quantum computing by lattice surgery. *New J. Phys., 14*, 123011.

Hsu, J. (2019, January 9). *IEEE Spectrum*. Retrieved from https://spectrum.ieee.org: https://spectrum.ieee.org/tech-talk/computing/hardware/race-for-the-quantum-prize-rises-to-national-priority

Kitaev, A. (2003). Fault-tolerant quantum computation by anyons. *Annals of Physics, 303*, 2.

Max, R., Kovacs, M., Zoller, P., Mlynek, J., & Calarco, T. (2019). Europe's Quantum Flagship initiative. *Quantum Science and Technology, 4*, 020501.

Mosca, M. (2013). *e-Proceedings of 1st ETSI Quantum-Safe Cryptography.*

Mosca, M., & Mulholland, J. (2017, January 5). *A Methodology for Quantum Risk Assessment.* Retrieved from Global Risk Institute: https://globalriskinstitute.org/publications/3423-2/

Mosca, M., & Piani, M. (2019). *Quantum Threat Timeline.* Global Risk Institute. Retrieved from https://globalriskinstitute.org/publications/quantum-threat-timeline/

Mosca, M., & Piani, M. (2021). *Quantum Threat Timeline Report 2020.* Global Risk Insitute.

National Academies of Sciences, Engineering, and Medicine. (2019). *Quantum Computing: Progress and Prospects.* Washington, DC: The National Academies Press.

Nielsen, M. A., & Chuang, I. (2002). *Quantum computation and quantum information.* Cambridge University Press.

Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. *Quantum, 2*, 79.

Raymer, M. G., & Monroe, C. (2019). The US National Quantum Initiative. *Quantum Science and Technology, 4*, 020504.

Ristè et al. (2020). Real-time processing of stabilizer measurements in a bit-flip code. *NPJ Quantum Information, 6*, 1.

Rivest, R. L., Shamir, A., & Adleman, L. (1978). A Method for Obtaining Digital Signatures and Public-key Cryptosystems. *Communications of the ACM*, 120.

Ryan-Anderson et al. (2021). Realization of real-time fault-tolerant quantum error correction. Retrieved from https://arxiv.org/abs/2107.07505

Sevilla, J., & Riedel, C. J. (2020). *Forecasting timelines of quantum computing.* Retrieved from https://arxiv.org/abs/2009.05045

Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review, 41*, 303.

Wu et al. (2021). Strong Quantum Computational Advantage Using a Superconducting Quantum Processor. *Phys. Rev. Lett., 127*, 180501.

Zhong et al. (2021). Phase-Programmable Gaussian Boson Sampling Using Stimulated Squeezed Light. *Phys. Rev. Lett., 127*, 180502.

A. Appendix

A.1 List of respondents

The respondents who have taken part in all our surveys so far, and whose opinions are tracking on multiple years, are listed at the top of this table, and their index has a grey background. Those who took already part in the 2020 survey but not in the 2019 one are listed immediately after (light-grey background for the respondent index).

A short description/bio that emphasizes the rationale for the inclusion of each respondent is provided after the table

#	Name	Institution	Country
1	Dorit Aharonov	Hebrew University of Jerusalem and QEDMA quantum computing	ISR
2	Dave Bacon	IonQ Inc.	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 Research	CHN
9	Martin Ekerå	KTH Royal Institute of Technology and Swedish NCSA	SWE
10	Artur Ekert	University of Oxford, and Centre for Quantum Technologies, National University of Singapore	GBR/SGP
11	Daniel Gottesman	University of Maryland	CAN
12	Jungsang Kim	IonQ Inc. and Duke University	USA
13	Ashley Montanaro	PhaseCraft and University of Bristol	GBR
14	Andrea Morello	UNSW Sydney	AUS
15	Yasunobu Nakamura	RIKEN Center for Quantum Computing	JPN
16	Tracy Northup	University of Innsbruck	AUT
17	Peter Shor	Massachusetts Institute of Technology	USA
18	Stephanie Simmons	Simon Fraser University and Photonic Inc	CAN
19	Krysta Svore	Microsoft	USA
20	Frank Wilhelm-Mauch	Institute for Quantum Computing Analytics, Jülich Research Center	GER
21	Shengyu Zhang	Tencent	CHN
22	Sergio Boixo	Google	USA
23	Fernando Brandão	California Institute of Technology	USA
24	Dan Browne	University College London	GBR
25	Eleni Diamanti	CNRS and Sorbonne University	FRA
26	Joe Fitzsimons	Horizon Quantum Computing	SGP
27	Yvonne Gao	Centre for Quantum Technologies, National University of Singapore	SGP
28	Winfried Hensinger	Sussex Centre for Quantum Technologies, University of Sussex	GBR

29	Elham Kashefi	School of Informatics, University of Edinburgh CNRS, LIP6, Sorbonne University	GBR/FRA
30	Sir Peter Knight	Imperial College London	GBR
31	Yi-Kai Liu	US National Institute of Standards and Technology (NIST)	USA
32	Klaus Moelmer	Aarhus Institute of Advanced Studies, Aarhus University	DNK
33	Bill Munro	NTT Basic Research Laboratories	JPN
34	Nicolas Menicucci	RMIT University	AUS
35	Kae Nemoto	National Institute of Informatics	JPN
36	John Preskill	California Institute of Technology	USA
37	Simone Severini	Amazon Web Services	USA
38	Lieven Vandersypen	QuTech and Kavli Institute of Nanoscience, TU Delft	NLD
39	David Wineland	University of Oregon	USA
40	James Daniel Whitfield	Dartmouth College	USA
41	Gregor Weihs	University of Innsbruck	AUT
42	Jun Ye	JILA, NIST and University of Colorado	USA
43	Jay Gambetta	IBM	USA
44	Justin Ging	Honeywell Quantum Solutions	USA
45	Chao-Yang Lu	University of Science and Technology of China	CHN
46	John Martinis	University of California, Santa Barbara	USA
47	Jacob Taylor	Joint Quantum Institute	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
An international expert in the theoretical and computational studies supporting the implementation of realistic quantum devices. He is the Associate Director of the UK National Hub on Networked Quantum Information Technologies, leading the package on quantum architectures, standards and systems integration.

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.

Fernando Brandão
Leading theoretical physicist specializing in quantum information theory. He is a Professor of Theoretical Physics at Caltech and the Head of Quantum Algorithms at Amazon Web Services.

Dan Browne
Professor of Physics at the University College London, where he has been also Director of the EPSRC Centre for Doctoral Training in Delivering Quantum Technologies. Among other contributions, he is renowned for his work on measurement-based quantum computation.

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.

Eleni Diamanti
A leading researcher at the French National Research Centre (CNRS) LIP6 Lab. Her work focuses on experimental quantum cryptography and communication complexity, and on the development of photonic resources for quantum networks.

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.

Joe Fitzsimons
A leading theoretical physicist and CEO of Horizon Quantum Computing. He is renowned for his contributions to blind quantum computing. His current goal is to develop programming tools that simplify software development for quantum computers.

Jay Gambetta
After major contributions to the theoretical study of superconducting systems, he joined IBM, where he is now Vice President of Quantum Computing, leading the effort to build a quantum computer based on superconducting qubits.

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

Justin Ging
He is the Chief Commercial Officer of Honeywell Quantum Solutions, which is focused on the development of quantum computers based on trapped ions.

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, associate director of the Networked Quantum Information Technologies and on the executive team of the Quantum Internet Alliance.

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)

Chao-Yang Lu
Professor of Physics at the University of Science and Technology of China, where is co-leads three teams working on quantum foundations and quantum technology. His results include the first optical demonstration of quantum supremacy, based on so-called boson sampling.

John Martinis
A worldwide leader in the development of the superconducting architecture for quantum computers, which also resulted in the first demonstration of so-called quantum supremacy.

Frank Wilhelm-Mauch
A leading theoretician working closely with experimentalists, he focuses on modelling and controlling superconducting circuits. He is the coordinator of the European project "OpenSuperQ", aiming at building a European quantum computer with 100 superconducting qubits in the next few years.

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.

Ashley Montanaro
An international expert on quantum algorithms and computational complexity, as well as quantum query and communication complexity, working on establishing fundamental limits and capabilities of quantum devices. He is the author of influential papers on quantum computational supremacy.

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.

Bill Munro
A distinguished scientist and group leader at NTT BRL. He was a leader in HP’s development of quantum enabled technologies and currently runs 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.

Tracy Northup
Leads the Quantum Interfaces Group at the University of Innsbruck. Her research uses optical cavities and trapped ions as tools to explore quantum-mechanical interactions between light and matter, with applications for quantum networks and sensors.

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.

Krysta Svore
She leads the Microsoft Quantum – Redmond (QuArC) group at Microsoft Research in Redmond, WA. Her research focuses on quantum algorithms and how to implement them fault-tolerantly, including coding them in high-level programming language and compiling them into fault-tolerant circuits.

Jacob Taylor

His research focuses on hybrid quantum systems, on applications of quantum information science, and fundamental questions about quantum behaviour. He was the assistant director for quantum information science at the White House from 2017 to 2020, leading the creation of the US National Quantum Initiative.

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.

James Daniel Whitfield
Leads a group at the Department of Physics and Astronomy of Dartmouth College. His research focuses on understanding the potential and the limitations of new and existing computers to perform physical simulations.

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.

David 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."

Jun Ye
A leading scientist, known for developing technologies in the areas of high-precision laser spectroscopy, atomic and molecular cooling and trapping, optical frequency metrology, quantum control, and ultrafast lasers.

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

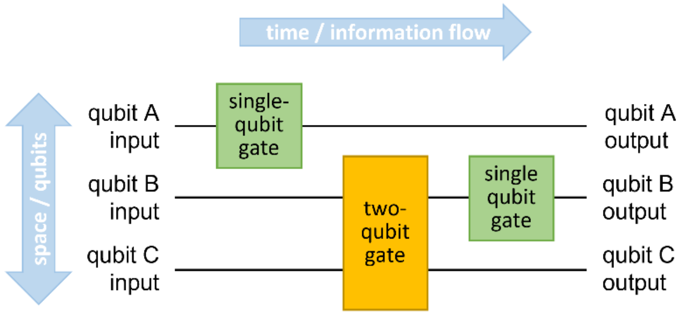

Figure 29 Illustration of the circuit/gate model for quantum computation. Each qubit corresponds to a horizontal line, so that multiple stacked lines illustrate many qubits. A qubit can be transformed individually by means of single-qubit gates, and two qubits can interact via a two-qubit gate. A given 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.

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 29). 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.

A.3 Error Correction

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

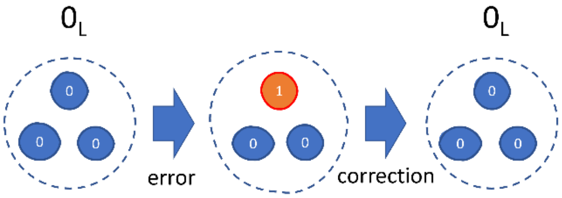

Figure 30 Example of classical information encoded logically. Several imperfect/error-prone physical bits (warped filled blue circles) are used to encode a logical 0, denoted 0_L (dashed perfectly round circle), by means of a repetition code: 0_L is encoded as 000 at the physical level. Errors can occur 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. As long as 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 30, 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, Mariantoni, Martinis, & Cleland, 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 & 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, Fowler, Devitt, & Van Meter, 2012).

A.4 Questions

Regarding the wording of the core questions, 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. In fact, we saw some respondents point out the need to make further assumptions and interpret in a specific way the questions in order to provide a sensible answer.

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 further respondents in 2020, and most recently in 2021.
- 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 about 30 minutes. In some cases, to secure responses to at least the major key questions revolving around timelines, and particularly for respondents who had taken part in the previous versions of the survey and were key in detecting trends in opinion, we gave the option to provide input about only key questions—the two explicitly stated in Section 3.1.
  NOTE: Given the latter flexibility, not all respondents have provided answers to all questions, some of which were optional to begin with. Nonetheless, almost all the 47 respondents provided input for the questions more essential to estimating the quantum threat timeline.
- 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; the results are still much more informative than what available prior to this series of surveys.

- We did keep several of the questions at the basis of the 2019 and 2020 reports unchanged, so to capture a change in trends (see Section 3.1 for more details).
- We modified to some extent the set of questions, 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);
  - the respondents' feedback from the 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 a more nuanced comment. This mitigated to some extent the issue of the experts potentially responding to the same questions under a different set of assumptions.

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

Q: *Please indicate how likely you estimate that a single fully controllable fault-tolerant (logical) qubit within a scheme / architecture viable for scaling will be demonstrated within the next 5 years, 10 years, 15 years, 20 years, and 30 years.*

Q: *What do you consider to be an important milestone to be achieved after realizing a scalable fault-tolerant logical qubit but before achieving the quantum factorization of 2048-bit numbers (as described in the previous questions)?*

Questions on "Factors that may impact the quantum threat timeline"

Q: *Assuming that the COVID-19 pandemic has had / will ultimately have impacted negatively the progress being made towards achieving the construction of a scalable fault-tolerant quantum computer, how large do you estimate such a slowdown to be?*

Possible answers: < 3 months, 3-12 months, 1-2 years, > 2 years, I prefer not to answer / I do not have an opinion

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 August 2020) 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 June 2022)*

We further asked the respondents to provide any information they were willing to share about their own research (either theoretical or experimental in nature)

Finally, we gave the option to freely comment on the state of the field

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

A.5 Responses and analysis

In this section of the Appendix we provide some informative quotes by the respondent, and provide some details on our methodology in handling and analyzing the responses.

Physical realizations

With respect to our questions about physical realizations, one respondent notes that[14],

> *[t]aking the analogy of the physical material as "hardware" and the fault-tolerant quantum algorithm as "software," it would be useful to ask [a similar question about the potential of] different "firmware" available for each hardware architecture – i.e., the low-level choice of how to encode quantum information in the hardware,*

making the point with the specific example of photonic implementations:

> *From a perspective that treats light modes (rather than photons) as the physical systems, there is nothing special or fundamental about single-photon optics. That is, photonic qubits are just one particular choice for [the] encoding of digital quantum information into an optical mode. Other choices – exactly the same as those available in microwave systems [..] – may offer better performance. This extra performance would be due to built-in error protection, the ability to design the bias of the noise they experience, and the different sets of gates that are simple in each architecture.*

Stephanie Simmons thinks that modularity and hybrid realizations are both key aspects to facilitate the realization of quantum computers, and in that they are associated:

> *It will be easier to build modular, distributed quantum computers than a single monolithic quantum supercomputer. With this in mind, integrated all-photonic approaches as well as hybrid solutions using telecom*[15] *photons rank highly. [In particular, I expect that] long-lived spins will be most competitive in a hybrid framework with telecom optical photons.*

She adds:

> *Although this is in some sense a race, there could be many winners. [..] What will ultimately decide the lead candidate(s) is when consensus begins to emerge around a dominant design; this will affect resource allocation which will accelerate some platforms preferentially.*

A similar sentiment about modularity and hybrid realizations appears to be shared by **Kae Nemoto**, who writes:

> *As [also the IBM roadmap states], some technology to connect different quantum chips, or qubits[,] would be necessary for almost any quantum computer hardware. [..] There are a few QC architecture[s] [..] fully hybrid and distributed, and they are highly scalable.*

[14] We note in passing that this is a comment that, by pointing out possible refinements of / additions to our questions, may inform changes in future installments of our survey.

[15] "Telecom photons" indicates photons whose wavelength is in the bands used in standard telecommunication infrastructure, e.g., for transmission through standard optical fibers.

With respect to the point that there could be many winners and that a clear winner is quite far from having emerged, **Tracy Northup** comments with a positive opinion on the richness and diversity of the quantum arena/"racetrack":

> *I think we're still far enough away from 100 logical qubits that we can expect the relative potentials of these various systems to shift back and forth a lot over the coming years. That's exciting and a testament to how much progress is being made in parallel on different platforms.*

She also comments on hybrid solutions, making it explicit that one rationale for them is that different platforms may excel at different hardware functions, and hybrid solutions try to make use to the best in all. She writes:

> *I expect hybrid solutions to be a part of digital quantum computers with 100 logical qubits[.] [..] I'm encouraged by recent experimental progress on optical and microwave interconnects and also on specific hardware functions (e.g., memory, fast computations) that might in the future be assigned to different platforms.*

Nonetheless, she warns that hybrid solutions come at a cost—that of combining the various pieces of the solution:

> *I do want to point out that [realizations of] hybrid solutions are challenging tasks that are going to take some time (and will be platform-specific), so we shouldn't expect a breakthrough from one year to the next that makes it all work.*

One respondent points out that, indeed, various realizations may have varying strengths and weaknesses, and makes it clear how they interpreted the various choices of "potential":

> *[C]onsiderations [of the potential] need to [take into account] speed, size, and reliability. Thus some [implementations] are "not promising" due to the very slow operation speed. Others are "somewhat promising" because overall system size or 2-qubit gates may be challenging. "Some promise" indicates that some miracles are still needed to reach that scale [within the next decade or shortly after].*

One respondent summarizes in this way their view of the recent progress exhibited by the various types of platforms:

> *Topological systems have taken a credibility hit this year. Superconductors, trapped ions and spins in silicon are making great progress. Integrated optics holds lots of promise, although its state of development is not as openly known as in other leading systems.*

In his first point, they are referring to a scientific debate regarding the unconfirmed detection of so-called Majorana fermions (Frolov, 2021), a new type of quantum particle[16] which could be exploited to implement topological quantum computation. In their last point, they instead refer to the fact that an

[16] In this case, we are not discussing fundamental particles. Rather, the proper design (e.g., the combination of layers) and manipulation of materials (e.g., choice of temperature and of electro-magnetic fields applied) can lead to the emergence of "effective" particles.

integrated-optics approach has been being developed particularly by new and somewhat more secretive start-ups.

Quantum factoring

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. Table 1 provides the raw aggregate counts of the responses.

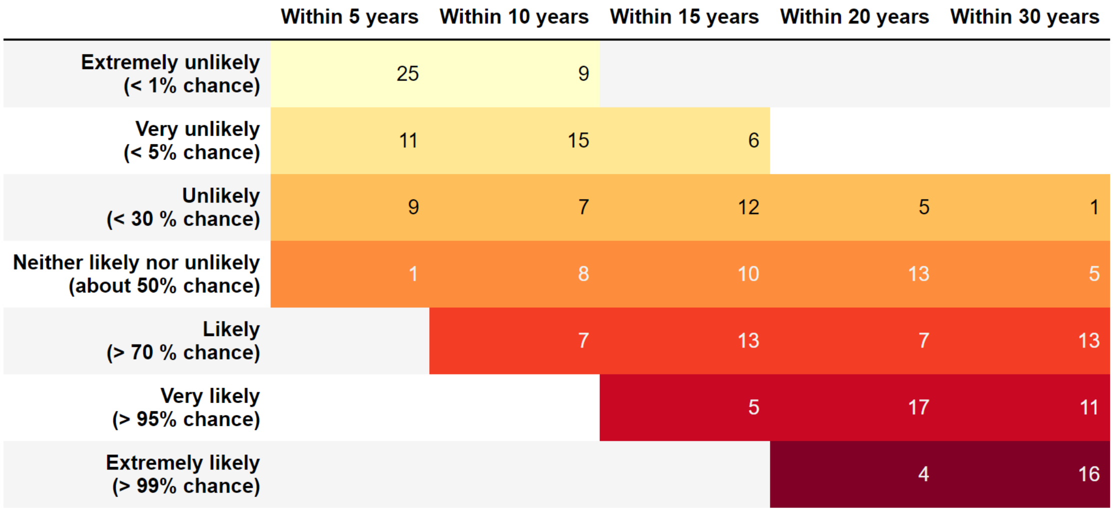

	Within 5 years	Within 10 years	Within 15 years	Within 20 years	Within 30 years
Extremely unlikely (< 1% chance)	25	9			
Very unlikely (< 5% chance)	11	15	6		
Unlikely (< 30 % chance)	9	7	12	5	1
Neither likely nor unlikely (about 50% chance)	1	8	10	13	5
Likely (> 70 % chance)		7	13	7	13
Very likely (> 95% chance)			5	17	11
Extremely likely (> 99% chance)				4	16

Table 1 Aggregate estimates for the likelihood of a quantum computer able to break RSA-2048 in 24 hours. What is shown is the number of experts (out of the 46 who provided such an estimate) who indicated one of the specific likelihood ranges (rows) for each of the time frames considered (columns). The colors displayed are associated to the likelihood and want to convey how a higher likelihood corresponds to a higher risk, at least from the perspective of cybersecurity.

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.

In order to derive from the responses the cumulative probability distributions as shown in Section 5.2, 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:

Optimistic assignment:

Extremely likely (> 99% chance)	100%
Very likely (> 95% chance)	99%
Likely (> 70 % chance)	95%

Neither likely nor unlikely (about 50% chance)	70%
Unlikely (< 30 % chance)	30%
Very unlikely (< 5% chance)	5%
Extremely unlikely (< 1% chance)	1%

Pessimistic assignment:

Extremely likely (> 99% chance)	99%
Very likely (> 95% chance)	95%
Likely (> 70 % chance)	70%
Neither likely nor unlikely (about 50% chance)	30%
Unlikely (< 30 % chance)	5%
Very unlikely (< 5% chance)	1%
Extremely unlikely (< 1% chance)	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 over the set of respondents for the sentiment/likelihood estimates ensures that outlier estimates (that is, estimates that are either too optimistic or too pessimistic) tend to have less of an effect, and may well cancel each other out. Nonetheless, such averages do 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.

Having made these general cautionary points, we can further try to understand the (relatively minor) change in outcome of our surveys for the fixed set of respondents (see Section 5.2.1).

We put forward this (non-exhaustive and not mutually-exclusive) list of potential explanations:

- 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 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.
- The effect of the present ongoing pandemic, which the experts estimate has slowed down progress to a significant extent (see Section 5.4.3).

Logical qubits and fault-tolerant schemes

Here some excerpts from the experts’ comments about the timeline for the realization of one or more fault tolerant qubit (see Section 5.3).

Stephanie Simmons

> *I've noticed that "logical qubits" is a term being applied very liberally lately—to mean any encoded quantum information with a lifetime beyond that of its constituent physical qubits. We will soon need to specify the total logical error rate for such questions to be correctly interpreted.*

Respondent

> *This question is a bit ill-defined, as it depends on your criteria for what you consider a fully-controllable fault-tolerant logical qubit.*

Ashley Montanaro

> *I answered [..] assuming that "fault-tolerant logical qubit" means a qubit suitable for inclusion in a fault-tolerant quantum computing scheme with improved error rates compared with no error correction, as opposed to a "perfect" qubit.*

Dave Bacon

> *Without an estimate of the fidelity of the scheme it is not clear what this question means. I have taken it to mean any [fault-tolerant] scheme that is below [the fault-tolerant] threshold and does all of the parts of a [fault-tolerant] scheme, including a universal gate set.*

Respondent

> *[One] must not only 'realise' logical qubits but also fully manipulate a plurality of them [..] and all within performance characteristics that are consistent with logical error rate falling with further scaling.*

Joe Fitzsimons

> *I am assuming the previous question means a quantum computer having qubits encoded to an extent that ensures a near zero error rate on non-trivial calculations.*

Respondent

> *The question assumes that a "nearly perfect" logical qubit will be demonstrated before modestly good qubits with reasonable number will be built. I do not necessarily see it that way.*

Respondent

> *The answer to the above question of course depends on what you mean by fully controllable fault tolerant logical qubit. The way I interpret this, this term means a qubit which can be plugged in into a full-fledged fault tolerant quantum computer with many qubits. In other words, the error rate had been reduced to negligible values filling say an operation of 1000 logical qubits performing a circuit of depth say 1000. Namely, the effective error rate of the logical qubit, with the help of the fault tolerance, needs to be computed as if the qubit lives in a full-fledged quantum computer.*

Andrea Morello points to scalability as an essential feature that, perhaps independently of the exact interpretation of the question, sets the bar relatively high:

> *The qualifier "within a scheme/architecture viable for scaling" is the key here. Some embryonic logic qubits are already out there.*

Bill Coish, in his interpretation, pins down a (minimum) number of physical qubits that he considers as necessary for building a fault-tolerant quantum computer:

> *I take the "viable for scaling" clause to be the limiting factor, where I imagine one would need to scale to >10 million physical qubits.*

One respondent warns about the timescale within which the field needs to prove the viability of fault-tolerance:

> *If it doesn't happen in 5 years the industry will likely collapse.*

Joe Fitzsimons is optimistic:

> *I think there are clear signs this is coming.*

One respondent points to an important possibility, especially in the context of this report, which is about the quantum threat timeline:

> *It might be that we achieve the factorization task before the fully fault tolerant single qubit is achieved, by using other methods. I believe there could be error mitigation techniques which could possibly enable achieving the task of 2048-bit factorization before fault tolerance is achieved.*

Most promising fault-tolerant schemes

Here are some excerpts about the issue of which fault-tolerant schemes are presently the most promising (see Section 5.3.1).

Daniel Gottesman provides a detailed breakdown of recent progress and its potential significance:

> *Based on extrapolations of the current state-of-the-art, I see three schemes which are comparably likely to win out. The first is fairly traditional surface-code architectures or similar topological codes. The second is based on high-rate low-density parity check (LDPC) codes. These will have much lower overheads than surface codes but may require systems with long-range gates or interconnects, favoring photonic or hybrid implementations with a photon component, but likely achievable in other implementations such as ion traps. (Superconductors might be possible as well, but the technology to do so is less developed.). Recent progress on LDPC codes has allowed the construction of codes with much higher distance than was previously known. This has not yet been applied to improve fault tolerance, and there remains a lot of work to be done in this regard. The third area which has seen recent progress is a combination of bosonic codes for directly encoding qubits in continuous variable systems, such as photonic modes and codes to allow higher error thresholds with biased noise dominated by phase damping, a very common type of error.*

Stephanie Simmons stresses the importance of exploring the promising features of LDPC codes:

> *It has become evident that high-performance [in term of threshold and encoding rate] quantum error correction codes require physical qubit connectivity far beyond that of the surface code [..] such as the recent LDPC codes making breakthrough after breakthrough over 2021. These high-connectivity codes naturally fit a distributed quantum computing framework. These codes don't yet all have full fault-tolerant universal gate sets worked through. [T]here is a chance that LDPC codes won't need costly processes such as magic state distillation to achieve logical fault-tolerant universality. More work is urgently needed in this area.*

David DiVincenzo emphasizes the leading role of superconducting qubits, but also the improvements at all levels of the stack needed to successfully run a quantum computation:

> *There is significant progress on 2D layouts, [and the results from China show] that the [work pioneered by Google] can be built upon. Thus, I think of some version of the transmon/superconducting platform as being the most promising. I also see that the control stack and software have matured considerably in this area.*

Frank Wilhelm-Mauch:

> *Surface code or variants thereof, implemented with superconducting qubits, as they seem to be most balanced - their biggest weakness (currently still a bit of lacking fidelity) is quite addressable.*

Respondent:

> *Some variation of the surface code [..] This when considering large-scale error-corrected systems based on superconducting qubits featuring nearest-neighbor connectivity, and in particular the systems being developed by Google.*
>
> *The above is in the long run, when designing schemes for fault tolerance for long-lived logical qubits to be used for, e.g., factoring 2048-bit RSA integers. To demonstrate a fault-tolerant logical qubit built from a few physical qubits, there are other possible options.*
>
> *As a caveat, quantum error correction is currently a very active research area. As time progresses, it is likely that we will see advances. This in particular as systems are scaled up.*

One respondent supports combining discrete- and continuous-variable approaches:

> *Cat or GKP qubits in a bosonic mode concatenated with a topological qubit-level code. These are potentially available in both superconducting microwave cavities and in continuous-variable optics (i.e., beyond single-photon-based architectures).*

In similar spirit, **John Preskill** bets on *"Continuous variable coding in superconducting circuits."*

Elham Kashefi points to novel approaches in photonic quantum computing, which make use of modular architectures.

Alexandre Blais admits that "[his] answer is naturally biased by what [he] know[s] best", something likely true in general, and highlights the importance of considering properties of the noise:

> *My impression is that hardware-efficient quantum error correction is still the most promising approach over more standard quantum error correction scheme. In the former approach, one first deals with the most likely errors while in the former all types of errors (even those that are not likely) are taken care of resulting in a larger number of physical qubits per logical qubit than strictly necessary. Recent modifications of the surface code to these "bias-noise qubits" appear particularly promising. Although these ideas could potentially be extended to other architectures, [they] are now more developed in the context of superconducting qubits.*

Kae Nemoto stresses the potential advantages of a distributed architecture, particularly with respect to the issue of scalability:

> Distributed architectures are more scalable, and hence more promising for a large-scale QC. However, the clock cycles for these architectures are slower than monolithic design such as superconducting qubits and silicon spin qubits. So, the answer is dependent on the goal. If the question is to build a large QC system for fault-tolerance, I think that the best scalability of the distributed architectures has the advantage.

Tracy Northup highlights recent promising progress, but cautions that systems and schemes that are being used for initial demonstrations may not be the same that bring us to hundreds of logical qubits:

> I'm most aware of efforts to achieve fault tolerance with current hardware, for example, the recent Honeywell realization with trapped ions [..], in which fault-tolerant quantum error correction is demonstrated [..] for a single logical qubit but not yet for operations between two logical qubits and not yet below the pseudo-threshold. I expect that in the near future, we'll see similar demonstrations on other platforms (e.g., recent work with NV centers [..]), along with extensions to entanglement of logical qubits. But the question of what scheme works best for these initial demonstrations is different from the question of what will be the most promising (and feasible) at the level of 100 logical qubits.

Dave Bacon is one of several respondents who have high hopes for better schemes in the (relatively near) future, also thanks to the insight developed so far:

> *I do not believe the most promising scheme has been [invented]! I do believe, however, that these will fit in a "middle way" between brute force and topological approaches. In particular we now have immense control over engineering small quantum systems, and we know that topological approaches are naturally robust. But the topological approaches seem to suffer from very challenging problems of going from a messy condensed-matter system to the actual Hamiltonian you want[..]. I think the most promising schemes will utilize the new engineering expertise to build naturally robust (encoded) qubits. [..] Topological quantum computing tells you [timing] precision is not necessary, and we need small naturally encoded qubits that utilize this insight.*

Andrea Morello:

> *I don't think the most promising scheme has been invented yet. Over the 10-20 years timescale it will take for the hardware to develop far enough, I expect some clever new scheme will come up.*

Respondent:

> *The most promising current scheme is surface codes with magic state factories. I have hope that better schemes will be developed.*

Respondent:

> *I think it is quite likely that the problem of the huge overhead in terms of number of qubits needed for current fault tolerant schemes, will be resolved theoretically in the coming few years, maybe with some price to pay in other directions. This might completely change the picture in terms of the goal you defined related to factorization of 2048 bit numbers.*

Ashley Montanaro:

> *At the moment, the most promising scheme for large-scale fault-tolerance appears to be the surface code combined with lattice surgery, though I expect that there will be significantly improved fault-tolerance techniques developed over the next 10-year period.*

Level of funding of quantum computing research

Here are some informative quotes about the level of funding, provided in the context of the question of Section 5.4.1.

David DiVincenzo

Government investment continues to ramp up in China and Europe.

Dave Bacon

1) There is still considerable appetite in a low interest rate environment for large investments in quantum computing (like 2 to 3 more unicorn[17] *type valuations).*

2) There doesn't seem to be a lowering of the temperatures between China, US, Europe and quantum feeds this narrative.

Ashley Montanaro

Government investment appears unlikely to increase significantly in the near future, given previous large-scale injections of funding; one exception may be China. Private investment will continue to grow as the quantum industry matures, as demonstrated by several recent high-profile events.

Respondent

The hype about quantum computing seems to continue unabated. I think that will stimulate increased investment for at least the next two years.

Simon Benjamin

I put 'increase', but I do think we are approaching a peak.

Respondent

We are still in an era of a plethora of start-ups, many of which may prove to be unviable in the next 2-3 years, leading to a consolidation of backing for a select few young companies. I would expect the amount of venture capital to remain steady or increase slightly in total, and I believe that government investment worldwide will increase. I think the major boom has already happened, and now we're in a sustained growth phase. Whether that continues or stalls will depend on whether there are sufficient breakthroughs in fault tolerance and scalability in the next 5-10 years.

Klaus Moelmer

I am sure that many late-comers (universities, countries, industries, ...) will feel compelled to enter the game at this stage, and this may increase the global volume of investments, while not necessarily bringing more ideas to the global efforts - I even see signs of a competition instead of collaboration, embargos, concerns for IP rights, ...

Respondent

I believe that in the [next 2-5 years] we will see new ideas of how to make use of quantum computers in which standard quantum error correction still cannot be applied. I think this will lead

[17] A *unicorn* type valuation regards a rare and very large valuation of a privately-owned start-up.

the public/governments/investors to believe that quantum computers might be useful even before fault tolerance is achieved; this is the reason why I [think] the investment will only increase.

Daniel Gottesman

I get the sense (which may be inaccurate) that private investment in new quantum computing start-ups has slowed a bit, so we may be reaching a plateau before a shake-out in the start-up scene sometime in the next few years. I view the status of the private funding landscape as pretty uncertain right now. It could go in any direction.

Bill Coish

I think the huge potential of quantum computing has now been realized by governments/industries around the world and the level of investment may (temporarily) be peaking until we have a very significant breakthrough.

Andrea Morello

I see lots of warnings about hype, and the related nervousness of investors. However, large government spending (>$1 billion per country) is only just starting. This will keep up the momentum for some time.

Frank Wilhelm-Mauch

I think that the structure of investments will change - short-term profit seekers will move on but the current successes of quantum computing will keep a number of patient investors in the game. Specifically, this should apply to China.

Shengyu Zhang

Just hope that the start-up companies hype less so that the overall expectation from governments and the public remains at a rational level, and the resources are allocated to the right places.

Respondent

I believe the proper amount of the investment strongly depends on the amount of human resource available in the region. A significant amount of the investment should also be spent on training of workforce in the field.

Kae Nemoto

The rapid increase of funding in the last few years caused an imbalance in demand/supply of QIP researchers. The total funding might increase slightly, but I think that that would not be necessarily directly to the development of quantum computing.

Respondent

There has been a steady increase in the size of investments in recent years, and no indication that this is slowing down. At the same time, we are still in a "blue skies" phase where concrete profitable outputs are not be expected yet by investors.

Respondent

> *I think funding will increase until one or several of the start-ups collapse. Then the investor capital will slow down.*

Tracy Northup

> *My answer is simply based on past trends: I continue to be surprised and encouraged by the momentum of investment in quantum computing and haven't seen any signs of that slowing down yet.*

Respondent

> *[I]nvestment will continue to increase for next few years. [W]ithout significant progress demonstrated [in, say, 5 years], investment may then begin to slow. Continued progress must be demonstrated to maintain/increase investment.*

Respondent

> *Public investment by governments around the world will continue. Private investments might peak out in the next year or so and might be concentrated on more promising technology leaders ("picking winners").*

Global race to build a fault-tolerant quantum computer

We report relevant quotes by the experts about the issue of funding (see Section 5.4.2)

Stephanie Simmons

> *Funding is entirely global and China, Europe and North America are the three major quantum talent hotspots – these do not need to overlap. For example, Asian-funded (and/or controlled) efforts will exist where talent lives around the globe, which pushes the interpretation of this question. Similarly, many technology companies will expand globally if only through increasing remote work arrangements in the scramble for quantum talent.*

Klaus Moelmer

> *I do not see sufficient investments and willingness to "pick a winner" in Europe. Both China and US have public and private actors with unlimited funds and an appetite on quantum technologies.*

Alexander Blais

> *We are definitively seeing an acceleration of the research in China.*

Bill Coish

> *My feeling is that investment in Europe is less focused (despite the presence of the quantum flagship). China has the willingness and ability to focus large sums of money on the problem and the involvement of large powerful tech companies primarily in the US will guarantee that North America stays a front-runner.*

Andrea Morello

> *China and North America are making the fastest progress, and have by far the most concentrated investments, albeit with radically different investment schemes. Europe is ramping up and is addressing the issue of fragmentation by creating concentrated hubs [..]. Other players are Japan and Australia, which were home to some of the earliest breakthroughs in superconductors, optics and silicon, and still host plenty of know-how and talent.*

Frank Wilhelm-Mauch

> *It will be interesting to see if there will be a strong program in India or Russia.*

Shengyu Zhang

> *US is simply leading [..] in terms of good teams, candidate schemes, infrastructure, and collaboration.*

Respondent

> *The amount of budget and the number of people involved in those regions are much larger than others.*

Tracy Northup

> *I'm not focusing on a particular company or research group here, but I do feel that taken as a whole, efforts in both Europe and North America are doing the most to push the field forward (and fostering a very fruitful sort of competition).*

Respondent

> *I think "global" companies are ahead right now, but we don't know enough from China to know their status/investment/progress. [We] should assume they can make rapid progress at any time.*
>
> *[I think the biggest progress] will come from industrial/govt partnerships.*

Respondent

> *For Europe, I wrote "possibly" given the fragmentation of investments.*

It is worth concluding with the opinion of **Ashley Montanaro**, who provided his rationale for abstaining from answering the questions, by opposing the idea that quantum computing research should be seen as a race rather than an international cooperative effort:

> *I'm abstaining from these questions because I don't believe it is helpful to associate quantum computing with nationalism - just one example of the result of this viewpoint is the recently proposed exclusion of certain countries from EU quantum research programmes. Advances in quantum technologies have been achieved by a diverse set of teams worldwide in both the public and private sectors. In my view, quantum computing should not be a race between countries, but a collaborative effort for mutual benefit.*

Next milestone

Here some opinions/desiderata/goals expressed by the experts about the next intermediate milestone on the way towards the realization of a cryptographically-relevant quantum computer (Section 5.7).

Sergio Boixo

Modular scalable system with ~10 qubits

David DiVincenzo

An important milestone would be to link this processor to a quantum optical network. Then various distributed computing schemes, or repeater schemes, could be successfully performed with a much smaller processor than the target 2048 processor.

Respondent

Two fault-tolerant logical qubits with arbitrary fault-tolerant two qubit gates

Dave Bacon

100 qubit, 10000 depth circuits algorithms (i.e., O(n^2) algorithms of un-simulate-able size)

Ashley Montanaro

Scaling up to 100-1000 logical qubits, which would be enough to address many exciting applications, e.g., in quantum simulation.

Respondent

Implement simulation of a realistic, useful problem

Simon Benjamin

[O]ne must not only 'realise' logical qubits but also fully manipulate a plurality of them [..], and all within performance characteristics that are consistent with logical error rate falling with further scaling.

Klaus Moelmer

One must demonstrate fault tolerant gates among few qubits and assess the consequences of a geometric lay-out with spatially separated qubits.

Elham Kashefi

In terms of cryptography other quantum cryptanalysis approaches could be explored on the quantum computer such as the one based on quantum machine learning that is already classically proven a very promising route to adjust system parameters.

Respondent

I have to say that it is not absolutely clear that the order is right. It might be that we achieve the factorization task before the fully fault tolerant single qubit is achieved, by using other methods. I believe there could be error mitigation techniques which could possibly enable achieving the task of 2048 bit factorization before fault tolerance is achieved.

Daniel Gottesman

My favored milestone would be running a fault-tolerant simulation of a spin system beyond what could be done at the time using a classical computer. A somewhat earlier milestone would be a quantum supremacy demonstration using purely fault-tolerant logical qubits.

Alexandre Blais

After demonstrating a single fault-tolerant logical qubit, we will need to learn to make a few of those and have them do toy computations. At that stage, the field might very well look like the first few years of quantum computation where executed, but this time on logical qubits rather than physical qubits. After that, demonstrating any quantum computation of scientific or commercial value will be interesting and useful.

Stephanie Simmons

Useful repeaters for global quantum communication

Respondent

Digital quantum simulation of condensed matter or molecular systems.

Bill Coish

A wide range of quantum simulations of physical systems (including dynamics and non-equilibrium properties) would be a huge leap forward. Fault tolerance is important to guarantee that the simulations accurately represent the underlying physical model -- e.g., simulating the Fermi Hubbard model to help to advance problems of high-temperature superconductivity. Simulations of the fractional quantum Hall effect could reveal non-Abelian excitations. There are many many many more examples that would probably require far fewer logical qubits than factoring a 2048-bit number, but that are currently beyond the reach of scientific computing today.

Andrea Morello

There's a big gap between those two milestones! To be filled with interesting near-term applications, probably in quantum chemistry or related topics

Respondent

Multiple logical qubits would need to be fabricated and integrated into modules of increasing scale to form a working large-scale quantum computer.

Needless to say, there are a number of additional technical hurdles that will need to be overcome to achieve the necessary scaling, besides the integration of the logical qubits into modules. In particular, improved error rates, larger dilution refrigerators, better methods for bringing signals in and out of the refrigerators, and methods for interconnecting refrigerators, come to mind.

Frank Wilhelm-Mauch

Fault-tolerant two-qubit gates[;] fault-tolerant simpler algorithms [for] theoretical chemistry.

Shengyu Zhang

1. Push the scale to the next level, such as 10 or 100 [logical qubits].

2. Continue to improve the quality of a single qubit and a single operation, so that obtaining logical qubits becomes easier.

Respondent

Having a scalable fault-tolerant logical qubit, we would find useful applications in quantum simulation and quantum chemistry using those low-error logical qubits. They could also be used in NISQ applications.

Kae Nemoto

Universal logical gate set.

[D]emonstration of multi-chip operation (if it is not included in the single logical qubit demonstration, assuming that the [quantum computer] has some distributed nature)

Respondent

Demonstration of an architecture that truly scales.

Respondent

Fault tolerant quantum "supremacy" (apologies for the awkward term) (i.e., order 50 logical qubits)

Respondent

Demonstration of practical industrial applications

Respondent

New physics insights from simulation of quantum many-body phenomena.

Respondent

Demonstration of scalable control. Demonstration of tiling of logical qubits and maintenance of same error rates.

Respondent

I believe scalable technology development for modest quality qubits (better than NISQ) will come before a single fault-tolerant logical qubit will be demonstrated.

Respondent

Scaling up the number of fault-tolerant logical qubits in practice will come before realizing any algorithm.

Respondent

Achieving a dozen fault-tolerant logical qubits and doing fault-tolerant computation with them.

Joe Fitzsimons

There are several clear milestones. The most obvious one is the demonstration of sustained entanglement between two or more logical qubits. Demonstration of high fidelity encoded classical computation will also be an intermediate milestone, which may manifest as oracles in Grover-like algorithms. Other examples are of course factoring of a 512-bit number or similar (e.g., a historic RSA key size, something that was once believed to be intractable for classical computers).

Tracy Northup

A crucial question will be how to go from the number of qubits required for the single fault-tolerant logical qubit to the (much greater!) number required for 2048-bit factorization, while holding on to the low error rates that made scalable fault tolerance possible for the one logical qubit. What "go from" means is highly dependent on the specific platform: for superconducting qubits, this may be a question of calibrating a much larger number of non-identical qubits; for ions or cold atoms, this may be a question of how to use lasers and/or shuttling to address much larger arrays with the same precision. But there isn't a single platform for which the route from one logical qubit to many is straightforward!