Quantum Threat Timeline Report

Executive summary

Quantum computers harness the computational power of quantum systems and offer the ability to solve computational problems previously thought to be intractable. The quantum features that quantum computers rely on are very difficult to preserve and control; this makes building a quantum computer a formidable task. However, when built, quantum computers will break some of the pillars of our cybersecurity infrastructure.

The quantum threat to cybersecurity can be mitigated by deploying new cryptographic tools (both conventional and quantum) that are believed or known to be resistant to quantum attacks. Nonetheless, the transition to quantum-safe cryptography is a challenge itself, as it requires the development and deployment of hardware and software solutions, the establishment of standards, the migration of legacy systems, and more.

The urgency for any specific organization to complete the transition to quantum-safe cryptography for a particular cyber-system relies on three simple parameters:

- the *shelf-life time*: the number of years the data must be protected by the cyber-system;
- the *migration time*: the number of years to migrate the system to a quantum-safe solution;
- the *threat timeline*: the number of years before the relevant threat actors will be able to break the quantum-vulnerable systems.

If the threat timeline is shorter than the sum of the shelf-life time and of the migration time, then organizations will not be able to protect their assets for the required years against quantum attacks. A better understanding of the threat timeline provides information on the time available to safely perform the transition to post-quantum cyber-systems.

Assessing the quantum threat timeline is very challenging because of the scientific and engineering obstacles involved in building a working quantum computer. Experts generally acknowledge that they still do not know when we will have quantum computers that can threaten cyber-systems. However, it would be very helpful to gain insights into the prospects of this threat becoming real in the short and medium term, into the rate at which progress is being made, and into the key milestones cyber-risk managers should pay attention to.

This study aims to provide such deeper insights into the threat timeline, by surveying an unprecedented breadth and depth of thought leaders with questions designed to help those managing the cyber-risk associated with quantum cryptanalysis.

We targeted a diverse set of 22 trusted thought leaders in key relevant areas of quantum science and technology. The respondent pool, from academia and industry, spans four continents. Employees of some major companies declined to take part in the survey at the time it was issued. In the future, as a greater proportion of activity in building scalable fault-tolerant computers likely moves to industry, it will be valuable to gain additional industry perspectives.

Expert opinions on the likelihood of the quantum threat to current public-key cryptosystems

The experts were asked to express their opinions about the development timeline for quantum computers. It is not surprising that opinions varied significantly, and several experts articulated the difficulty inherent in to making such predictions. Nonetheless, some valuable patterns emerged.

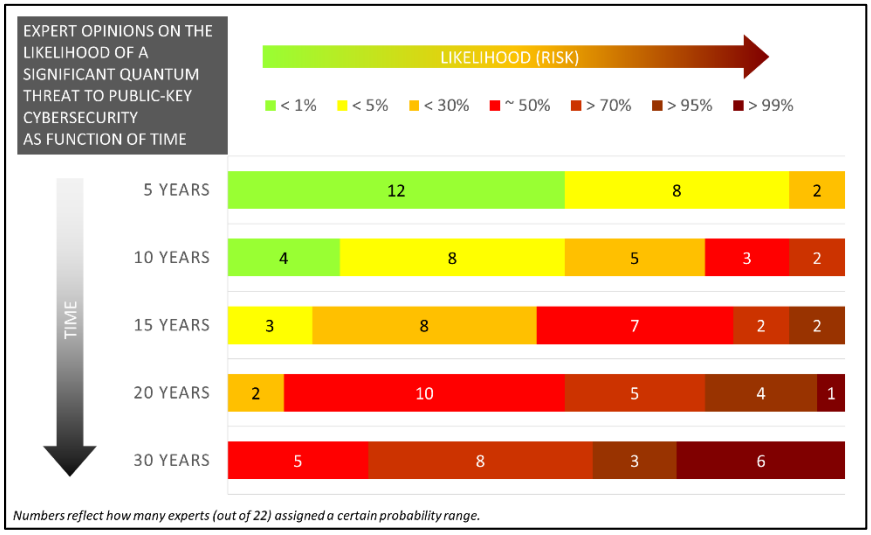

NEXT 5 YEARS: Most experts (12/22) judged that the threat to current public-key cryptosystems in the next 5 years is "<1% likely". The rest selected "<5%" (8/22) or "<30%" (2/22) likely, suggesting there is a small non-negligible chance of a short-term surprise.

NEXT 10 YEARS: Still more than half of the respondents (12/22) judged this was "<1%" or "<5%" likely, but 5/22 felt it was "about 50%" or ">70%" likely, suggesting that the quantum threat could very well become concrete in this timeframe.

NEXT 15 YEARS: Half (11/22) of the respondents indicated "about 50%" likely, or more likely, with two experts indicating a ">95%" likelihood.

NEXT 20 YEARS: About 90% (20/22) of respondents indicated "about 50%" or more likely, with 5/22 feeling it was ">95%" or ">99%" likely. This suggests that the chances of the quantum threat are more than even at the 20-year mark.

NEXT 30 YEARS: *All the experts* responded that the quantum threat has a chance of "about 50%" or more, with 17 out of 22 experts indicating that the quantum threat will be likely (">70%"), very likely (">95%") or extremely likely (">99%").

Expert opinions on the technical realization of quantum computers

A major challenge in building a quantum computer is that of creating reliable fundamental components, so-called physical qubits, whose number can be scaled while maintaining control and quality. In this respect, the experts indicated that the most promising physical platform for the realization of a cryptographically relevant quantum computer is presently offered by superconducting systems, followed relatively closely by trapped ions, and with several other physical implementations having significant potential.

A very important step forward will be the experimental demonstration that error-correcting schemes improve the reliability of so-called logical qubits as compared to physical qubits. For this to happen, it must be possible to prepare, manipulate, and measure the underlying physical qubits well enough. How 'well' this 'well enough' needs to be depends on the best known error-correcting schemes, which may themselves be superseded by new and better schemes.

Another milestone will be the demonstration of so-called "quantum supremacy", that is of the ability for a quantum device to perform some computation that would be practically impossible, even for the most powerful classical supercomputer, independent of the usefulness of such computation. While the achievement of quantum supremacy will not necessarily lead to decisive progress towards a cryptographically relevant quantum computer, it will signify having achieved a relatively high level of control on a relatively large number of physical qubits, which is a necessary ingredient for quantum computing. The experts agreed that this milestone is likely to be passed in the next couple of years.

From the threat timeline to the migration timeline

The expert opinions collected in our survey, and summarized in this report, offer unique insight into the quantum threat timeline. Depending on its own specific shelf-life times and migration times, each organization will have a longer or shorter time at its disposal to implement post-quantum cryptographic solutions. The Global Risk Institute and evolutionQ Inc. have already made available a quantum risk assessment methodology for taking estimates of the threat timeline and assessing the overall urgency of taking action (Mosca & Mulholland, 2017).

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

Contents

1 Introduction / background

1.1 Quantum computing

Quantum mechanics is our best description of the inner workings of nature. The framework of quantum mechanical laws 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. On the other hand, classical mechanics provides a great deal of descriptive and predictive power at the level of macroscopic objects. The effectiveness of classical mechanics is partly explained by the fact that quantum phenomena are more directly manifest at a microscopic scale than at the macroscopic, everyday-life scale, because of the magnitude of relevant physical constants. Another reason for the approximate validity of classical physics is that quantum phenomena are inherently fragile: the uncontrolled interaction of a quantum system with its environment tends to 'wash out' quantum features, a process often described as *decoherence*. This point is of the utmost importance when we consider that quantum computing is about preserving and controlling quantum behaviour at a level and with a precision that has no precedence in human history.

Quantum computing (Nielsen & Chuang, 2002) was born from the recognition that, since information must be stored and processed in physical systems, and since physical systems are fundamentally quantum, it might be possible to process information quantumly and even extend the definition of what information is. The basic unit of quantum information is the quantum bit, or *qubit*. While a standard *bit* can store a binary value, either 0 or 1, a qubit may store a *superposition*—technically, a linear combination—of 0 and 1. By exploiting quantum features, quantum computers will be able to solve computational tasks much faster than standard—for comparison, referred to as "classical"—computers, to the extent of being able to solve problems that are practically intractable by classical computers.

1.2 Quantum threat to cybersecurity

There are mathematical problems whose intractability by classical computers is at the basis of the security of several protocols for public-key cryptography, like the Rivest–Shamir–Adleman (RSA) cryptosystem (Rivest, Shamir, & Adleman, 1978), which is based on the difficulty of finding the prime factors of large numbers.

Such protocols would be broken by the existence of quantum computers, through the implementation of, e.g., the factoring algorithm by Shor (Shor, 1999). The ability of a quantum computer to search through a solution space with 2^n values (i.e., all the possible combinations of n bits) in roughly $2^{n/2}$ steps (Grover, 1996) would also weaken symmetric-key cryptography.

The quantum threat can be mitigated by deploying new cryptographic tools (both conventional and quantum) that are believed or known to be resistant to quantum attacks. Unfortunately, the transition to quantum-safe cryptography is a demanding and delicate task in itself (Mosca M. , 2013).

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 to 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.

If $T_{\text{SHELF-LIFE}} + T_{\text{MIGRATION}} > T_{\text{THREAT}}$, then organizations will not be able to protect their assets for the required $T_{\text{SHELF-LIFE}}$ years against quantum attacks. Organizations need to assess $T_{\text{SHELF-LIFE}}$ and T_{THREAT}, and the difference $(T_{\text{MIGRATION}})^{\text{MAX}} = T_{\text{THREAT}} - T_{\text{SHELF-LIFE}}$ is the **maximum available migration time**, that is, the maximum amount of time they have at disposal to safely realize the transition. Attempting to rush this migration will likely lead to serious security risks that can be exploited even without quantum computers, arising from gaps and omissions, design flaws, and implementation errors.

While the security shelf-life $T_{\text{SHELF-LIFE}}$ is generally established a business decision or a regulatory requirement, assessing the threat timeline T_{THREAT} is very challenging, because of the scientific and engineering obstacles involved in building a large-scale working quantum computer. While such challenges may delay the deployment of quantum computers for a long time, it also means that technical progress and scientific / engineering breakthroughs may suddenly speed up such development. It is also a matter of investments: the more resources are directed towards the development of quantum technologies—and of quantum computers, in particular—the sooner these will become a reality. Investments in the area have grown enormously in recent times (see also Section 1.3.4).

1.3 Realization of quantum computers

In the same vein in which classical information, intended in terms of bits, can be encoded in various classical physical systems—e.g., whether a switch, or a light bulb, is on or off—so quantum information can be encoded and processed in many different quantum physical systems. The latter include, e.g., quantum spins, or the polarization of quanta of light—photons. All physical realizations of quantum computers must deal with the issue of the 'fragility' of quantum properties, due to the inevitability of uncontrolled interaction with the environment leading to decoherence. Most crucially, it is difficult to have, at the same time, good preparation, control, and measurement of a physical system—all steps needed to perform computations—and isolation from the environment.

Apart from the issue of the physical realization (see Section 1.3.2), there are various *models* of computation. While many models are known to be computationally equivalent—roughly speaking, they allow one to solve the same class of problems with similar efficiency—each may offer different insights into the design of solution algorithms, or provide an advantage in / be more suitable for a particular physical realization. One such model is the *circuit* model or *gate* model, where transformations are performed on single and multiple qubits in sequences of so-called gates. From the perspective of

[1] These parameters have respectively been called also *x*, *y*, *z* in literature; see e.g., (Mosca M. , 2013).

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

In order to have universal quantum computation—that is, the ability to perform arbitrary computations—in the circuit model, it is enough to be able to realize a finite set of *universal gates* into which the computation can be decomposed. Such set necessarily includes at least a multiqubit—typically a two-qubit—gate.

Historically, the following criteria—part of a larger set of desiderata—listed by DiVincenzo in (DiVincenzo, 2000) and hence 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.*

One important parameter is the *fidelity* with which gates are applied, that is, to what extent the transformation that is actually implemented is close to the intended one. A related parameter is the physical error rate with which gates are applied. In a sense, this parameter is the 'opposite' of fidelity. Most research groups use either the "fidelity" or the "error rate" when characterizing the gate quality.

1.3.1 Error correction and fault tolerance

Errors in the manipulation of (quantum) information and decoherence may be reduced by improving the physical implementation, including qubit control, but cannot be eliminated entirely. A reliable computation can, nonetheless, be achieved by employing error-correction strategies. These correspond to encoding *logical* qubits into multiple *physical* qubits, so that errors 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—the so-called *fault-tolerance threshold*—then it is possible to devise 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.

Surface codes (Fowler, Mariantoni, Martinis, & Cleland, 2012) are currently among the leading candidates for large-scale quantum error correction. A single logical qubit is encoded into a square array of physical qubits. A detection & correction algorithm must be run at regular intervals in order to track the propagation of physical qubit errors and correct them to prevent logical errors. Another type of code, the *color code* (Bombin & Martin-Delgado, 2006), is a generalization of surface codes that provides the error-protection of the surface code with increased ease in logical computation, at a price of less efficient detection & correction algorithm. More details on these codes can be found in the Appendix.

1.3.2 Physical realizations

The various physical implementations of quantum computers each have advantages and disadvantages, e.g., related to the typical decoherence time, to how fast and how precisely gates can be applied, to how easy it is to scale to more and more qubits. The following are some physical realizations:

- Quantum optics, meaning that information is stored and manipulated in states of light; this includes, e.g., polarization states or photon-number states.
- Superconducting systems, meaning that information is stored and manipulated in electric circuits that make use of 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 a non vanishing total electric charge) that are confined by electro-magnetic means.
- Quantum spin systems, meaning that information is stored and manipulated in the internal degree of freedom called quantum spin.

1.3.3 "Quantum supremacy"

"Quantum supremacy"[2] (Preskill, 2018) is the ability for a quantum device to perform some computation that would be practically impossible for classical computers, independent of the usefulness of such computation. Criteria for quantum supremacy are challenging to define. One reason is that one must prove that no classical means—including even the most powerful existing classical supercomputer, and the best possible classical algorithm—would allow one to perform the same computation in a 'reasonable' time. In addition, even if one was content with known—rather than 'possible'—algorithms, quantum supremacy can be considered a moving target: classical computers and known classical algorithms—also those meant to simulate quantum systems—improve over time.

This said, quantum supremacy is a natural goal for so-called *noisy intermediate-scale quantum (NISQ) systems* (Preskill, 2018). These are systems composed of tens to hundreds of physical qubits, of quality not high enough to allow full quantum computation but still potentially useful to outperform classical computers greatly in some tasks. It is generally expected that quantum supremacy is within reach, and it could be achieved within the next couple of years.

1.3.4 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 pool of respondents. This interest has been supported and boosted by several national and transnational

[2] This terminology is somewhat controversial because some people find that 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 complete superiority of quantum computers over classical computers in some strictly technical sense. *Note added:* Shortly after completing this report, statements that quantum supremacy has purportedly been achieved started to circulate, see, e.g., this Financial Times article. We note that, at the time of writing this note, such statements do not refer to an official communication from the interested parties, nor to a peer-reviewed article.

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 specializing in various aspects of quantum computing research have been established, often supported by venture-capital investments. A detailed description of such a flourishing quantum landscape is beyond the scope of this report, but it is important to stress the following: While the challenge to create a fully scalable, fault-tolerant quantum computer is enormous, the investments in the area have never been stronger.

2 Scope of this report

This document reports the results of a survey conducted by evolutionQ Inc. among 22 internationally leading experts on quantum computing research, who were asked to complete an online questionnaire on the state of development in this field. 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 (see Section 3.1) 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 has already been considered, see, e.g. (National Academies of Sciences, Engineering, and Medicine, 2019), later referred to as the "NAS report". Within its more-than-200 pages, the NAS report articulated an opinion on when quantum computers would threaten RSA-2048:

> *[…] it 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; furthermore, it does not consider that much has happened since the NAS report was compiled.

The present report differs from previous ones, 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 can not provide definitive answers, we aim to depict a much better picture of what people think about this question. Chances of 1%, 10% and 49% are meaningfully different shades of "unexpected". Do all experts agree, or is there a wide variance in opinions even amongst experts? What about 5 years, 15 years, 20 years? What is this timeline depending mostly on?
- Also, there are many fast-moving parts in the field, and much changes in just one year. What are risk managers supposed to do now, in 12 months, in 24 months, etc., with advice that back in 2018 a committee of experts thought a quantum attack on RSA-2048 by 2028 was "highly unexpected"? For example, since 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, 2019; Gheorghiu & Mosca, 2019), and other players have joined the quest to build quantum computers.
- The scope of our survey/report is much tighter and more focused than that of other reports.
- We ask specific questions individually to several leading researchers, and compile relevant statistics.
- We aim to track how the opinions of experts evolve over time. We will run a similar survey again in a year, with the potential to continue running it as long as the community finds it helpful.

3 Survey design and methodology

There was a range of non-trivial considerations in the phrasing of the questions. It was most important to understand the perspectives of the diverse range of people who would be asked to complete the survey, and how the target audience would interpret the questions and possible answers.

Most importantly, we wanted to avoid having the various respondents interpreting the questions very differently. E.g., 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?

Given the goals of our survey:

- We kept the questions focused on the issue of the implementation of fault-tolerant quantum computers that would be able to run quantum algorithms that could pose an actual threat to cryptosystems;
- We sought a range of relevant perspectives. 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;
- 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;
- Given the inherent uncertainty in the progress towards realizing a quantum computer, we gave respondents the opportunity to indicate how likely something was to happen in a relatively coarse-grained fashion, but still much more informative than what is available in previous reports.

Answering some of the questions was optional, specifically when it came to sharing details about personal research activities of the respondents and to providing more free-form opinions on the state of the field.

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

3.1 Questions

A list 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).

We posed a similar question about "quantum supremacy":

> *Please indicate how likely you think it is that a quantum computer / device will demonstrate so-called "quantum supremacy" (that is, the ability to perform some computation practically impossible for classical computers, including classical supercomputers, irrespective of the utility of such a computation) within the next 1 year, 3 years, 5 years, and 10 years,*

with potential answers about the likelihood following the same classification as above.

We asked the respondents to also provide an opinion on the best schemes for fault-tolerance, as well as on recent and near-future expected progress that has been and will be, respectively, key in the development of a quantum computer.

Finally, we asked the respondents to share some information about their own research—if willing to do so—and to express opinions about the general state of the field of quantum computing research.

3.2 Participants

We selected about 80 potential candidate respondents. From this list, we shortlisted about 30 people who 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. We contacted shortlisted candidates in the first wave of invitations, inquiring about their willingness to take part in our survey. Those who accepted were asked to complete the online questionnaire in about two weeks. Some candidate respondents simply 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 and legal advice. In order to achieve a target number of 20 or more respondents, we proceeded to invite more people. We contacted or attempted to contact 48 people in total.

In about a month and half, we were able to collect responses from 22 respondents (see

Appendix for a complete list). Figures 1-3 summarize classifications of our respondents in terms of:

- country where they work (Figure 1),
- kind of (primary) organization they belong to (Figure 3),
- kind of activity they lead (Figure 2).

The respondent pool comprised a diverse set of expertise and nationalities, and a mix of university and company researchers, representative of the diversity of the quantum computing community among its top players. We note that employees of some major companies declined to take part in the survey at the time it was issued. In the future, as a greater proportion of activity in building scalable fault-tolerant computers moves to industry players, it will be valuable to gain additional industry perspectives.

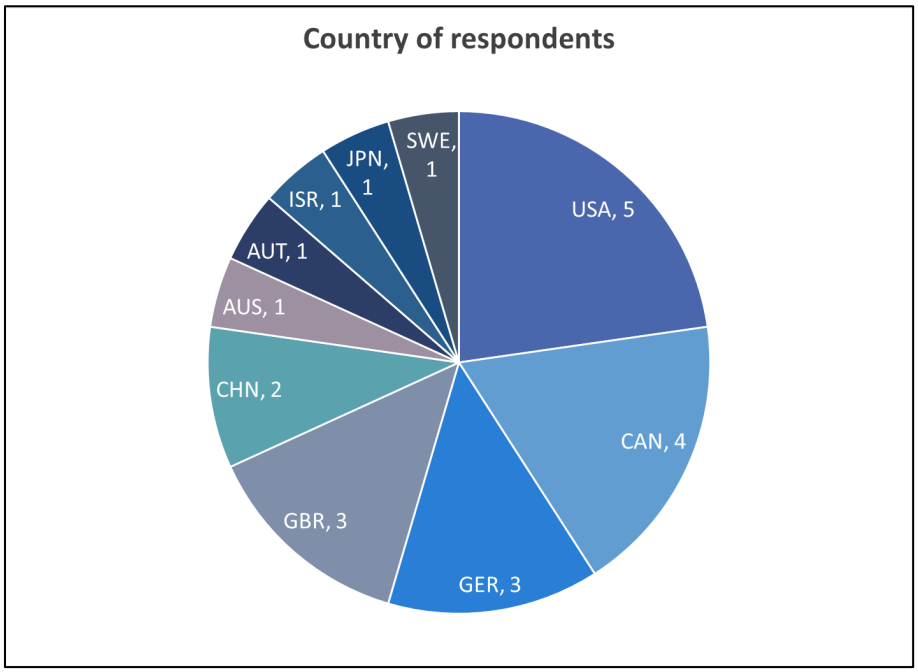

Figure 1: Our respondents constitute a very international mix, with higher representation from countries (like Canada, USA, and China) and geographical areas (like Europe) where the efforts to develop quantum computers and quantum technologies are very strong.

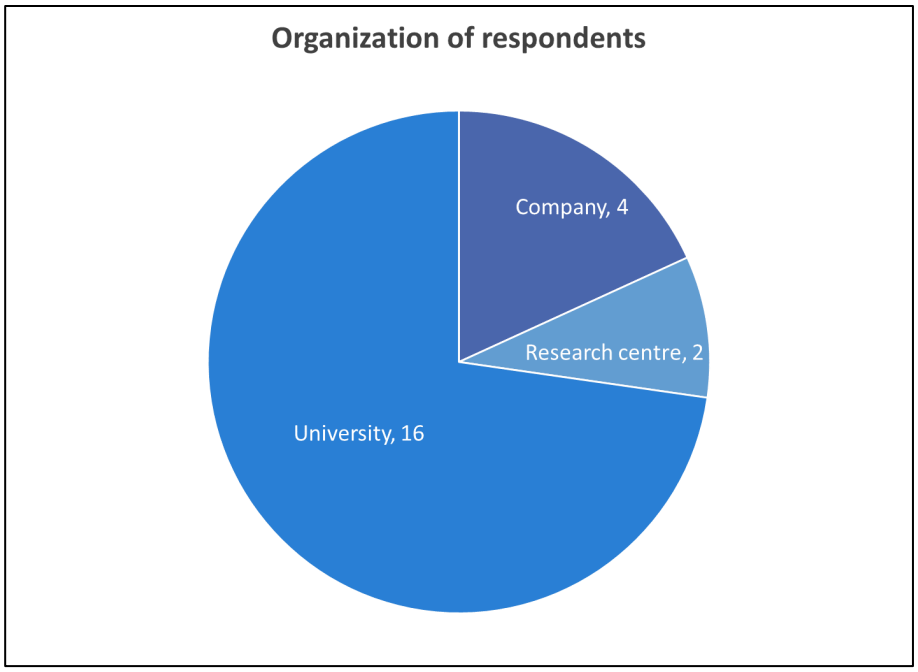

Figure 3: Most of the respondents work at universities, but some work at companies or research centres. Some researchers/academics may have some role in, or collaborate with, external companies; this is not illustrated in the above classification.

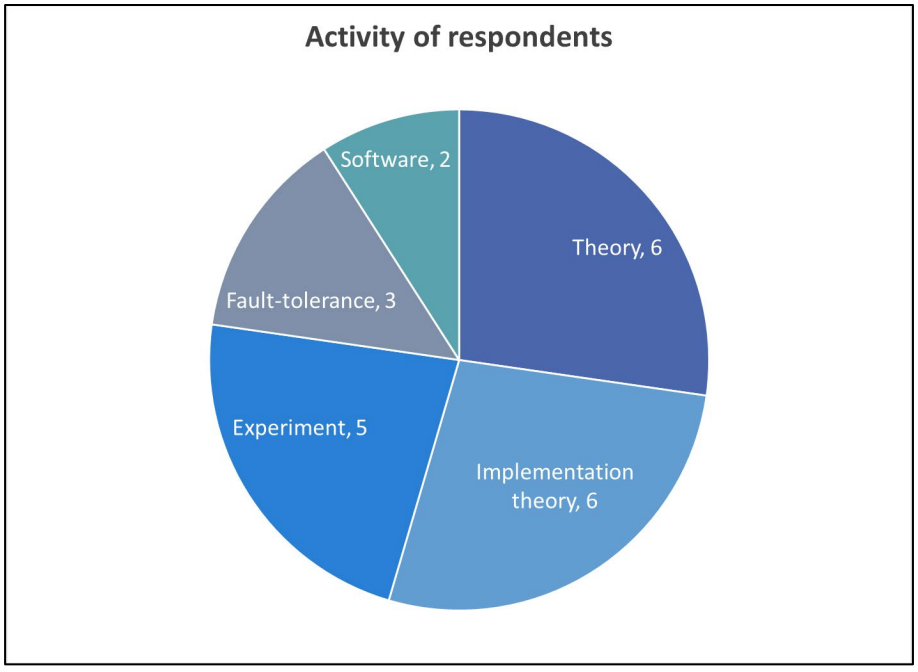

Figure 2: Our respondents cover a wide range of research activities. While the major division is between non-experimental research and experiment, research that is not directly experimental can be very different. E.g., implementation theory focuses on guiding, supporting, and, in general, facilitating experimemental effort. Respondents are classified under simply "theory" if their theoretical activity is not specificically related to experiments or implementations.

4 Survey results

The results of our survey reported here comprise:

- an aggregate analysis of the key responses about the quantum threat timeline;
- a list of key recent developments in the field of quantum computing research, as highlighted by the respondents;
- a selection of opinions about fault-tolerant schemes;
- a list of near-future (that is, approximately, by the end of 2020) developments that the respondents see as essential on the path to developing a fully scalable fault-tolerant quantum computer;
- a collection of notable remarks made by the respondents.

4.1 Aggregated analysis of responses

In the aggregated analysis of the responses we typically indicate how many of the respondents chose a specific answer among the many possible ones, when dealing with multiple choices. In general, on one hand one can appreciate the variability/spread in the answers while, on the other hand, general trends tend to emerge.

In assessing the timeline of the actual threat to RSA-2048, we provide a more detailed picture of the responses, also plotting the answers of single respondents (kept anonymous). This emphasizes the differences between 'optimistic' participants, who feel quantum computers are relatively close to becoming a reality—and, from the perspective of this report, a threat—and 'pessimistic' participants, who tend to believe that building a quantum computer will take a really long time.

We also find it interesting to separately consider the responses of those participants who are close/closer to experiments, a group that comprises both experimentalists and theorists who contribute to experiments or are in some way concerned with actual implementations—a kind of theoretical activity we refer to as *implementation theory.* Conceivably, such a group has a very informative vantage point when it comes to judging the hurdles of building a quantum computer in the lab.

4.1.1 Physical realizations

With respect to the physical realizations of quantum computers, we asked the respondents:

- to indicate the generic potential of several physical implementations as candidates for fault-tolerant quantum computing,
- to rank physical implementations for the specific goal of realizing a digital quantum computer with 100 *logical* (rather than physical) qubits in the next 15 years.

The responses indicate a fairly general consensus that the present leading platforms are superconducting systems and trapped ions. Daniel Gottesman wrote:

> *Right now, the main contenders I see are superconducting qubits, ion traps, and optical systems. Superconducting qubits are slightly ahead at the moment and have an easier path for manufacturing large numbers of qubits but ion traps probably have*

better prospects for lower noise rates and greater potential for long-range coupling of traps. Optical systems are a somewhat distant third.

One respondent wrote:

At the moment, superconducting qubits are, in my impression, leading when it comes to progress towards [quantum error correction] and fault tolerance. Trapped ions could very well catch up.

With respect to spin qubits, one respondent wrote:

Semiconductor spin qubits, while promising, are still far. For example, progress towards high-fidelity two-qubit gates have to be made.

It is worth mentioning that one respondent judged that the questionnaire should not have grouped all spin systems together, as they are very different and have very different potential.

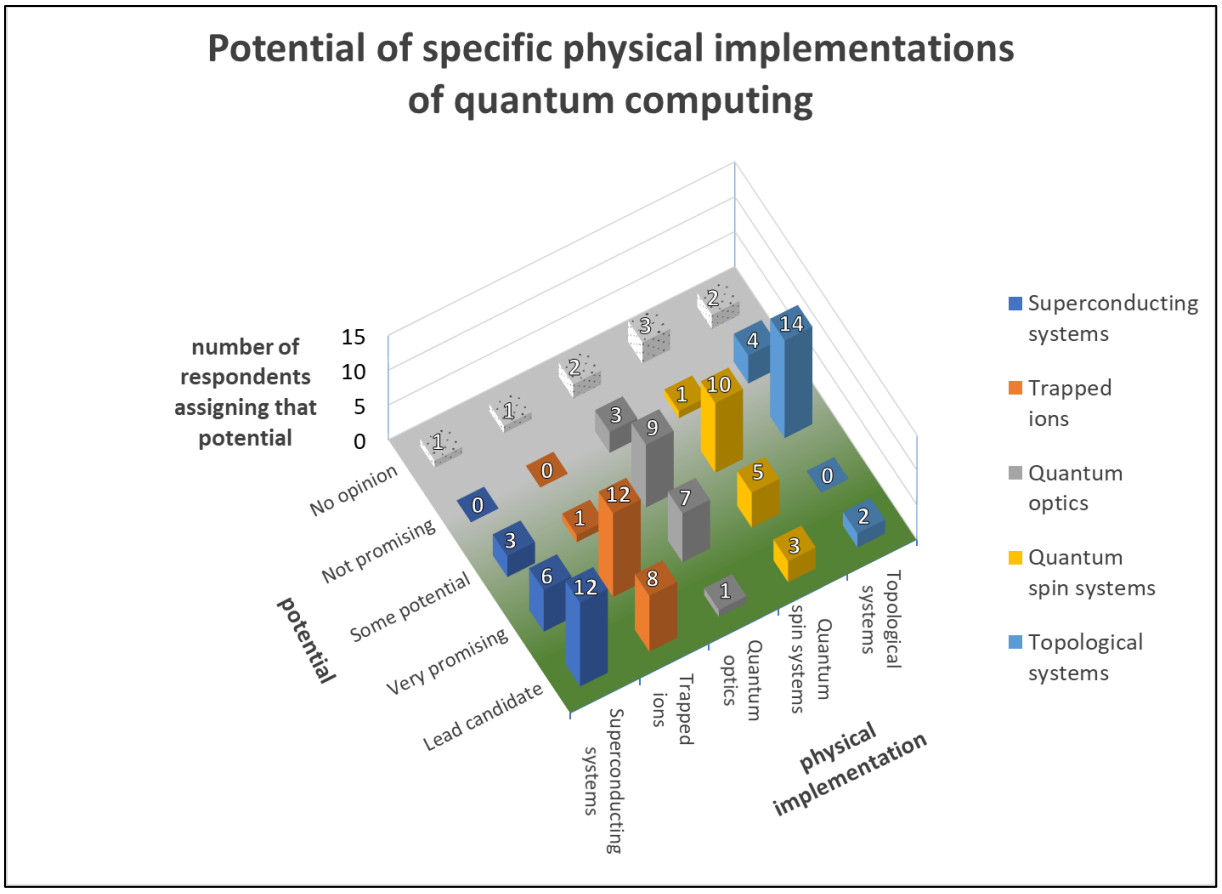

Figure 4: Superconducting implementations as well as ion-trap implementations are perceived as presently having an edge over other physical realizations.

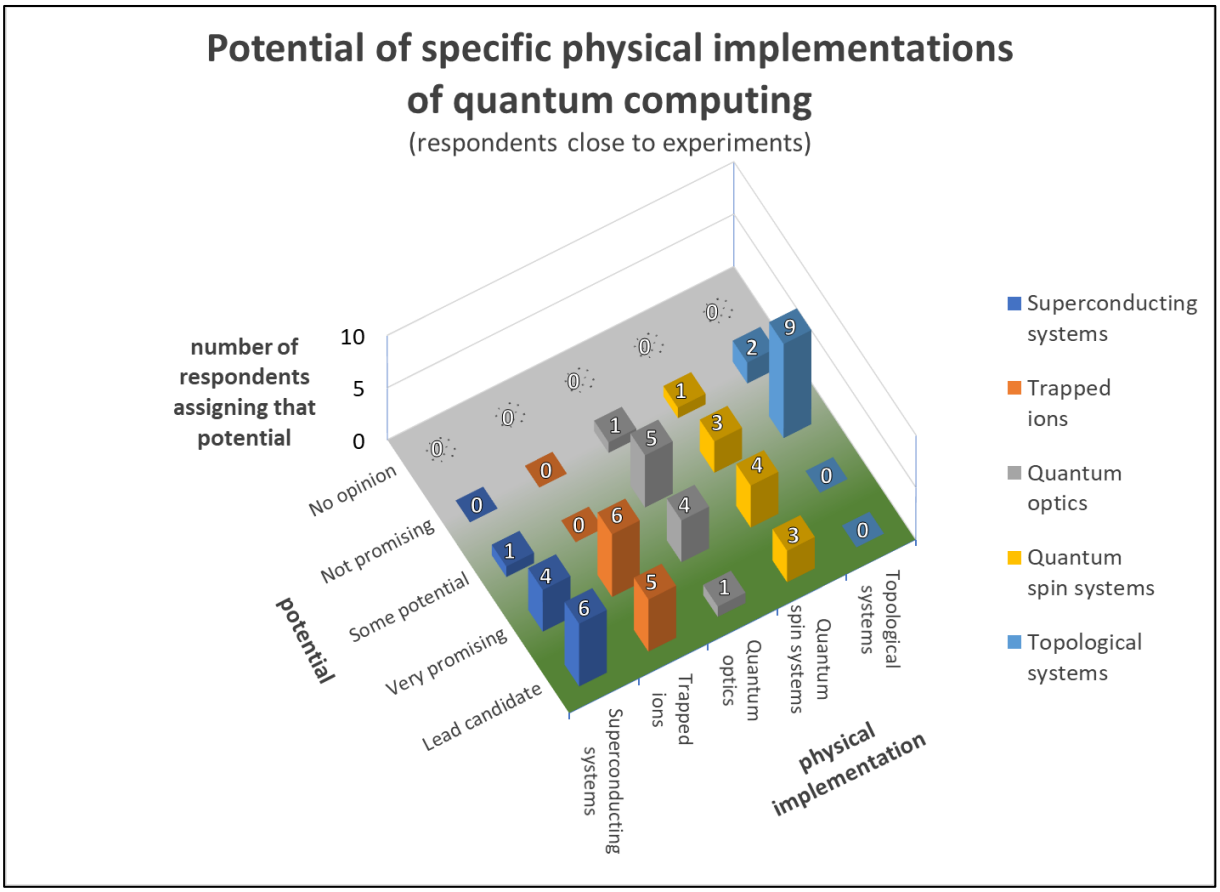

Figure 5: The respondents close to experiments have indicated that trapped ions and superconducting systems are very similar in their potential. Interestingly, among such practitioners spin systems rank somewhat higher than among the general pool of respondents, at least for our set of respondents.

Speaking of silicon spin qubits, Stephanie Simmons wrote:

> *I believe silicon will, in the long run, create the largest controllable quantum systems due to both its native quantum metrics as well as its commercial familiarity and deployability. […] The major bottleneck is the identification of a scalable system design: a design that ideally offers ultra-low-error multi-qubit non-interacting modules which can be efficiently entangled with one another. […] In silicon the constituent quantum ingredients are just about the best nature has to offer; what our community lacks is the recipe to make the most of them.*

Another respondent seems to agree with Simmons' take on silicon spin qubits:

> *I expect that progress […] will not be a steady increase of qubit numbers, but rather an explosion: until we can fabricate qubits in a proper foundry, we will keep pushing up to numbers perhaps around 10. Once the foundry processes deliver useful qubits, the numbers can escalate very quickly […] Ascertaining whether semiconductor*

foundry processes are amenable to produce good-quality spin qubits in silicon […] will require very significant investments of time, money and industrial resources.

With most (9/11) implementation-oriented respondents viewing topological implementations as having "some potential", such kind of implementations appear to be behind, both in terms of present development status and of respondents' expectations. One of them wrote:

As for topological implementations, it is not even clear when, if ever, the operation of a single qubit will be demonstrated.

Some caution towards topological systems is shared by Daniel Gottesman, who, nonetheless, emphasizes some key advantages such systems may have if realized:

[The] reason they are still worth considering is that if they can be made to work, it is possible that they can immediately jump to low error rates due to a degree of intrinsic fault tolerance.

Such potential advantage of topological systems towards fault-tolerance is further stressed by another respondent (see Section 4.2).

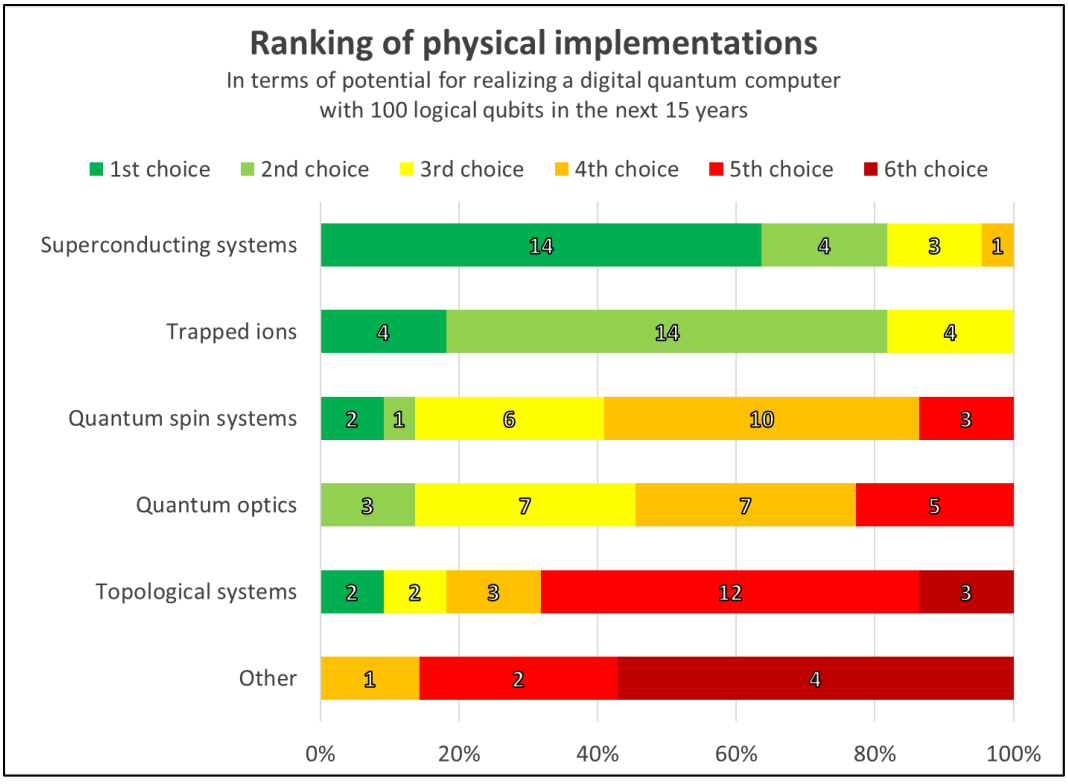

Figure 6: Superconducting systems and trapped ions are perceived as having an edge over other physical implementations not only in terms of generic potential for quantum computing (see Figure 4), but also when it comes to the more specific goal of developing a digital quantum computer with 100 logical qubits within the next 15 years. Numbers indicate how many respondents assigned that ranking. Not all respondents weighed in on an "other" system.

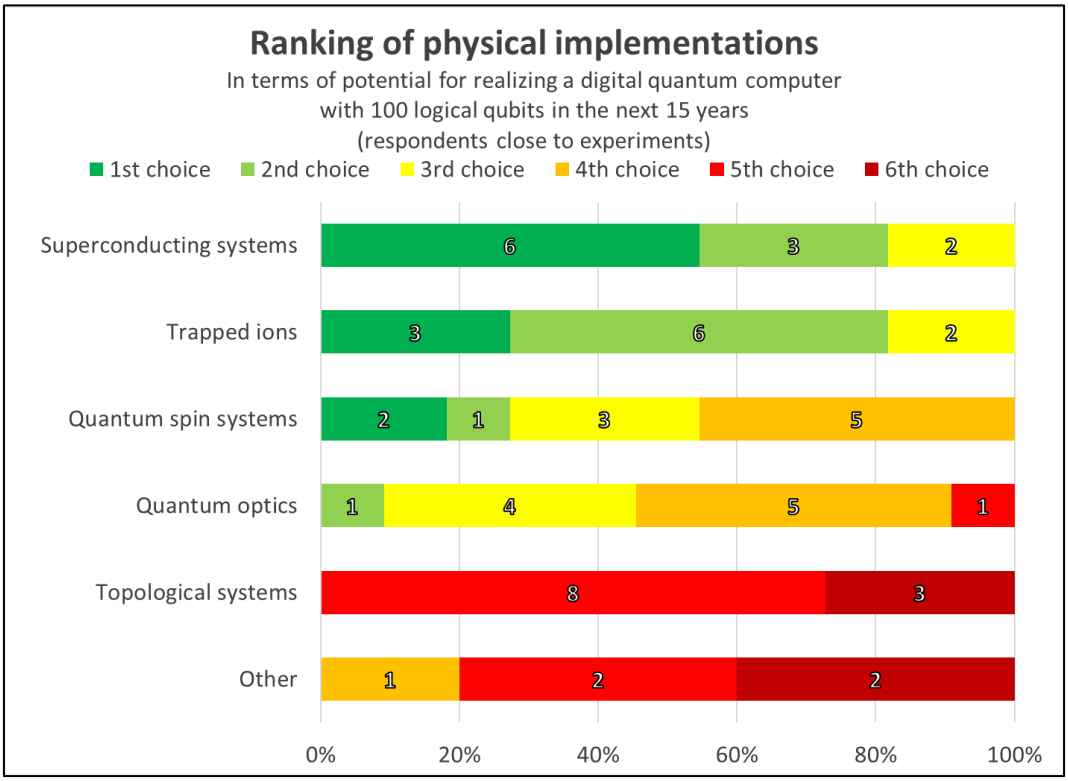

Figure 7: Quantum spin systems are "closer" to the two leading physical implementations (superconducting systems and trapped ions) if we restrict the pool of respondents to those closer to experiments, at least for our set of respondents. Numbers indicate how many respondents assigned that ranking. Not all respondents weighed in on an "other" system.

Trapped neutral atoms and Rydberg atoms were among physical implementations that we did not list explicitly, but were indicated by the respondents as "other". Some respondents think such systems have substantial potential, particularly since it is relatively easy to realize many physical qubits.

4.1.2 Quantum supremacy

There seems to be a general consensus that the demonstration of quantum supremacy is quite close. Most respondents estimate the chance that it will happen within a year to be about 50% or higher. Bill Coish, together with other respondents, thinks that it will be claimed by the end of 2020, but warns that such a claim will "likely [be] controversial".

The words of another respondent help to make clearer where the controversy may come from:

> *[T]he theory behind quantum supremacy in noisy devices is still far behind, and at this point we only have theoretical understanding of demonstrating supremacy […] with extremely accurate devices; however, I personally believe that this will be overcome.*

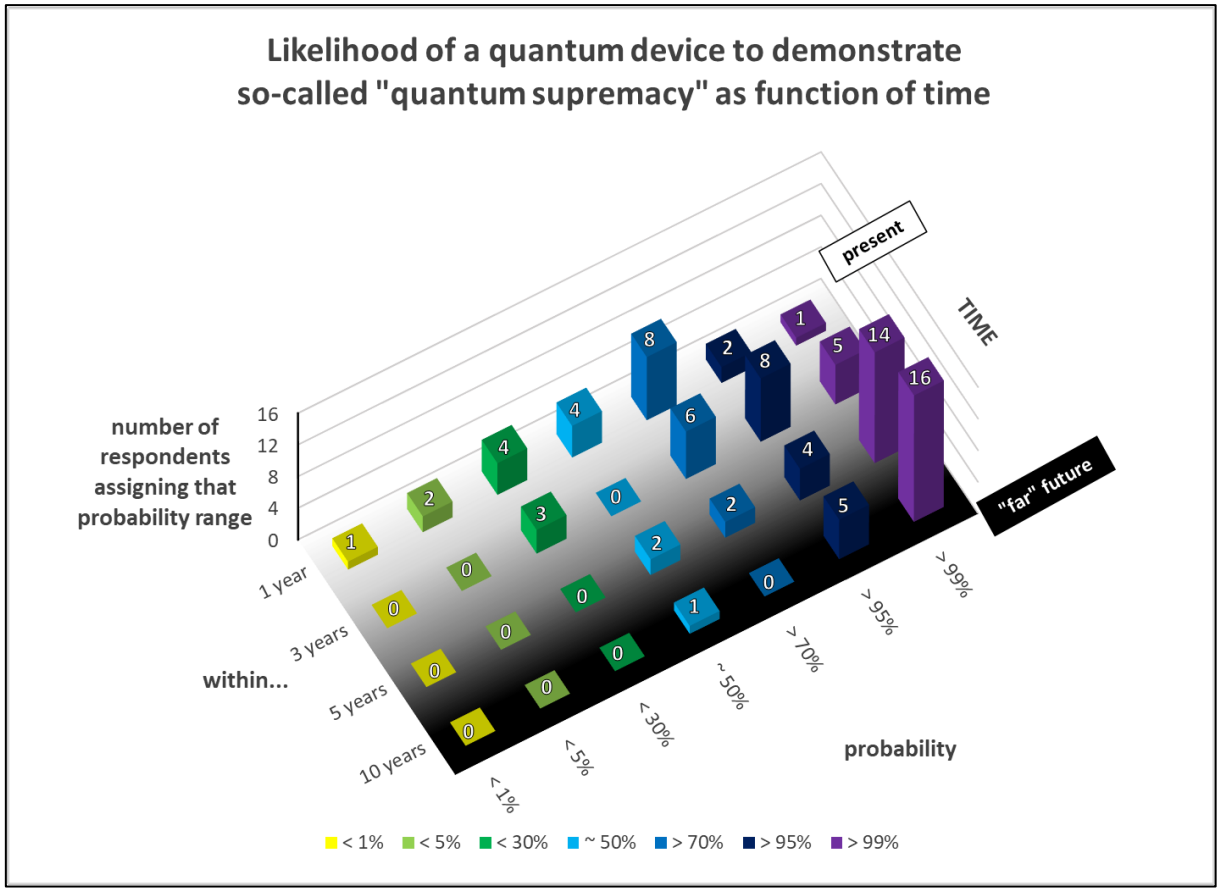

Figure 8: We asked our respondent to assign a likelihood to the demonstration of quantum supremacy. The consensus appears to be that this will happen relatively soon.

Several respondents have mentioned that achieving quantum supremacy is among the next essential steps towards evolving to a fault-tolerant quantum computer; one respondent argues:

> *"[Q]uantum supremacy" although not directly aimed at [fault tolerance], needs to be possible if we are to make complex 50+ qubit systems.*

Another respondent wrote:

> *[E]ven though the experimental platform [demonstrating quantum supremacy] will be intrinsically noisy, that should give us important insights about the challenges on the way to fault tolerance.*

With respect to quantum supremacy as a 'precursor' of full quantum computers, one respondent warns:

> *I believe that the NISQ [(noisy intermediate-scale quantum)] era is a very different era than the era of full-fledged quantum computers; I think quantum supremacy with noisy low-depth quantum circuits will probably be demonstrated rather soon in NISQ devices. However, supremacy demonstrations, if achieved, will be done for certain*

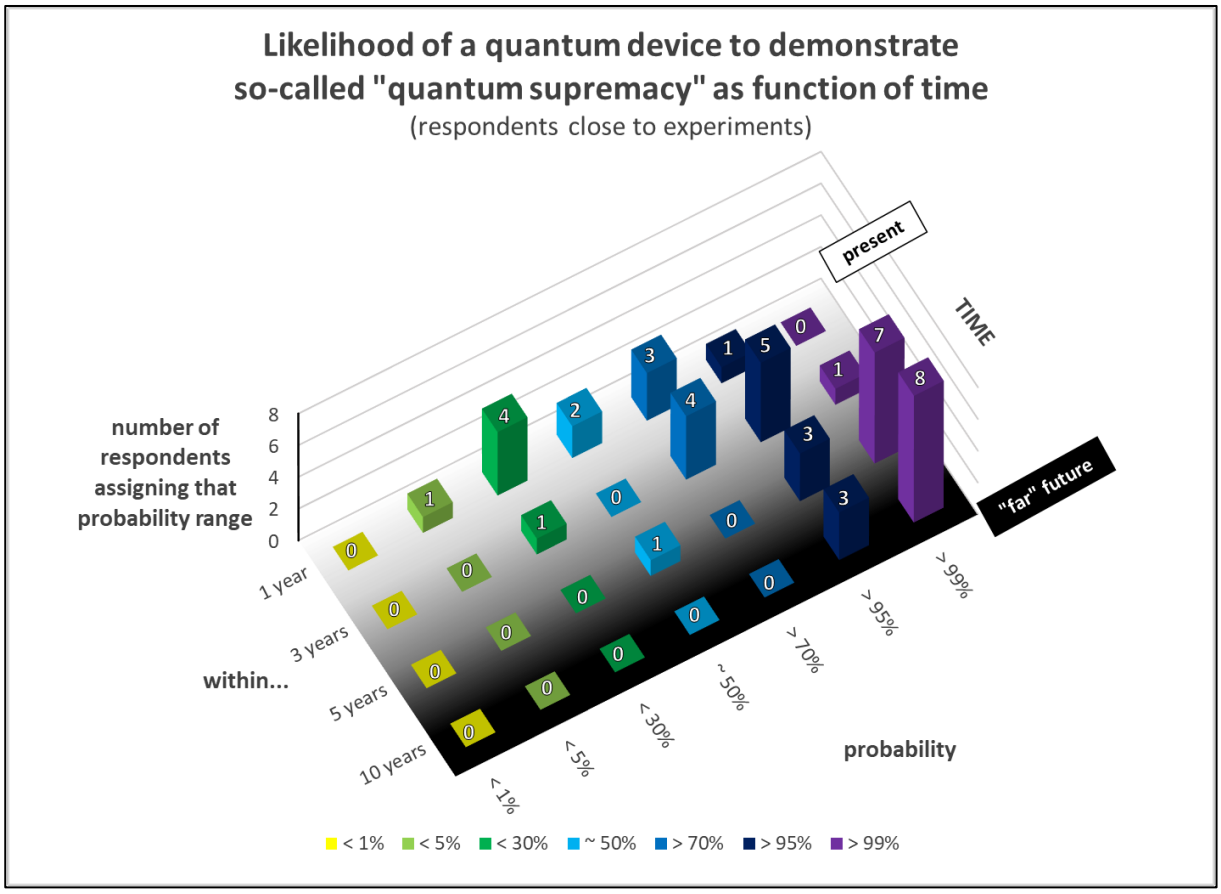

Figure 9: Respondents closer to experiments have slightly more pessimistic expectations about the timeline towards quantum supremacy.

> *sampling tasks which are very far from the true goal of quantum computers. Once we have demonstrated quantum supremacy, it will be a major quantum leap towards truly controllable quantum computers which can operate in the presence of noise. This phase will require a big change in attitude of industry; the reason is that at that point, we will have no way around it: we will have to handle the challenge of error correction.*

Although not necessarily statistically relevant, responses of those experts close to experiments seem to show that they are slightly more pessimistic than the rest of respondents about the timeline for achieving quantum supremacy.

4.1.3 Quantum factoring

Given our interest in gaining insights on the quantum threat timeline, the most interesting responses were those about estimating 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). Recent estimates on the practical requirements to achieve such a feat, taking into account the imperfections of physical implementations, were presented in (Gheorghiu &

Mosca, 2019) and in (Gidney & Ekerå, 2019) (the second author of the latter paper is part of our pool of respondents).

In Figure 10 and Figure 11 we provide a graphical representation of the estimates made by the individual respondents. It is possible to appreciate the ample range of opinions, with some experts being very optimistic and some others being relatively pessimistic about the rate of development of quantum computers. E.g., three respondents judge that within 15 years there is still less than a 5% chance that such a computer will exist, while two other experts judge it very likely (more than 95% chance). These are 'extreme' opinions, and most of the experts estimate some intermediate likelihood. Also, respondents closer to experiments appear to be slightly more pessimistic.

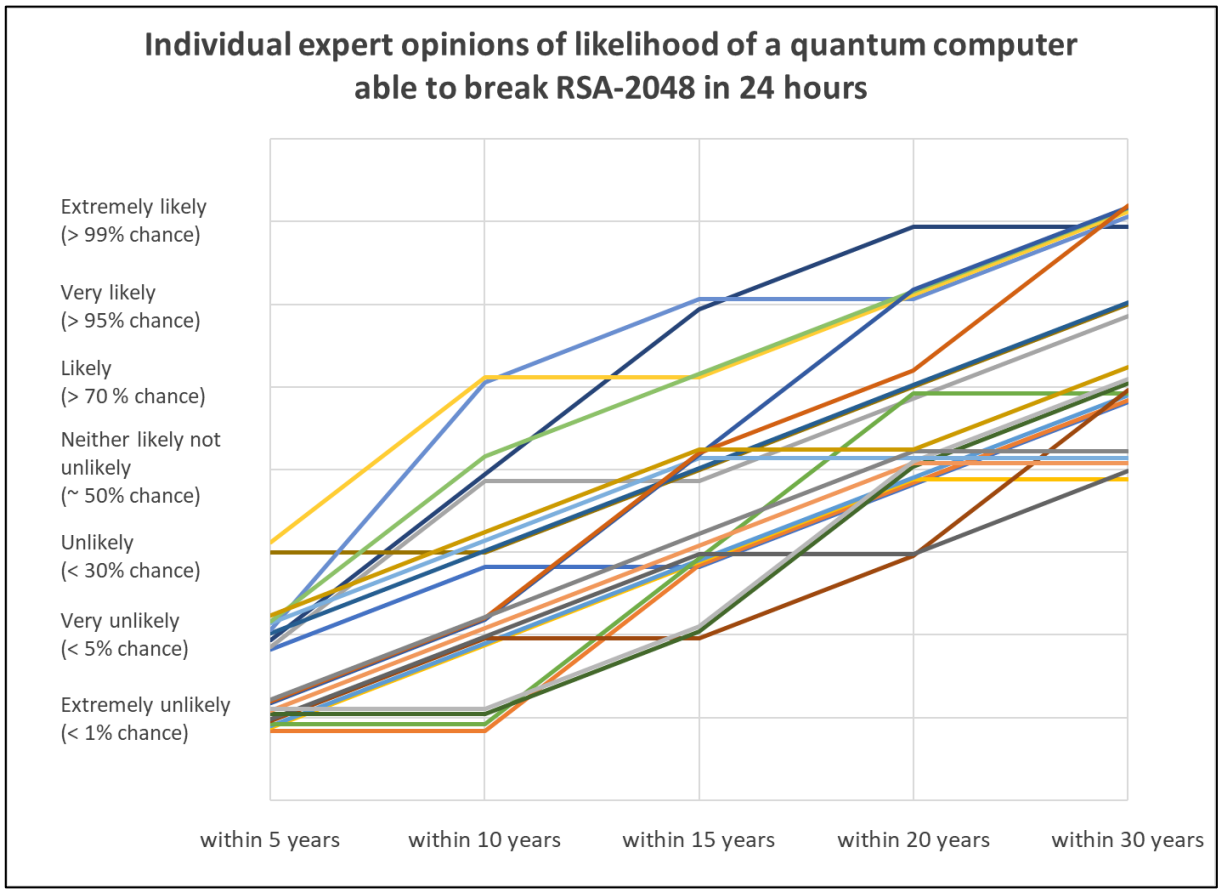

Figure 10: Opinions of individual experts about the likelihood of having a quantum computer able to factorize a 2048-bit number—that is, able to break RSA-2048—in, at most, 24 hours. Each line represents the opinion of one expert about the evolution of such a likelihood in time.

The experts' responses are analyzed in a more coarse-grained fashion in Figure 12 and Table 1, and in Figure 13 and Table 2. Despite the great variability of the responses, some valuable patterns emerged.

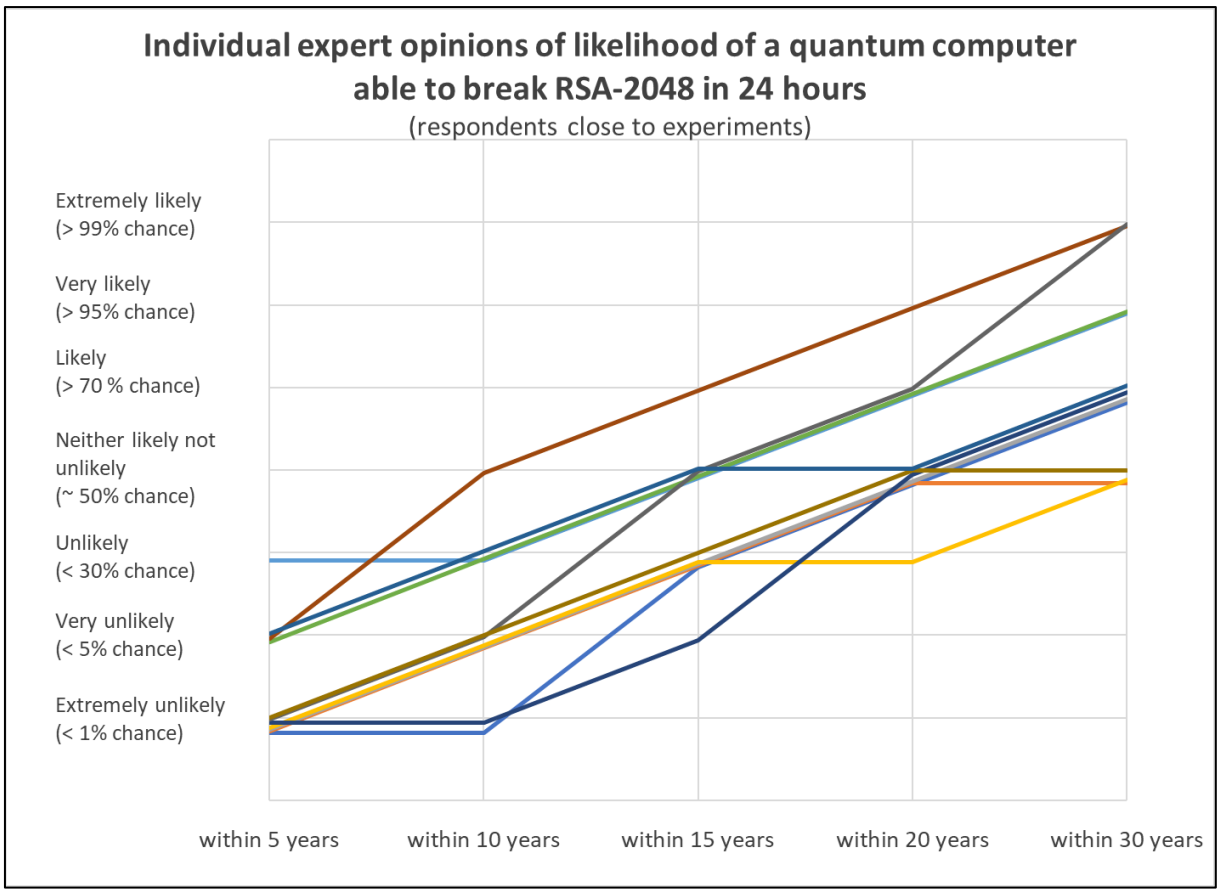

Figure 11: Opinions of only the individual experts close to experiments about the likelihood of having a quantum computer able to factorize a 2048-bit number—that is, able to break RSA-2048—in, at most, 24 hours. Each line represents the opinion of one expert.

NEXT 5 YEARS: Most experts (12/22) judged that the threat to currently public-key cryptosystems in the next 5 years is "<1% likely". The rest selected "<5%" (8/22) or "<30%" (2/22) likely, suggesting there is a small non-negligible chance of a short-term surprise.

NEXT 10 YEARS: Still more than half of the respondents (12/22) judged this was "<1%" or "<5%" likely, but 5/22 felt it was "about 50%" or ">70%" likely, suggesting that the quantum threat could very well become concrete in this timeframe.

NEXT 15 YEARS: Half (11/22) of the respondents indicated "about 50%" likely, or more likely, with two experts indicating a ">95%" likelihood.

NEXT 20 YEARS: About 90% (20/22) of respondents indicated "about 50%" or more likely, with 5/22 feeling it was ">95%" or ">99%" likely. This suggests that the chances of the quantum threat are more than even at the 20-year mark.

NEXT 30 YEARS: *All the experts* responded that the quantum threat has a chance of "about 50%" or more, with 17 out of 22 experts indicating that the quantum threat will be likely (">70%), including almost half feeling that it was very likely (">95%") or extremely likely (">99%").

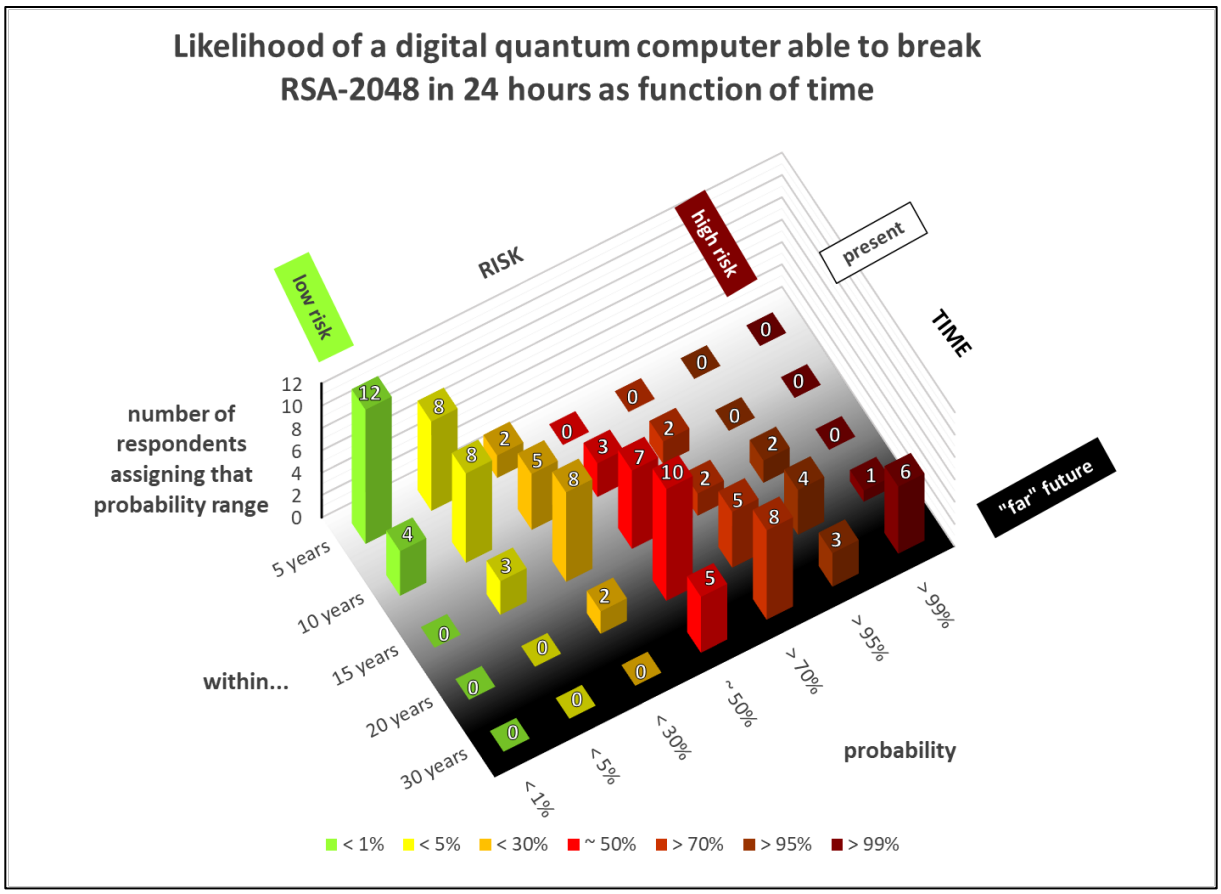

Figure 12: Number of respondents that have indicated a certain likelihood that a quantum computer able to factorize a 2048-bit number in less than 24 hours will be built within a certain period of time. See also Table 1 for data and informal 'likelihood' wording.

	Period of time				
How likely	***5 years***	***10 years***	***15 years***	***20 years***	***30 years***
Extremely unlikely (< 1% chance)	12 (55%)	4 (18%)	0 (0%)	0 (0%)	0 (0%)
Very unlikely (< 5% chance)	8 (36%)	8 (36%)	3 (14%)	0 (0%)	0 (0%)
Unlikely (< 30 % chance)	2 (9%)	5 (23%)	8 (36%)	2 (9%)	0 (0%)
Neither likely nor unlikely (about 50% chance)	0 (0%)	3 (14%)	7 (32%)	10 (45%)	5 (23%)
Likely (> 70 % chance)	0 (0%)	2 (9%)	2 (9%)	5 (23%)	8 (36%)
Very likely (> 95% chance)	0 (0%)	0 (0%)	2 (9%)	4 (18%)	3 (14%)
Extremely likely (> 99% chance)	0 (0%)	0 (0%)	0 (0%)	1 (5%)	6 (27%)
Total number of respondents (percentage)	*22 (100%)*	*22 (100%)*	*22 (100%)*	*22 (100%)*	*22 (100%)*

Table 1: Number (percentage) of respondents that have indicated a certain range of likelihood that a quantum computer able to factorize a 2048-bit number in less than 24 hours will be built within a certain period of time. See Figure 12 for an intuitive graphical representation.

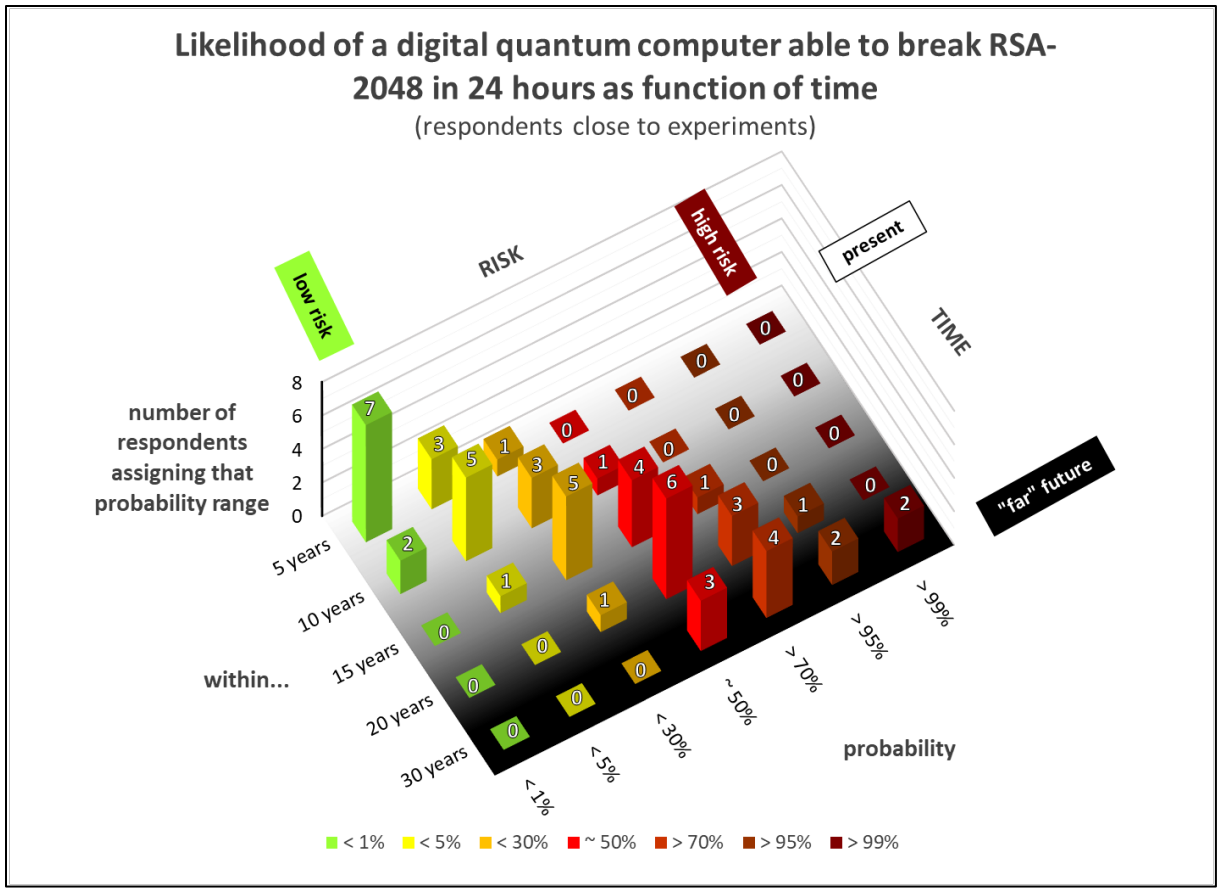

Figure 13: Number of respondents who are close to experiments and have indicated a certain likelihood that a quantum computer able to factorize a 2048-bit number in less than 24 hours will be built within a certain period of time. See also Table 2 for data and informal 'likelihood' wording.

	Period of time				
How likely	***5 years***	***10 years***	***15 years***	***20 years***	***30 years***
Extremely unlikely (< 1% chance)	7 (64%)	2 (18%)	0 (0%)	0 (0%)	0 (0%)
Very unlikely (< 5% chance)	3 (27%)	5 (45%)	1 (9%)	0 (0%)	0 (0%)
Unlikely (< 30 % chance)	1 (9%)	3 (27%)	5 (45%)	1 (9%)	0 (0%)
Neither likely nor unlikely (about 50% chance)	0 (0%)	1 (9%)	4 (36%)	6 (55%)	3 (27%)
Likely (> 70 % chance)	0 (0%)	0 (0%)	1 (9%)	3 (27%)	4 (36%)
Very likely (> 95% chance)	0 (0%)	0 (0%)	0 (0%)	1 (9%)	2 (18%)
Extremely likely (> 99% chance)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	2 (18%)
Total number of respondents (percentage)	*11 (100%)*	*11 (100%)*	*11 (100%)*	*11 (100%)*	*11 (100%)*

Table 2: Number (percentage) of respondents close to experiments that have indicated a certain range of likelihood that a quantum computer able to factorize a 2048-bit number in less than 24 hours will be built within a certain period of time. See Figure 13 for an intuitive graphical representation

To further summarize the results of the survey in a meaningful way, we have transformed the expert responses into an average cumulative probability distribution, assigning an actual probability to each likelihood classification. There is a degree of arbitrariness in such an assignment, given the broad likelihood ranges which were given as options to choose from. We have taken the most conservative approach, by considering both the highest possible probability and the lowest possible one compatible with the ranges indicated by the responses. The resulting probabilities for each expert were then simply averaged. See more details about the method used to produce Figure 14 in the Appendix.

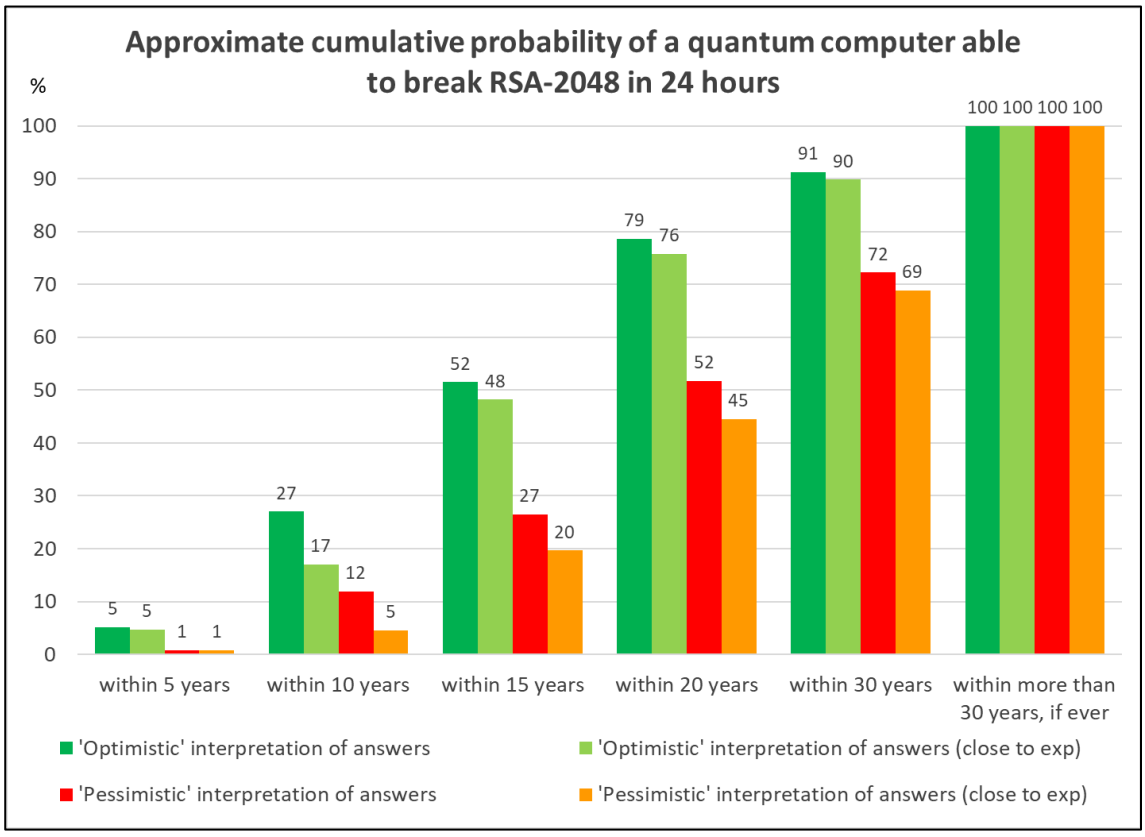

Figure 14: Cumulative probability for the creation of a quantum computer able to break RSA-2048 in 24 hours in a certain number of years in the future, based on the data of Table 1. The series correspond to 'optimistic' and 'pessimistic' interpretations of the responses, given the probability/likelihood intervals given. We have considered, separately, the answers of those respondents closer to experiments. See

The Figure 14 chart should be interpreted cautiously, but it provides insight into the expert opinions, and, hence, on the quantum threat timeline itself. For example, even in a 'pessimistic' interpretation of responses (lowest compatible probability), the approximate probability of a disruptive quantum threat is already 12% in the next 10 years, and growing steadily in the years that follow; the estimated probability is more than 50% by the 20-year mark, even in the pessimistic interpretation. In Figure 14 we have also plotted the cumulative probability distribution obtained by averaging the probabilities associated with the opinions of only the respondents close to experiments. As observed before, those respondents tend to be more pessimistic about the timeline for the development of a quantum computer, a tendency more evident in their answers about the medium term (i.e., 10-20 years).

4.2 Fault-tolerant schemes

Several respondents think that the ultimate choice of the error-correcting mechanism will depend on the underlying physical platform, and different error correction codes may be (and may need to be) combined. Stephanie Simmons wrote:

> *Different physical layers will require a range of approaches to be able to meet the overall goal of fault-tolerance. […] Assuming the physical qubit layout can be made to mitigate the damage done by dynamic crosstalk at the physical layer, certain logical-level fault-tolerance schemes may work more naturally than others. I anticipate that logical-level schemes, such as surface codes, colour codes, magic-state distillation, and higher-dimensional codes will be combined across and within different logical levels in order to obtain a universal gate set fault-tolerantly. I don't suspect a single such scheme will be used in isolation: concatenated coding shows substantial promise.*

Nonetheless the responses by the experts point to surface-code based architectures (or variants such as the color code) as the leading candidates for large-scale, fault-tolerant quantum error correction. One respondent wrote:

> *The surface code seems to remain the most plausible scheme at the moment. […] I am acutely aware of the manufacturing challenges involved in the construction of a large array of qubits in two dimensions. What gives me hope, however, is the incontrovertible fact that the semiconductor industry does have the capability to place a billion nanoscale silicon transistors in a 2D grid. Our challenge is to convert such transistor structures into qubits with sufficient gate fidelities and controllable interactions.*

Some of our respondents point out that improvements (e.g., on surface codes) are likely to be found. Dave Bacon wrote:

> *The most promising approach right now is to use surface codes plus ["T-state factories"] in a two-dimensional architecture of superconducting qubits. However, I also expect that the most promising architecture really has not been found.*

Daniel Gottesman, one of the pioneers of fault-tolerant quantum error correction, literally echoes this latter point:

> *The most promising fault tolerant scheme is most likely one that has not been invented yet.*

Gottesman goes on to add:

> *I think the most promising [schemes for fault-tolerance] in the long-run are schemes based on high-rate [low-density parity check] codes […]. However, there is not yet a practical protocol based on these codes.*

With respect to future progress in fault-tolerance, Gottesman writes:

> *I consider getting a better understanding of how to control correlated errors such as cross-talk errors in a quantum computer one of the [biggest] current open problems in fault tolerance.*

Some respondents suggested that for some highly specialized task or specific physical implementations such as ion traps, other types of quantum error-correcting codes, such as ones with longer-range interactions, may overperform surface codes. Tailoring to the platform may go as far as focussing on correcting errors that are likely in the given platform so that:

> *[…] fewer resources (i.e. qubits, gates and measurements) are wasted taking care of errors that are unlikely than in the standard [quantum error correction] approaches that are completely platform agnostic. In practice, a combination of these tailored codes and more generic codes, such as the surface code, might be a winning strategy towards fault tolerance and scaling up.*

One respondent emphasized how topological qubits are very promising for fault-tolerance:

> *The most promising scheme for fault tolerance is a system based on topological qubits, due to the ability for such a system to dramatically reduce the error rates on the qubits. In addition, such a system would not require off-chip communication, as required in other competing technologies, which may lend itself further to better scalability and fault tolerance.*

4.3 Recent developments

When asked to name important recent developments (approximately since the beginning of 2018), many respondents referred to the progress made in superconducting systems. Particularly, Frank Wilhelm-Mauch pointed out, this progress has led to "[r]eaching very consistent high fidelities of two-qubit gates across a large chip".

One respondent wrote:

> *Since 2018, we have seen how super-conducting systems have been scaled up to a size where, provided the error rates can be reduced to acceptable levels, it should be feasible to experiment with error correction for the purpose of producing logical qubits that could be used to perform some form of useful limited-depth computations. Further into the future, such systems may conceivably allow quantum supremacy to be demonstrated.*

Recent theoretical and experimental efforts are, in general, seen by our respondents as very promising for error correction and fault tolerance, including the development of new and practical error-correcting codes.

Speaking, instead, of the importance of "quantum supremacy" and of recent related developments in that direction, Dave Bacon wrote:

> *The definition of quantum supremacy has given a concrete target for implementations to shoot for. […] I believe this focus is incredibly useful for experimental teams, who often instead focus on optimizing components of quantum computer at the expense of an architecture where all things work together.*

One respondent wrote that, from a conceptual standpoint, the most important insight recently gained is that:

> *[t]he errors in digital quantum simulations can be understood within the same theoretical framework used to study the onset of quantum chaos […] I am absolutely convinced that understanding and controlling the onset of chaos in quantum computers will be a matter of "life or death" for the field.*

Rather than pointing to specific breakthroughs, Ashley Montanaro points to a change in the quantum research landscape (see also Section 1.3.4), that is, to:

> *sustained progress in experimental quantum computing research, especially from commercial quantum computing vendors.*

It is worth noticing that commercial companies in the quantum area have created what we could call 'cloud quantum computing prototypes', that is, in the words of one respondent:

> *systems that are not tailor-made or tuned for one application but that can be used (ideally by a remote user) for arbitrary algorithms.*

4.4 Next big step

When asked about the next big step towards the realization of a fault-tolerant quantum computer—something achievable conceivably by the end of 2020—many of our respondents mentioned the experimental implementation of some form of error correction that goes beyond "break-even", in the sense that the encoding of, and operations on, logical qubits will offer an actual improvement over lifetime and error rates compared to direct physical encoding.

A respondent set the bar higher:

> *An essential step will be the demonstration of an encoded qubit with an output error rate close to the level required for realization of […] Shor's algorithm on a [2048-bit] number.*

On the other hand, Dave Bacon would like to see:

> *[the demonstration of] a scalable architecture for one of the leading contending architectures. That is an architecture which can have thousands to hundreds of thousands of qubits and operate at least at some non-trivial fidelity. This is distinct from showing a fault-tolerant scheme and instead focuses on getting the building blocks correct in a way that doesn't blow out heat budgets, wiring, costs, etc.*

One respondent emphasizes that scalability while preserving quality is key:

> *I think that it is of utmost importance to show that one can build bigger and bigger devices with a constant error per gate (i.e., the errors in a gate is independent of, or converges with the total number of qubits). If this is not achieved, fault tolerant error correction will not work.*

Frank Wilhelm-Mauch 'sets an agenda' for various physical implementations:

> *For trapped ions: Reaching high fidelity and low heating for two-dimensional surface traps, similar to those achieved in 1D Paul trapped; for superconducting qubits: even higher two-qubit gate fidelities; for semiconductors: a path to avoid charge noise or its impact; in Rydberg atoms: more reliable trapping.*

The demonstration of quantum supremacy will also be a decisive step, according to other respondents, but its actual realization and, even more importantly, its significance on the path towards a cryptographically relevant quantum computer will have to be evaluated in light of the points raised in Section 4.1.2.

4.5 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[3]. 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, including growing interest from the private sector;
- the challenge that building a quantum computer constitutes;
- the dangers of hype and of high (and potentially, too short-term) expectations from funders, government and the public;
- the difficulty of making predictions about the rate of development in the field.

[3] Given that most of the text is quoted, we refrain from formatting quotes in the same fashion as quotes reported in other sections.

Here are the excerpts.

Dave Bacon: "Roughly I would define [a Noisy Intermediate-Scale Quantum (NISQ) device] as 50-250 qubits with fidelities of 99.9% or higher. While we believe such quantum computers can achieve quantum supremacy we do not have any idea if there are practical algorithms. Because this is expected to be a 'discovery' era, we are focusing on delivering software that provides workflows that are fast and iterative. […] There is a need for a lot more tooling in this era, and often the theoretical ideas needed are very far divorced from practical considerations."

Bill Coish: "By analyzing/understanding the detailed dynamical models governing qubit readout (and improving the physical systems used to perform readout), we can improve inference methods used to extract information from a string of qubit measurements. […] There's a strong move in the direction of exploiting machine-learning methods and other 'black-box' techniques, which are powerful, but when these come at the expense of exploiting additional knowledge of the physical error model, they will not realize their full potential. Combining these two approaches will be essential to reaching fault-tolerance and scalability."

Daniel Gottesman: "Experimental progress is always slower than I would like, and this has been true for 20 years. I see no real change in this, so, even though there is a lot of current excitement about the prospect of building large quantum computers, I still expect it to take a significant amount of time (20 or more years), although there is always the chance of a breakthrough that would accelerate matters. However, progress to date has been more and more convincing in showing that there is no insurmountable difficulty ahead, and that really it is only a matter of time until one is built. The most likely problem that would prevent that is a loss of interest (or patience) from funders."

Frank Wilhelm-Mauch: "[The] tight integration of quantum control and calibration with hardware […] is a largely untapped resource that keeps current hardware (specifically those without strong separation between energy scales, like superconducting qubits) [from] realising its full potential. This goes along with a tight integration of the quantum computer with its firmware.
[…]
A theoretical problem is […] the question how high coherence can be maintained in large processors, which simply means that our understanding of decoherence, albeit far developed, needs to improve even more."

Frank Wilhelm-Mauch: "Quantum computing is going through a great run right now and I believe that the influx of new people will help to accelerate even more. It is important to take research projects to the next level in scale and sharing of labour. In a few years, we will see if there are hard limits to building quantum computers—right now we are clearing out the not-so-hard ones. So I guess I am optimistic."

Ashley Montanaro: "Quantum computers are almost at the point where they can go beyond classical supercomputers for some carefully chosen problems, but it remains a significant challenge to demonstrate a quantum advantage for problems of practical interest. We may see this within a few years' time, depending on the eventual performance of NISQ machines, or it could take substantially longer. Improving the runtime of quantum algorithms for practically relevant problems, and also designing new quantum algorithms, continues to be of fundamental importance for the field."

Stephanie Simmons: "The time between the disclosure of the fission of uranium and the demonstration of a working atomic bomb was a mere six and a half years. The development of quantum computers—the information-security equivalent of the nuclear bomb—will at some point have a similar breakthrough moment which unlocks an effort equivalent to the Manhattan Project. Correspondingly, academic predictions based upon extrapolated past academic progress are somewhat meaningless, as it is very much a question as to when such breakthroughs will occur. Given that there are so many physical systems which could support this technology, which are each facing their own independent challenges, one could be excused for assuming that one such breakthrough will occur "soon". Such optimism draws upon the historical trajectories of other watershed technologies rather than the insider knowledge of a given physical platform. After all, if a convincing blueprint large-scale quantum computer already existed, the question would be cost, not time, and such systems would [be] manufactured behind the closed doors of a number of clandestine government facilities as quickly as humanly possible."

Respondent: "We need better gates and more qubits. These are coming, but slower than I would expect."

Respondent: "Research efforts should be more down-to-earth and research results should be less hyped."

Respondent: "[Building a quantum computer] is really tough, and we are not there yet. This will require a lot more fundamental research and collaboration between multiple disciplines. However, year after year we see real progress and this is encouraging."

Respondent: "The field is very overhyped, so that investors and the public feel we are on the verge of game-changing levels of quantum computing power. In practice we ARE at an exciting time when practical theory and experiment is co-evolving, but it is sobering to note that for a group that works on novel algorithms […] it is still FAR better to have access to a classical emulator than ANY of the various prototype quantum devices […] A key milestone that we can hope for in the next couple of years, related to quantum supremacy, is the point when serious quantum algorithm development needs access to prototype quantum computers."

Respondent: "The most important lesson to keep in mind while observing and judging the status of quantum computers development is that there are often big surprises. The literature in

the 1990s is wonderfully instructive: some of the most admired luminaries gave scathing commentary on the plausibility of quantum computers, providing seemingly unassailable arguments. Then quantum error correction came... Then, people retorted that the physical error rates required by quantum error correction were unrealistic. Then the surface code came... and so on. We need to remain alert of the hype (or outright misrepresentation of the facts), but also never forget that what we are doing is not like, say, developing a better kind of car. We are creating something that never existed before. Surprises will meet us around every corner."

Respondent: "I think it is a great area of research. I am afraid that after the hype triggered by the public announcements of several private companies and other public institutions may affect the scientific community if they decide to quit at some point."

Respondent: "I think that it is very difficult at this point in time to make a prognosis for the development of large-scale quantum computers capable of breaking currently widely deployed asymmetric cryptographic schemes based on [some] conjectured computational intractability […]

I am not yet entirely convinced that we will witness the materialization of such quantum computers. It may for instance turn out that the engineering challenges associated with scaling up the systems are simply too great, or that we first need technological breakthroughs to achieve scaling, leading to something akin to a quantum winter. Furthermore, it may turn out that the interest in funding research for the development of such computers dries up as time progresses and people move from current asymmetric schemes to post-quantum secure schemes. It may even turn out that that we will witness the development of efficient classical algorithms capable of solving some of the problems that underpin the current asymmetric schemes.

[…] It is hard to develop credible estimates at this point in time, given the great number of unknowns. […] [A]t the same time [I] recognize that it is conceivable that quantum computers capable of breaking current asymmetric schemes may materialize sometime within say a 10-25 year time period, with the probability becoming much greater towards the further end of the interval.

With respect to protecting confidentiality in the long term, it is important to be conservative and to assume such a scenario […].

In my opinion, there is currently something akin to a hype surrounding quantum computing. It is important therefore to properly and responsibly describe the current status of quantum computing, and the uncertainty associated with making prognoses for the future development of quantum computing. This should be done without detracting from the fact that mitigating action is needed if current asymmetric schemes are used to protect confidentiality in the long term."

Respondent: "We need to continue to reduce the resource requirements for fault tolerant quantum computing on a commercially relevant quantum computer, both through the

advancement of improved quantum error correction and distillation techniques, as well as through the development of optimized quantum algorithms."

Respondent: "We are still far away from fault tolerance and from factoring 2048-bit numbers! And there is certainly a lot of concern in the community currently that we are living in a bubble and that once funding agencies and industry leaders realize that a useful quantum computer is not around the corner, we'll be in trouble. Nevertheless, I am optimistic when I look at the progress that's been made recently on many fronts, in areas that I certainly wouldn't have anticipated. In particular, I think the close collaborations between experiment and theory are essential (and are happening in a way that wasn't possible ten or fifteen years ago, when experimenters were still demonstrating basic qubit functionalities). I also believe that quantum funding initiatives worldwide as well as the many new companies (and established companies that are now doing quantum research) will provide a lot of momentum and lead the field in surprising directions."

Summary and outlook

The quest to build a fully scalable, fault-tolerant quantum computer able to run the presently known quantum algorithms—and new algorithms that may be developed in the future—is a formidable one. It has often been described as a 'quantum race' (Hsu, 2019), with competition at the level of nations as well as of private companies. It has also been described as a marathon, rather than a sprint race. Nonetheless, to keep the analogy, the racecourse is partially unknown, and the finish line—in our context, the creation of a quantum computer that poses a significant risk to cyber-security—might be closer than some may think.

Indeed, in the aggregate, our experts judged that developing a quantum computer that could break a scheme like RSA-2048 within the next 10 years was unlikely, but also that such a possibility was far from negligible. Even more importantly, the experts indicated a very significant chance of a quantum threat emerging within 15 years. Thus, the opinions collected and analyzed in this report may be seen as 'optimistic' for the future of quantum computing, and 'worrying' for cyber-security.

Our respondents are generally devoting their careers to quantum information science and quantum computing. Does this mean they are implicitly biased toward believing in the earlier achievement of large-scale quantum computers? Or are they implicitly biased toward 'under-promising and over-delivering' and not setting up expectations that may not be achieved? We are confident this distinguished cohort of leading scientists would not deliberately overestimate or underestimate the chances of building a quantum computer for a given timeframe, and that the opinions the experts expressed are genuine best estimates based on their own deep and relevant expertise and knowledge. The logical possibility that meaningful quantum cryptanalysis is, for some reason, infeasible or impossible is captured in the non-negligible likelihood that quantumly breaking RSA-2048 will take over 30 years. The salient question for cyber-risk managers is not whether there is a non-negligible chance that quantum cryptanalysis is 20 or more years away, but whether there is a non-negligible chance that quantum cryptanalysis is less than 20 (or 15 or 10 or 5) years away. Lines of reasoning for the impossibility or infeasibility of large-scale quantum computation are typically based on the fact that quantum properties are very hard to preserve and control, and that the present proof of the possibility of fault-tolerant schemes are based on reasonable assumptions that may nevertheless turn out to not capture key aspects of noise and errors. Quantum computing researchers, including our respondents, are very aware of these issues and of the great challenge that building a quantum computer constitutes; they are hard at work testing the assumptions and overcoming the obstacles. On the other hand, for people managing cyber-risk, there is nothing close to a scientifically convincing or established argument for why the efforts currently underway are highly likely to fail 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.

Our respondents provided several indications about what to 'aim for'—or 'watch out for', in terms of judging progress towards a quantum computer that may disrupt cybersecurity—in the near future that will constitute an important step forward, or milestone. This includes improving the quality of quantum

gates, particularly two-qubit gates, and demonstrating experimentally that error correction can be used to prolong the storage and manipulation of logical qubits.

Another milestone will be the achievement of quantum supremacy. It will signal that there has been great progress in our ability to build and operate quantum devices, and it will certainly receive the attention of news outlets. On the other hand, it will only be a relatively small step towards a cryptographically relevant quantum computer, which requires a much higher level of sophistication, specifically in relation to using error correction, to achieve fault-tolerance.

The expert opinions collected in our survey and summarized in this report offer unique insight into the quantum threat timeline. Depending on its own specific shelf-life times and migration times, each organization will have a longer or shorter time at its disposal to implement post-quantum cryptographic solutions. In the words of one of the respondents:

> *Mitigating actions should be taken now by standardizing and rolling out post-quantum secure asymmetric [cryptographic] schemes, or symmetric keying, whenever feasible, for the purpose of complementing or replacing current asymmetric schemes. Taking mitigating actions now, in good order, provides an affordable insurance should large-scale quantum computers capable of breaking current asymmetric schemes materialize in the future.*

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

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

References

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

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

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.

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

Gidney, C., & Ekerå, M. (2019). How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. arXiv:1905.09749.

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

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/

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.

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

Shor, P. W. (1999). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review, 41*, 303.

Appendix

List of respondents

Name	Institution	Country
Scott Aaronson	University of Texas at Austin	USA
Dorit Aharonov	The Hebrew University of Jerusalem	ISR
Dave Bacon	Google	USA
Simon Benjamin	University of Oxford	GBR
Alexandre Blais	Université de Sherbrooke	CAN
Ignacio Cirac	Max Planck Institute of Quantum Optics	GER
Bill Coish	McGill University	CAN
David DiVincenzo	Forschungszentrum Jülich	GER
Runyao Duan	Institute for Quantum Computing, Baidu Research	CHN
Martin Ekerå	KTH Royal Institute of Technology and Swedish NCSA	SWE
Artur Ekert	University of Oxford	GBR
Daniel Gottesman	Perimeter Institute for Theoretical Physics and Quantum Benchmark Inc	CAN
Jungsang Kim	Duke University	USA
Ashley Montanaro	University of Bristol	GBR
Andrea Morello	UNSW Sydney	AUS
Yasunobu Nakamura	The University of Tokyo	JPN
Tracy Northup	University of Innsbruck	AUT
Peter Shor	Massachusetts Institute of Technology	USA
Stephanie Simmons	Simon Fraser University	CAN
Krysta Svore	Microsoft	USA
Frank Wilhelm-Mauch	Saarland University	GER
Shengyu Zhang	Tencent	CHN

Scott Aaronson
A leading computer scientist in the study of the capabilities and limits of quantum computers; he devised the boson-sampling problem as one of the first computational problems where quantum devices could prove superior to classical computers

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.

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

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

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

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

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

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

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

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.

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.

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.

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.

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.

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.

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

Questions

Besides identification questions and preliminary questions about familiarity with various subfields of quantum information and quantum computing research as well as with experimental implementations (QUESTIONS 1 to 7), the following were the major questions appearing in the online questionnaire.

QUESTION 8

Please indicate the potential of the following physical implementation as candidates for fault-tolerant quantum computation.

Physical implementations: Superconducting systems, Trapped ions, Quantum optics, Quantum spin systems (quantum dots, NV centers, ...), Topological systems

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

QUESTION 9-10

Rank the following quantum implementations in terms of their potential for realizing a digital quantum computer with 100 logical qubits in the next 15 years.

Physical implementations: Superconducting systems, Trapped ions, Quantum optics, Quantum spin systems (quantum dots, NV centers, ...), Topological systems, Other

QUESTION 11

Please indicate how likely you think it is that a quantum computer / device will demonstrate so-called "quantum supremacy" (that is, the ability to perform some computation practically impossible for classical computers, including classical supercomputers, irrespective of the utility of such a computation) within the next 1 year, 3 years, 5 years, and 10 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)

QUESTION 12

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.

Classification as for Question 11.

QUESTION 13

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

QUESTION 14

What has been the most significant recent (that is, approximately since the beginning of 2018) achievement in the progress towards building a fault-tolerant quantum digital computer?

QUESTION 15

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 the end of 2020)

QUESTIONS 16-17

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

QUESTION 18

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

Some details on the analysis methods

As mentioned in Section 3.1 and in this appendix, we asked the respondents to provide an informative but rough estimate of the likelihood of quantum supremacy being demonstrated (of a quantum computer able to factorize a 2048-bit number in less than 24 hours, respectively) within a certain number of years. In order to derive from such responses the cumulative probability distributions as shown in Section 4.1.3, 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)	1
Very likely (> 95% chance)	0.99
Likely (> 70 % chance)	0.95
Neither likely nor unlikely (about 50% chance)	0.7
Unlikely (< 30 % chance)	0.3
Very unlikely (< 5% chance)	0.05
Extremely unlikely (< 1% chance)	0.01

Pessimistic assignment:

Extremely likely (> 99% chance)	0.99
Very likely (> 95% chance)	0.95
Likely (> 70 % chance)	0.7
Neither likely nor unlikely (about 50% chance)	0.3
Unlikely (< 30 % chance)	0.05
Very unlikely (< 5% chance)	0.01
Extremely unlikely (< 1% chance)	0

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

In order to generate the graph of Figure 14, the resulting cumulative probabilities of the experts have simply been averaged for both the optimistic assignment and the pessimistic assignment.

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) in order 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.