QUANTUM THREAT TIMELINE REPORT 2020

JANUARY 2021

Authors: Dr. Michele Mosca, *Co-Founder & CEO, evolutionQ Inc.*
Dr. Marco Piani, *Senior Research Analyst, evolutionQ Inc.*

GLOBAL RISK INSTITUTE

evolution

EXECUTIVE SUMMARY

We are pleased to provide an update to the first report issued in November 2019. This report documents a shift in the opinions from the last report due to the advances and changes in the quantum computing landscape.

Quantum computers use quantum systems to run computations that go beyond what is achievable by standard computers. They do this by exploiting quantum features that are difficult to preserve and control; this makes building a quantum computer an immense challenge. Despite some skepticism about their realizability, no fundamental roadblock has been identified, and relatively small prototypes have already been built.

Once available, full-fledged quantum computers will be able to solve computational problems previously thought to be intractable, hence breaking several elements of the current 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: 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 depends 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, and
- 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 against quantum attacks for the years required.

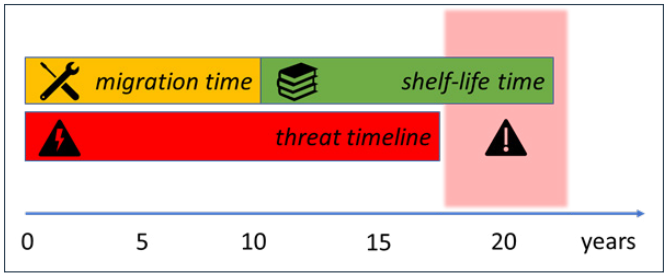

A better understanding of the threat timeline provides information on the time available to safely perform the transition to post-quantum cyber-systems.

Assessing the quantum threat timeline is challenging because building a working quantum computer requires pushing beyond the limits of what is known scientifically and beyond what is presently possible from an engineering perspective. Experts generally acknowledge that we do not know when we will have quantum computers that can threaten cyber-systems, as predicting progress in the field is not straightforward. However, this series of reports aims at providing (1) insight into the prospects of this threat becoming real in the short and medium term, (2) the rate at which progress is being made, and (3) the key milestones cyber-risk managers should pay attention to.

In 2019 we surveyed an unprecedented breadth and depth of 22 thought leaders with questions designed to help those managing the cyber-risk associated with quantum cryptanalysis. This year we have asked the same experts for an update on their opinions and have received 21 responses to key questions. Furthermore, we have enlarged our respondent pool with 23 new respondents. Forty-four leading experts, worldwide, have provided input this year, doubling last year's number.

The pool of respondents is comprised of experts from academia and industry, working on several aspects of quantum computing research from across four continents. We note that employees from some companies declined to take part in our survey, but we nonetheless secured a significant representation of some major private players in the field.

Expert Opinions on the Likelihood of the Quantum Threat to Current Public-Key Cryptosystems

The experts were asked to provide estimates on the development timeline for quantum computers, specifically for quantum computers powerful enough to pose a threat to cybersecurity. Several respondents articulated the difficulty inherent in making such predictions, and it is not surprising that opinions varied significantly. Despite such variance, the opinions expressed indicate that the quantum threat will quickly become non-negligible in the future, and it could well become concrete sooner than many may expect. Here are some trends.

Next 5 years:
Most experts (27/44) judged that the threat to current public-key cryptosystems in the next 5 years is "<1% likely". A quarter of them (11/44) judged it relatively unlikely ("<5% likely"). The rest selected "<30%" (3/44) or "about 50%" (3/44) likely, suggesting there is a non-negligible chance of a short-term and impactful surprise.

Next 10 years:
Still more than half of the respondents (23/44) judged this was "<1%" or "<5%" likely, but 11/44 felt it was "about 50%" or ">70%" likely, suggesting that the quantum threat could very well become concrete in this timeframe. In addition, 9/44 experts judged the likelihood smaller than 30%.

Next 15 years:
Slightly more than half (23/44) of the respondents indicated "about 50%" likely or more likely, among whom 5 indicated a ">70%" likelihood, and 7 an even higher ">95%" likelihood. Still, half of the respondents (21/44) indicated "<30%" or even less likely.

Next 20 years:
About 86% (38/44) of respondents indicated "about 50%" or more likely, with 12 feeling it was ">95%" or ">99%" likely. This suggests that the near consensus is that the quantum threat is likely to occur before the 20-year mark.

Next 30 years:
All but one among the experts responded that the quantum threat has a chance of "about 50%" or more, with 36 out of 44 experts indicating that the quantum threat will be likely (">70%"), very likely (">95%") or extremely likely (">99%").

If we compare this year's opinions to those of last year, despite judging it somewhat less likely *in the very short term (i.e. within 5 years)*, the experts seem to assess quantum computing as being closer to becoming a reality than what one could have expected from the simple fact that one year has passed since the previous survey. These somewhat conflicting tendencies (that is, less likely in the very short term, but more likely in the short/medium term) could stem from a better understanding of the hurdles towards building a quantum computer and approaches for circumventing them. This better understanding might increase confidence in its eventual realization and reduce

concerns about fundamental new hurdles or show-stoppers potentially emerging. At the same time, greater understanding of the challenges might also reduce the expectation of a short-term breakthrough. Another factor that may affect the short-term predictions is the effect of the present ongoing pandemic, which the experts estimate will slow down progress in the short term.

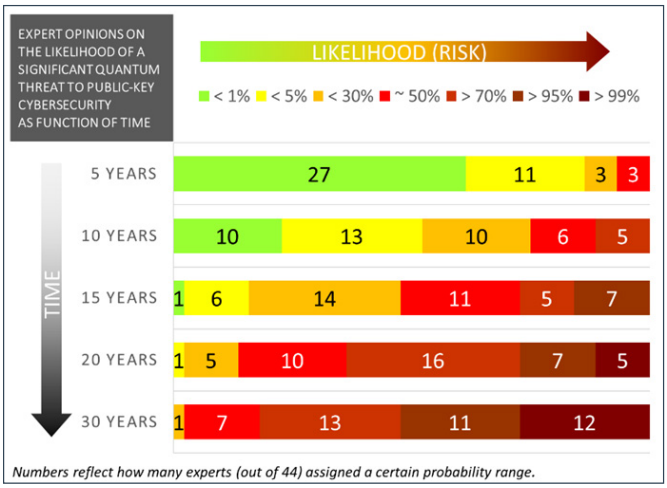

Expert Opinions on the Technical Realization of Quantum Computers

Creating reliable fundamental components, so-called (physical) *qubits*, the number of which can be scaled while maintaining control and quality, presents a major challenge in building a quantum computer. As was the case last year, 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. Several other physical implementations were noted as having significant potential.

Physical errors cannot be eliminated completely. Imperfect physical qubits can nonetheless be combined into so-called logical qubits. Moving forward, It will be very important to demonstrate, experimentally, that error-correcting schemes improve the reliability of 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.

From the Threat Timeline to the Migration Timeline

The expert opinions collected in our surveys 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 quantum-safe 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, A Methodology for Quantum Risk Assessment, 2017).

The Global Risk Institute and evolutionQ Inc. hope to 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

Note of correction, January 2022

A previous version of the 2020 report contained a calculation error that affected Figure 15, where it comes to the average probability associated to an optimistic interpretation of the answers provided by the respondents. The figure has been amended.

1. INTRODUCTION / BACKGROUND

We provide here the background necessary to understand how quantum computers pose a threat to cybersecurity, and how building such computers is an incredible scientific and technological challenge. The content of this introduction, as well as some of the content of Sections 2 and 3, is similar to content in our 2019 report and is provided to make the present report comprehensive. Nonetheless, Sections 1.3 and 3 are new or modified to take into account both changes in the field of research and in our survey questions.

1.1. Quantum Computing

Quantum mechanics is our best description of the inner workings of nature. It allows us to explain the behaviour of matter and energy at small physical scales, including the behaviour of fundamental particles like electrons, or of atoms and molecules. 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 typical quantum phenomena are more directly manifest at a microscopic scale than at the macroscopic, everyday-life scale, because of the magnitude of the 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 precision, that has no precedence in human history.

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

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

The basic unit of quantum information manipulated by a quantum computer is the quantum bit, or *qubit*. Unlike a standard bit, a qubit can store not only the two values 0 and 1, but also a *superposition* — technically, a linear combination — of them.

Not only will quantum computers allow us to simulate quantum systems as proposed by Feynman, but, by exploiting quantum features like superposition, they will be able to tackle a number of mathematical, optimization, and search problems much faster than conventional computers.

1.2. Quantum Threat to Cybersecurity

Widely used public-key cryptographic schemes rely on mathematical problems which are intractable for classical computers, the best known example probably being 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 schemes can be broken by quantum computers. For instance, RSA can be attacked by implementing Shor's algorithm (Shor, 1997). Furthermore, the ability of a quantum computer to search through a solution space with 2^n values (i.e., all the possible combinations of n bits) in roughly $2^{n/2}$ steps (Grover, 1996) would also weaken symmetric-key cryptography.

The threat posed by quantum computers can be mitigated by adopting new cryptographic tools which are designed to be resistant to quantum attacks. These so-called *quantum-safe* cryptographic tools can be conventional or quantum in nature, the first kind amounting to basing security on problems that are hard or believed to be hard also for quantum computers, and the second kind being based, for example, on quantum key distribution.

However, transitioning to quantum-safe cryptography is both arduous and delicate (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, and
- T_{THREAT} (**threat timeline**): the number of years before the relevant threat actors will be able to break the quantum-vulnerable systems.

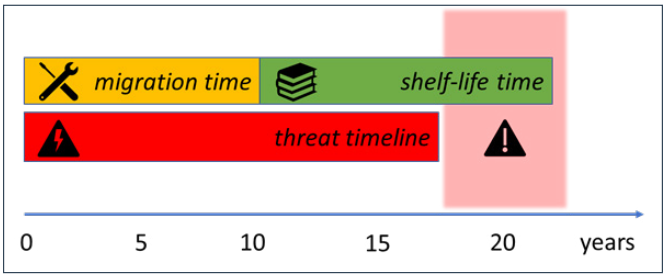

Figure 1: The timeline for the development of quantum computers that may pose a threat to cybersecurity should be compared with the time needed to migrate the cyber-system to post-quantum security combined with the shelf-life time of the data to be protected. See main text for details.

If $T_{\text{SHELF-LIFE}} + T_{\text{MIGRATION}} > T_{\text{THREAT}}$, then an organization will not be able to protect its assets for the required $T_{\text{SHELF-LIFE}}$ years against the quantum threat (see Figure 1). Organizations need to assess $T_{\text{SHELF-LIFE}}$ and T_{THREAT}. Their difference, $T_{\text{THREAT}} - T_{\text{SHELF-LIFE}} =: (T_{\text{MIGRATION}})^{\text{MAX}}$ is the **maximum available migration time**, that is, the maximum amount of time they have at disposal to safely realize the transition. A key point is that

Rushing the process of migration might itself create security issues which could be exploited even with standard computers.

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

While the security shelf-life $T_{\text{SHELF-LIFE}}$ is generally a business decision or dictated by regulations, assessing the threat timeline T_{THREAT} is a much less straightforward task. This is mostly because there are numerous scientific and engineering related challenges to overcome before building a quantum computer capable of breaking existing cryptographic schemes. While these challenges imply that the deployment of cryptographically relevant quantum computers is likely to only happen in many years, it also means that technical progress and scientific/engineering breakthroughs may suddenly accelerate such a deployment. Investments into the development of quantum computers and, more generally, quantum technologies, also play a major role in the speed at which advances are made. Investments in the area have grown enormously in recent times (see also Section 1.3.4), coming from all sources: governments, private companies, and venture capitalists.

1.3. Realization of Quantum Computers

Quantum information can be encoded and processed in many different physical systems that behave quantumly. The latter include, e.g., quantum spins, or the polarization of quanta of light—photons.

Regardless of implementation, a common issue to contend with is that of the previously mentioned decoherence, due to the interaction with the environment and leading to the loss of the quantum features used to encode and process proper quantum information. It is vital to ensure adequate preparation of the physical system, maintain control of it, and measure it, while isolating it from the surrounding environment. Given the miniature scale of the systems at play and the numerous potential sources of decoherence, this is a daunting task.

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

Apart from the issue of the physical realization, 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 offers different insights into the design of algorithms, or can be more suitable for a particular physical realization. One such model is the *circuit* model or *gate* model, where transformations are sequentially performed on single and multiple qubits. From the perspective we are interested in — analysing the quantum threat timeline — it is useful to focus on the circuit model as there is a well-articulated path to implementing impactful cryptanalytic attacks.

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

The following criteria, which are part of a larger set of desiderata, were listed by DiVincenzo in (DiVincenzo, 2000) and are hence known as *DiVincenzo's criteria*. Historically they 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, and*
5. *a qubit-specific measurement capability.*

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

1.3.1. Physical Realizations

The various physical implementations have their own advantages and disadvantages in relation to factors such as (but not limited to): building and controlling larger and larger quantum devices with more and more qubits, also known as *scalability*; compatibility with/ease of implementation of different computational models; typical decoherence time; speed and precision with which gates can be applied. 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, and can be implemented also on-chip using waveguides.
- 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 non-vanishing total electric charge) that are confined by electro-magnetic fields.
- Quantum spin systems, meaning that information is stored and manipulated in the internal degree of

freedom called quantum spin; such systems may be realized in silicon, like standard microchips are, or in less conventional systems, like diamonds with point defects known as nitrogen-vacancy (or NV, in short) centers.

- Cold atoms gases, where neutral atoms (rather than ions) are cooled down to close to absolute zero; while ions repel each other because of their electric charge, neutral atoms do not, and can be trapped and arranged in very regular arrays via the use of laser beams, and controlled at single-site level.

1.3.2. "Quantum Supremacy"

"Quantum supremacy"[2] (Preskill, 2018) may be generally described as 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 firmly establishing whether a device has achieved quantum supremacy are somewhat 'fuzzy'. 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 also 'possible', algorithms, quantum supremacy can be considered a moving target, because classical computers and known classical algorithms improve over time.

This said, quantum supremacy has been 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 a quality not high enough to allow full quantum computation but still potentially useful to greatly outperform classical computers in some tasks.

Google argued to have achieved quantum supremacy in (F. Arute et al., 2019), which appeared around the time the 2019 Quantum Threat Timeline report was issued. That report described quantum supremacy as one key milestone to track and the experts, then surveyed, expressed it was well within reach in a relatively short amount of time.

Google's claim was challenged, exactly because of the lack of clear-cut criteria for quantum supremacy. Nonetheless the consensus is that (F. Arute et al., 2019) represents a landmark result. In the present survey we have asked the experts to weigh in on its significance for the field (see Section 4.2).

1.3.3. Error Correction, Fault Tolerance, and Logical Qubits

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*. *Logical* qubits are encoded into multiple *physical* qubits, so that errors can be detected and corrected, and logical information be protected (an illustration of these concepts for classical bits is provided in Figure 2). 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.

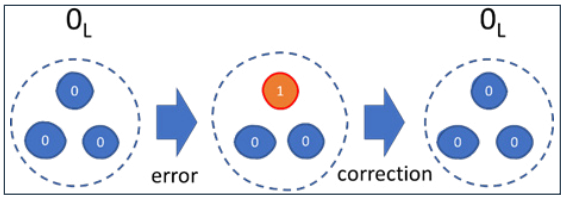

Figure 2: Example of classical information encoded logically. Several imperfect/error-prone physical bits (warped filled blue circles) are used to encode a logical 0, denoted 0_L (dashed perfectly round circle), by means of a repetition code: 0_L is encoded as 000 at the physical level. Errors can occur at the level of the physical bits, but they can be corrected, in this case by a simple majority voting scheme, so that the logical bit is preserved. As long as the probability of a physical bit flipping is small enough, the probability of a logical bit being affected by an error (in this case, going from 0_L to 0_L) is less than the probability of a physical flip.

2 This terminology is somewhat controversial because it recalls, e.g., racial supremacy. Nonetheless it has been widely used in literature, in the same way in which, e.g., "air supremacy" may be used in warfare jargon; in our context, "quantum supremacy" indicates superiority of quantum computers over classical computers in some strictly technical sense.

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 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.

Lattice surgery is a technique to merge and split surface codes to implement fault-tolerant interactions between qubits encoded in separate surface codes (Horsman, Fowler, Devitt, & Van Meter, 2012).

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 also evident by the affiliations of our pool of respondents. This interest has been supported and boosted by several national and transnational initiatives, like the National Quantum Initiative in the United States (Raymer & Monroe, 2019) and the Quantum Flagship Initiative in the European Union (Max, Kovacs, Zoller, Mlynek, & Calarco, 2019), with investments in the field of quantum technologies seen as strategic. In addition, many start-ups specializing in various aspects of quantum computing research have been established, often supported by venture-capital investments. Some are directly represented among our pool of respondents.

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.

Nonetheless, in the 2019 report, several respondents indicated that the risk created by a combination of hype and high expectations for the field, could lead to reduced funding and future investments, if those high expectations are not met due to slow or slower-than-anticipated progress. Such a "quantum winter" scenario, as called by some, could trigger a vicious feedback loop between slow progress and a decrease in funding, leading to a substantial slowdown in the development of a cryptographically relevant quantum computer — a 'stretching' of the quantum threat timeline.

To better understand the likelihood of this scenario, we asked this year's respondents to express whether they see funding in the field increasing, decreasing, or staying stable.

2. SCOPE OF THIS REPORT

This document reports the results of a survey conducted by evolutionQ Inc. among 44 leading experts, internationally, on quantum computing research. Experts were asked to complete an online questionnaire on the state of development in this field. This follows a similar survey conducted in 2019 among 22 experts, 21 of whom have participated again this year.[3] This report aims not only to provide a snapshot of experts' opinions, but also point to shifts in opinions. Essentially the entire cohort of past respondents has taken part in the present survey which allows us to provide a consistent analysis.

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 3.2 below) speaks explicitly of breaking RSA-2048, which security is based on the difficulty of factoring a 2048-bit number.

3 Some of the past respondents only provided input to a subset of this year's questionnaire.

The threat that quantum computers pose to RSA-2048 has already been considered in, 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.

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

- We seek a fine-grained picture of what leading experts think with respect to the timing and likelihood of the quantum threat. While our survey obviously cannot provide definitive answers, we aim to depict a much better picture of what 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?

- There are many fast-moving parts in the field, and much can change in just one year, as development in the field has been characterized in the years by steady progress combined with breakthrough results. Risk managers need to know how the expected quantum threat timeline is affected by such changes. For example, in the year after the NAS report, in the public literature the overall cost estimates of breaking RSA-2048 have 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.

- We aim to track how the opinions of experts evolve over time. This year's report is the second one in this series, with the potential for us to continue running such reports as long as the community finds them helpful.

- The scope of our survey/report is much tighter and more focused than that of other reports.

- We ask specific questions to several leading researchers, individually, and compile relevant statistics.

- We ask the respondents to indicate what they judge as the most important milestones to pass, or the necessary steps in the creation of a quantum computer.

- We give the respondents the chance to provide free-reign comments on the status and expected evolution of the field, in doing so gaining substantial insight on what to expect and the outlook for the future.

Note added:

Subsequent to conducting our survey, during the period we were compiling findings, a significant pre-print appeared (Sevilla & Riedel, 2020). In this work, not yet peer-reviewed, the authors try to forecast progress in the domain of quantum computing, not by polling experts as in our series of reports, but by extrapolating progress in the field and looking at relevant metrics. Sevilla & Riedel focus on superconducting implementations, and on the task of breaking RSA-2048, as we do. Their estimates for when (super-conducting) quantum computers could achieve such a feat, are described by the authors as "one piece of relevant evidence that can supplement expert opinion" and "more pessimistic but broadly comparable to those produced through the survey of experts in (Mosca & Piani, Quantum Threat Timeline, 2019)". They write that a cryptographically-relevant quantum computer could be built earlier if progress is faster than what they forecast through extrapolation of current—and still limited in temporal span—trends.

Other significant news that appeared after the completion

of our survey was the announcement by IBM of a relatively aggressive roadmap for the development of their family of superconducting quantum chips (Gambetta, 2020). Such a roadmap includes, for example, a 1,000-qubit machine in 2023, and machines with millions of qubits realizing a fault-tolerant quantum computer after that.

3. SURVEY DESIGN AND METHODOLOGY

We phrased our questions using a range of non-trivial considerations. It was most important to understand the perspectives of the diverse range of people asked to complete the survey and how the target audience would interpret the questions and possible answers.

This year we also wanted to strike a balance between keeping a fixed subset of questions, so as to track the shift of opinions in time, and replacing/adding some questions to capture the changes in the field.

In wording the core questions, we aimed to avoid having the various respondents interpret them 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, including that of tracking evolving opinions:

- We kept the questions largely focused on the issue of implementing fault-tolerant quantum computers that would be able to run quantum algorithms posing an actual threat to cryptosystems.

- We sought a range of relevant perspectives. In 2019, we engaged a select number of respondents with authoritative and profound insights. They provided a great variety of expertise on the most recent developments and the next steps needed towards the realization of fault-tolerant quantum computers. The same philosophy guided our selection of the additional 23 respondents of this year.

- Considering the quality of the pool of respondents — all busy professionals and researchers — we kept the questions limited in number so that the estimated time to complete the questionnaire was about 30 minutes. We gave the option to those respondents who took part in our 2019 survey (and were key in assessing trends around timelines) to provide input confined to a limited number of questions.

- Given the inherent uncertainty in progress towards realizing a quantum computer, we gave respondents the opportunity to indicate the likelihood of something happening in a relatively coarse-grained fashion, but still much more informative than what has been available prior to this series of surveys.

- We kept several of the basic questions from the 2019 report unchanged to capture a change in trends (see Section 3.2 for more details).

- We modified the set of questions accounting for recent developments in the field, such as the demonstration of so-called quantum supremacy, and the efforts shifting towards quantum error correction and the realization of logical qubits.

Some of the questions were optional, specifically when it came to sharing details about respondent's personal research activities, and to provide more free-form opinions on the state of the field.

Finally, to facilitate frank answers we analyzed estimates in an aggregate, anonymous fashion which we shared in advance with the respondents. 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. Looking for Trends

Our goal is and always has been to provide a useful assessment of the quantum threat timeline based on the opinion of experts. There is a natural and unavoidable uncertainty about such a timeline, which is affected by several factors, ranging from potential breakthroughs in science and the technology being developed, to changing

levels of investment by countries, institutions, and companies, to relatively unexpected events that affect society and economy as a whole, e.g., the COVID-19 pandemic. All these factors also inform the opinion of experts, be it consciously or unconsciously, and one may expect a corresponding shift in opinions over time.

Tracking changes in the opinions may be considered as important as taking snapshots of such opinions, because they may point to speed-ups or slow-downs in the development of quantum computers.

3.2. Questions

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

The key survey question 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 classifications 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 the realization of a fully controllable fault-tolerant qubit:

Please indicate how likely you estimate that a single fully controllable fault-tolerant (logical) qubit will be demonstrated 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 opinions on:

- which physical platforms are the most promising,
- schemes for fault-tolerance,
- recent and near-future expected progress that has been and will be, respectively, key in the development of a quantum computer,
- the impact of the COVID-19 pandemic, and
- the expected trend for the investments in the field.

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.

3.3. Participants

Most importantly, for the sake of consistency and tracking trends, we aimed at securing the participation of the same respondents of the 2019 survey. We were pleased that this was possible for the vast majority — 21 out of 22.

From a list of around 75 additional potential candidate respondents, we contacted several who, like the original respondents, were intended to provide a balanced and insightful range of opinions on the state of development of the field. Those who accepted were asked to complete the online questionnaire over two weeks.

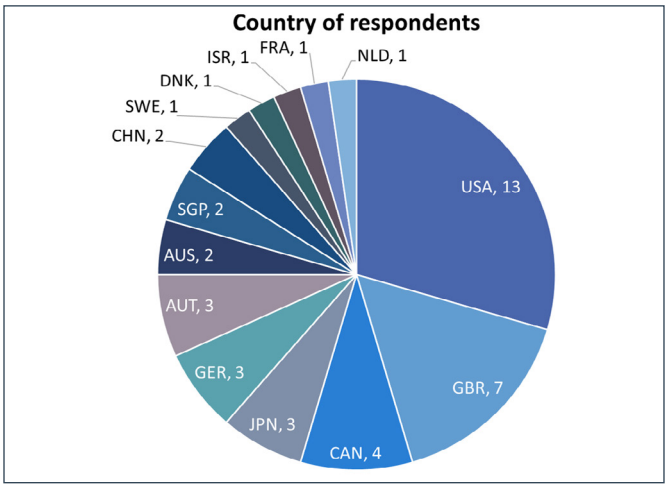

Figure 3: Our respondents constitute a very international mix, with higher representation from countries (like Canada, China, Japan, and USA) and geographical areas (like Europe) where the efforts to develop quantum computers and quantum technologies are very strong.

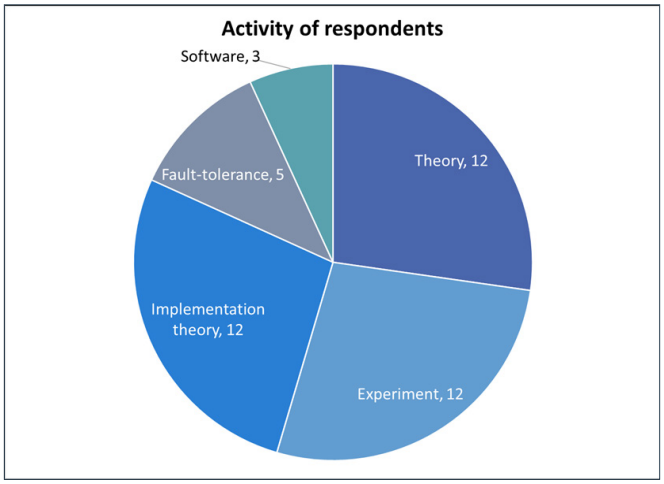

Figure 4: Our respondents cover a wide range of research activities. While the major division is between non-experimental research and experiment, research that is not directly experimental can be very different. E.g., implementation theory focuses on guiding, supporting, and, in general, facilitating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

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.

In about two months, we were able to collect responses from 44 respondents (see Appendix for a complete list). Figures 1-3 summarize classifications of our respondents in terms of:

- country where they work (Figure 3),
- kind of activity they lead (Figure 4), and
- kind of (primary) organization they belong to (Figure 5).

The respondent pool comprised a diverse set of expertise and nationality, and a mix of university and company researchers, representative of the diversity of the quantum computing community among its top players. Employees of some major companies did not respond or declined to take part in the survey. Nonetheless, we secured a significant representation of some private players in the field.

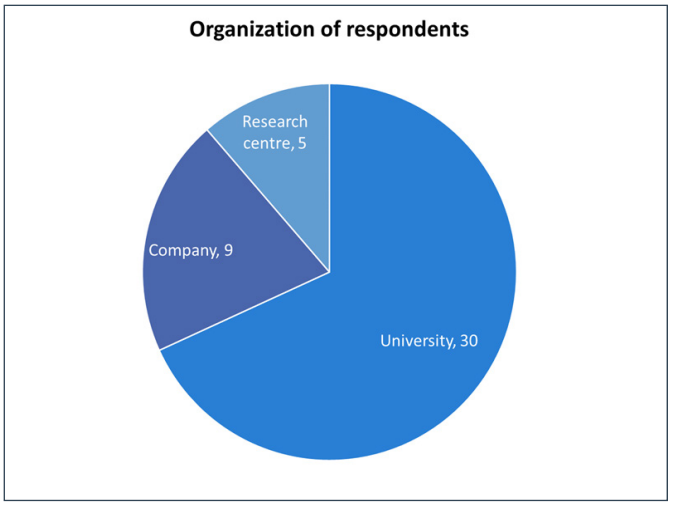

Figure 5: Most of the respondents work at universities, but some work at companies or research centres. Some researchers/ academics may have some role in, or collaborate with, external companies; this is not illustrated in the above classification.

4. SURVEY RESULTS

The survey results reported here comprise:

- an aggregate analysis of the key responses about:
    - physical implementations/platforms for quantum computing,
    - the quantum threat timeline,
    - the timeline for the implementation of a fully controllable fault-tolerant qubit, with experts' comments about fault-tolerant schemes, and
    - the expected change in funding in support of quantum computing research;
- a selection of opinions about:
    - key recent developments in the field of quantum computing research, as highlighted by the respondents,
    - how the COVID-19 pandemic has impacted, and is expected to impact, the progress of the field, and
    - near-future (that is, approximately, by the end of 2021) developments that the respondents see as essential on the path to developing a fully scalable fault-tolerant quantum computer;
- a collection of other notable remarks made by the respondents.

Where we deemed appropriate, we analyzed shifts in the responses as compared to last year.

4.1 Aggregated Analysis of Responses

In the aggregated analysis of the responses we indicate how many of the respondents (alternatively, what percentage of them) chose a specific answer among the many possible ones when dealing with multiple choices. There is, in general, substantial variability/spread in the answers, but some 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 'trajectories' of each respondent, that is, how the probability of the threat increases in time according to each respondent. 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 long time. It also shows that the pace of progress could be different, as some respondents estimate a low probability initially which then grows faster than in the responses of other experts.

Similar to what was done in 2019, we separately considered the responses of those participants who are close/closer to experiments. This group is comprised of 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 an 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, and
- 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.

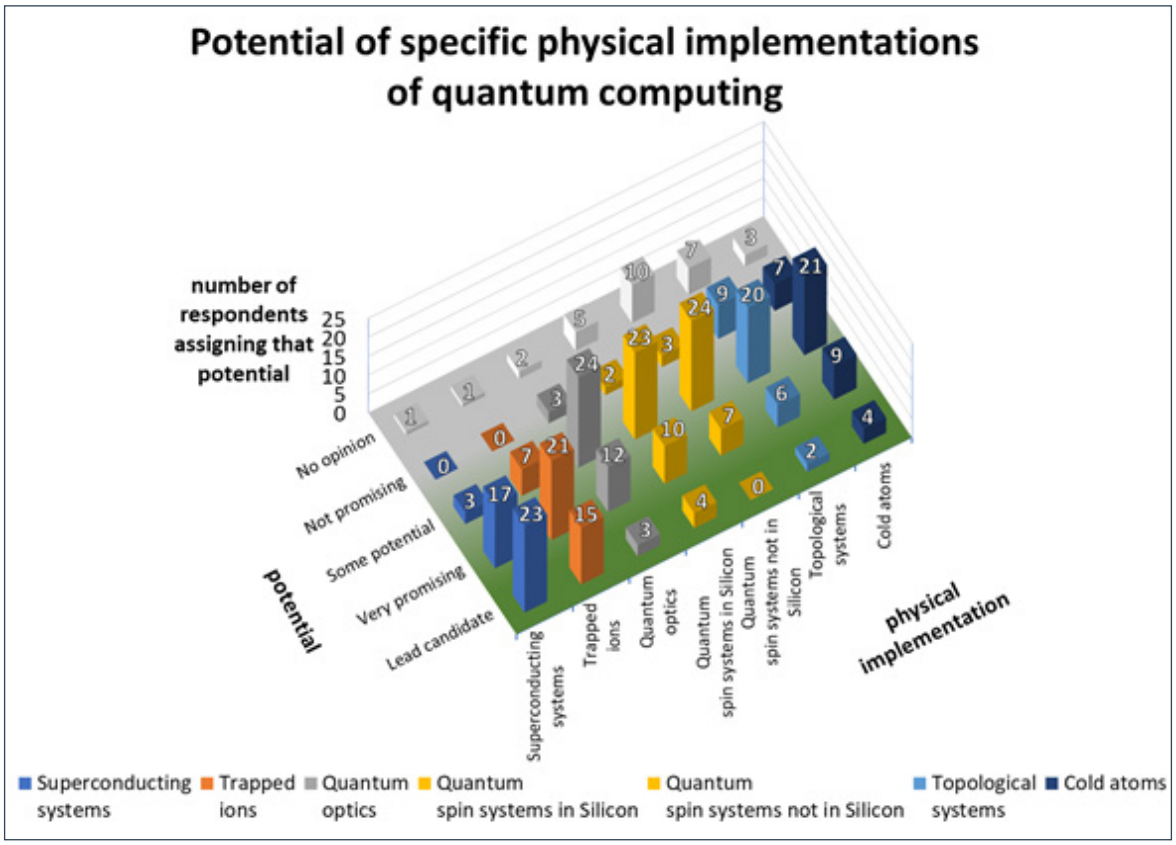

Figure 6: Superconducting implementations, followed by ion-trap implementations, are perceived as presently having an edge over other physical realizations. Compared to last year, we have added some categories; check main text for details.

With respect to the 2019 survey we have:

- introduced an additional platform, cold atoms, and
- split the category of silicon implementations into two more fine-grained categories.

Such changes were based on the feedback obtained from the experts during the preparation of the 2019 report.

The responses indicate a significant consensus that the present leading platforms are superconducting systems and trapped ions. This is consistent with the opinions from last year, and probably cemented by significant results obtained in both platforms during the last year (this includes quantum supremacy achieved by superconducting systems; see Sections 1.3.2 and 4.2).

Here is a selection of input from the experts that provides further insights on the potential of various platforms.

Winfried Hensinger stresses how ion-trap systems

[do not suffer from any] fundamental barriers. Especially since the development of ion microchips and the use of global microwave fields to execute quantum gates (rather than laser beams), the key challenges are engineering ones which are not simple [but also] are definitely not insurmountable.

He also highlights the potential modularity of such systems, which could be scaled to millions of qubits.

Other respondents praise neutral-atom platforms, with one saying:

Neutral-atom experiments with Rydberg gates have finally entered as a real contestant for quantum computing and may surpass trapped-ion performance.

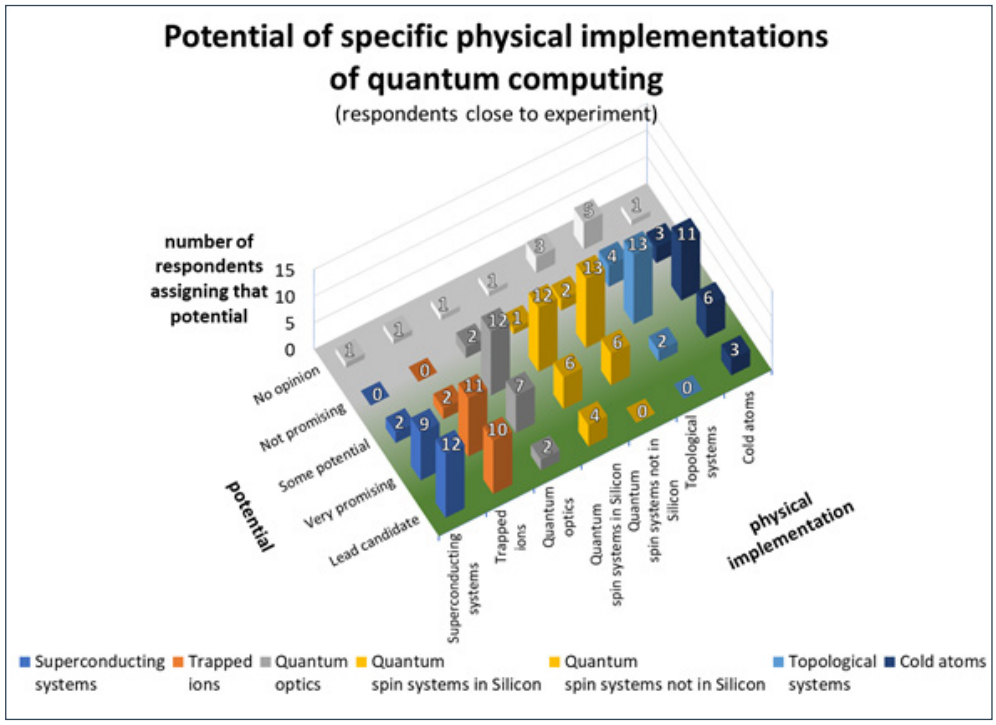

Figure 7: Similar to last year, respondents close to experiments judge trapped ions and superconducting systems as more similar in their potential than the general cohort.

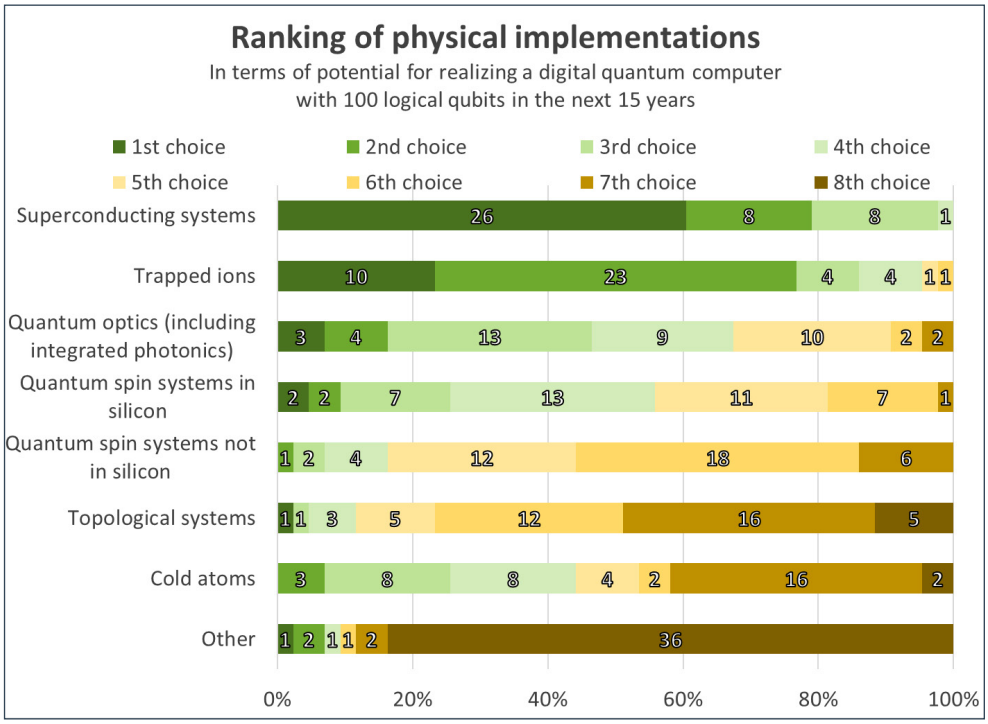

Figure 8: Superconducting systems and trapped ions are perceived as having an edge over other physical implementations not only in terms of generic potential for quantum computing (see Figure 6), but also when it comes to the more specific goal of developing a digital quantum computer with 100 logical qubits within the next 15 years. Numbers indicate how many respondents assigned that ranking. Not all respondents weighed in on an "other" system.

The challenge that topological qubits face compared with other platforms leads Joe Fitzsimons to write:

I think it is more likely that an as yet unknown system will achieve 100 logical qubits within 15 years than it is that topological qubits will reach this milestone.

On the other hand, Winfried Hensinger points out how the topological approach could ultimately pay off:

Topological qubits are high risk - high gain. They are unlikely to work but if they do as anticipated, then they will lead.

A respondent comments:

It is difficult to differentiate between the potential of cold atoms and trapped ions. I also believe that both quantum optics and spin systems in silicon may create a big surprise due to a breakthrough although at this point in time they appear less promising as fewer advanced results have been demonstrated with these systems.

Some respondents point out how hybrid systems could combine the benefit of different platforms, and that it the source of "Other" ranking high in Figure 8. Joe Fitzsimons elaborates:

The most obvious hybrid systems to me are those that couple matter qubits (whether [quantum] dots, ions, [NV centers] etc.) using entangling optical measurements. This is clearly one path to build a modular system, and the technology has been demonstrated in a range of systems since 2007. The relatively low entangling rate would make communication between modules relatively slow, but there are many benefits in being able to physically separate hardware components. [..] Quantum computing is currently trapped in an effort to build [quantum computers] as monolithic processors built from a single technology. If interoperability became more easily achievable between quantum components (for example superconducting systems and trapped ions or optical-frequency photons) this would open up exciting new paths for QC architectures.

Tracy Northup agrees:

I expect that a digital quantum computer with 100 logical qubits on any hardware platform will include hybrid solutions in the form of microwave and/or optical interconnects or, in the case of superconducting qubits, memories.

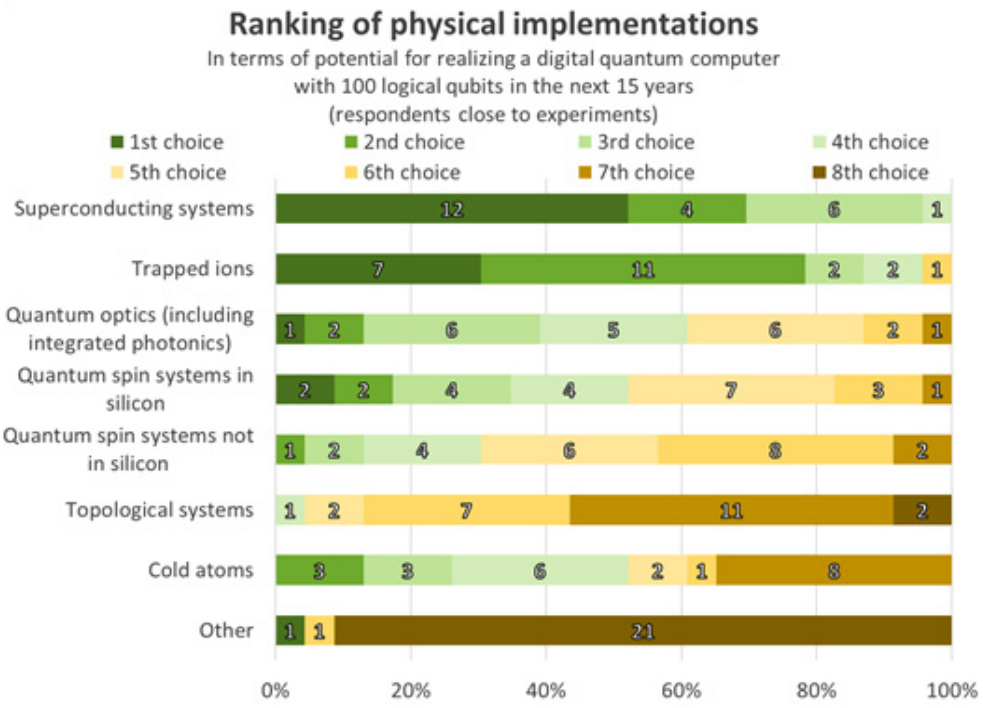

Figure 9: Superconducting systems and trapped ions are "closer" in potential if we restrict the pool of respondents to those closer to experiments. Numbers indicate how many respondents assigned that ranking. Not all respondents weighed in on an "other" system.

Similarly, Stephanie Simmons agrees:

It is conceivable that it will be easier to build modular, networked quantum computers than a monolithic quantum supercomputer.

Sir Peter Knight thinks the 'competition' among the various platforms is still quite open, writing:

[it is] still too early to have too much confidence in down-selecting at least for the next three years.

Another respondent points to the fact that, if we think of some platforms as more promising than others, it is perhaps because we have learned more about them, and there is room for other technologies to catch up:

we know how to build quantum computers using trapped ions, primarily because we know a lot about that kind of technology. For technologies that are less mature, like quantum dots in silicon, I think it is almost impossible to judge their potential for quantum computing.

Andrea Morello warns about the high number of physical qubits that appear to be needed in creating logical qubits in matter-based physical qubits:

With our current understanding, 100 logical qubits with matter-based physical qubits will require of order a million physical qubits. This is why I am not placing superconductors and trapped ions at the top—their physical-qubit footprint is concerningly large.

Instead, he says he has:

great hope in integrated photonics, especially to the extent that it can exploit silicon-foundry manufacturing processes.

Another respondent stresses that in optical implementations there might be big differences:

I don't see photonic qubits as all that promising [..]. In optics, I see continuous-variable approaches as having more potential.

4.1.1.1. Comparison with Last Year

As mentioned, last year all spin systems were combined in just one category, and cold atoms had not been explicitly included as a potential platform (but some respondents

had indicated it as “other”, hence the decision to include them explicitly this year).

The comparison seems to indicate that, while superconducting systems have maintained the lead, there might have been enough progress in trapped-ion systems and in quantum optics that such categories have seen their perceived potential increase. Cold atoms have also made a relatively strong entrance, which justifies tracking them as their own category.

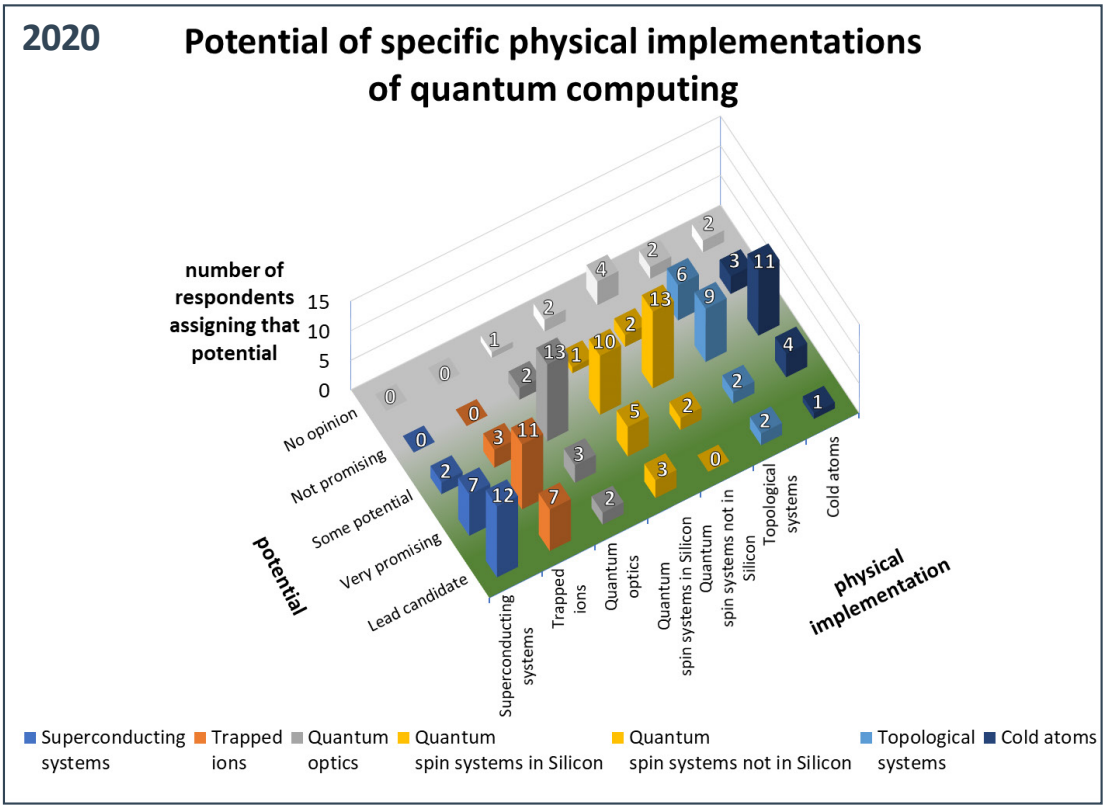

Figure 10: Potential of various physical implementations of quantum computing: comparison of the opinions expressed this year (2020, top panel) by almost all the respondents (21/22) who had already participated in the survey in 2019, with their past aggregated opinions show in the bottom panel. Note the slight change of categories (see main text for details).

4.1.2. Quantum Factoring

The most directly relevant information about the quantum threat timeline comes from the experts' responses 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 3.2 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 11 and Figure 12 we provide a graphical representation of the estimates made by the individual respondents, with the second figure comprising only responses from respondents closer to experiments. It is possible to appreciate the ample range of opinions, with some experts being optimistic and some others being relatively pessimistic about the rate of development of quantum computers.

The experts' responses are analyzed in a more coarse-grained fashion in Figure 13 and Table 1, and in Figure 14 and Table 2. Despite the great variability of the responses, some valuable patterns emerged.

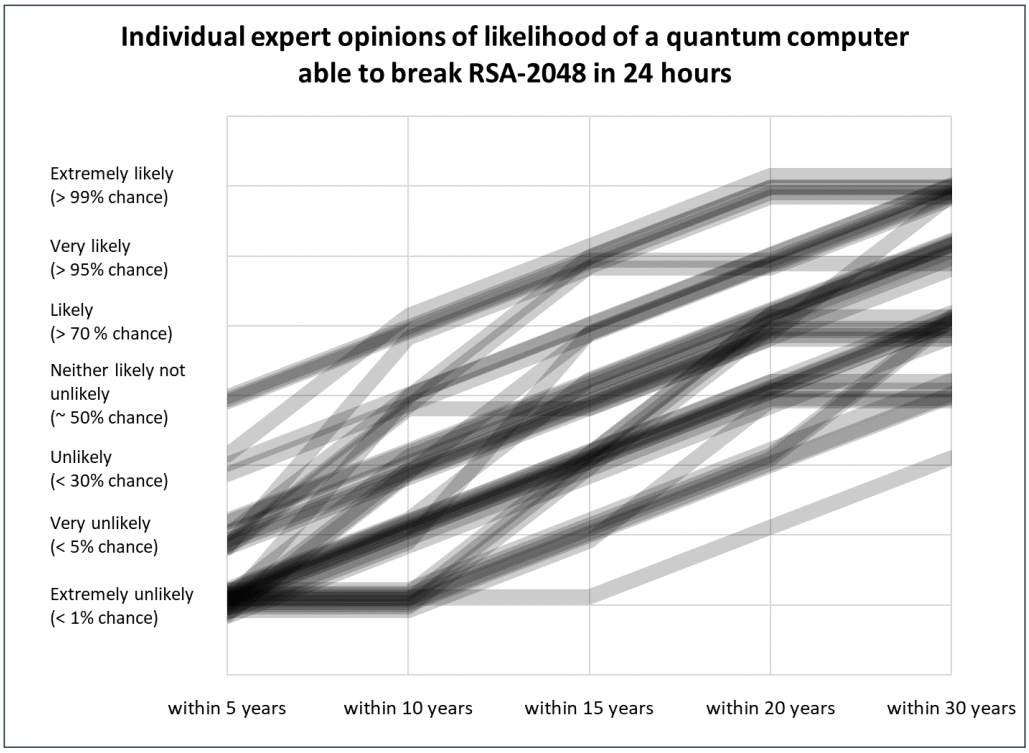

Figure 11: Opinions of individual experts about the likelihood of having a quantum computer able to factorize a 2048-bit number—that is, able to break RSA-2048—in, at most, 24 hours. Partially opaque lines represent the evolution of the likelihood assigned by each individual expert, with resulting total opacity dependent on how many opinions evolve according to that 'trail' in going from short-term to long-term.

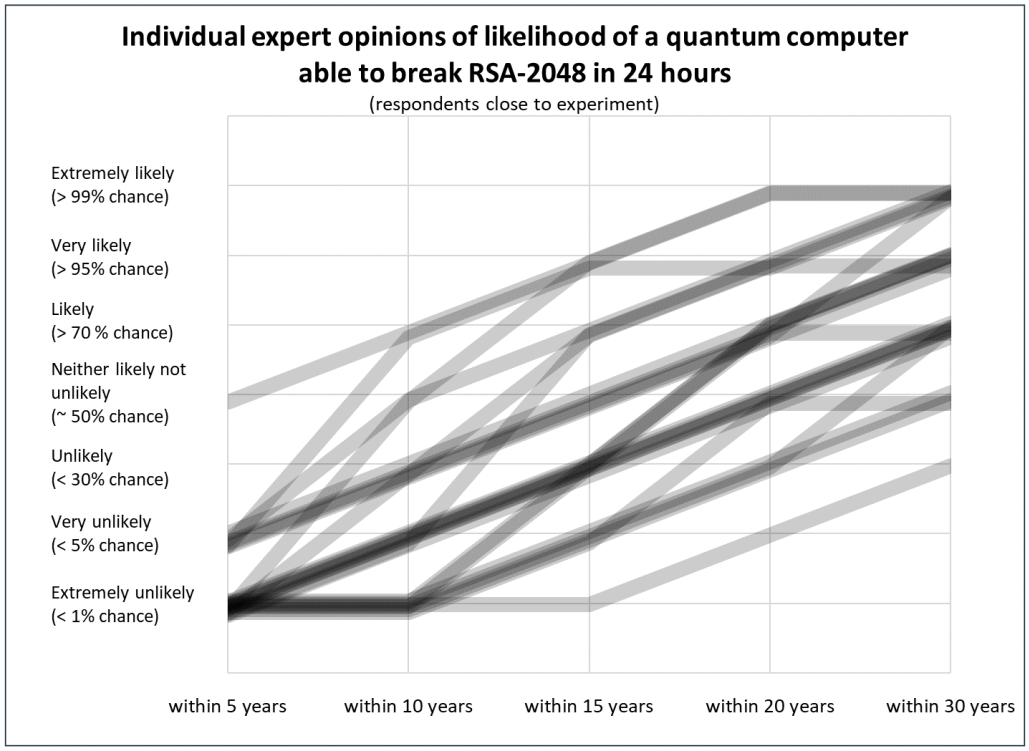

Figure 12: Opinions of only the individual experts close to experiments about the likelihood of having a quantum computer able to factorize a 2048-bit number—that is, able to break RSA-2048—in, at most, 24 hours. Each line represents the opinion of one expert.

Next 5 years:
Most experts (27/44) judged that the threat to current public-key cryptosystems in the next 5 years is "<1% likely". A quarter of them (11/44) judged it relatively unlikely ("<5% likely"). The rest selected "<30%" (3/44) or "about 50%" (3/44) likely, suggesting there is a non-negligible chance of a short-term and impactful surprise.

Next 10 years:
Still more than half of the respondents (23/44) judged this was "<1%" or "<5%" likely, but 11/44 felt it was "about 50%" or ">70%" likely, suggesting that the quantum threat could very well become concrete in this timeframe. In addition, 9/44 experts judged the likelihood "<30%".

Next 15 years:
Slightly more than half (23/44) of the respondents indicated "about 50%" likely or more likely, among whom 5 indicated a ">70%" likelihood, and 7 an even higher ">95%" likelihood. Still, almost half of the respondents (21/44) indicated "<30%" or even less likely.

Next 20 years:
About 86% (38/44) of respondents indicated "about 50%" or more likely, with 12 feeling it was ">95%" or ">99%" likely. This suggests that the chances of the quantum threat are more than even before the 20-year mark.

Next 30 years:
All but one among the experts responded that the quantum threat has a chance of "about 50%" or more, with 36 out of 44 experts indicating that the quantum threat will be likely (">70%"), very likely (">95%") or extremely likely (">99%").

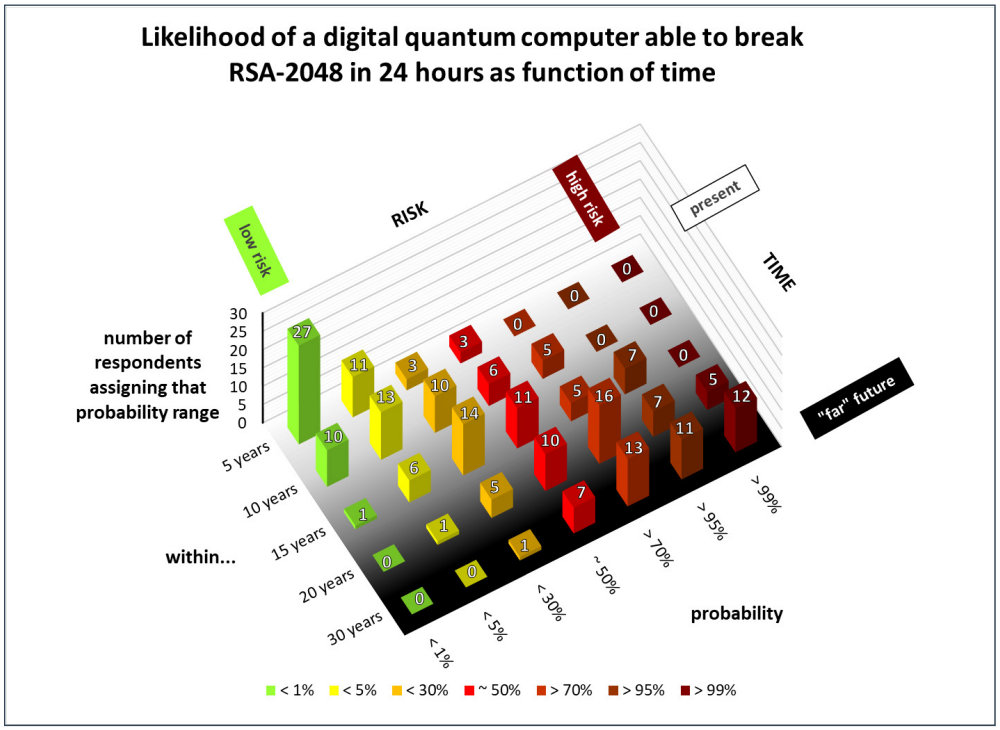

Figure 13: Number of respondents that have indicated a certain likelihood that a quantum computer able to factorize a 2048-bit number in less than 24 hours will be built within a certain period of time. See also Table 1 for data and informal 'likelihood' wording.

	Period of time				
How likely	***5 years***	***10 years***	***15 years***	***20 years***	***30 years***
Extremely unlikely (< 1% chance)	27 (61%)	10 (23%)	1 (2%)	0 (0%)	0 (0%)
Very unlikely (< 5% chance)	11 (25%)	13 (30%)	6 (14%)	1 (2%)	0 (0%)
Unlikely (< 30 % chance)	3 (7%)	10 (23%)	14 (32%)	5 (11%)	1 (2%)
Neither likely nor unlikely (about 50% chance)	3 (7%)	6 (14%)	11 (25%)	10 (23%)	7 (16%)
Likely (> 70 % chance)	0 (0%)	5 (11%)	5 (11%)	16 (36%)	13 (30%)
Very likely (> 95% chance)	0 (0%)	0 (0%)	7 (16%)	7 (16%)	11 (25%)
Extremely likely (> 99% chance)	0 (0%)	0 (0%)	0 (0%)	5 (11%)	12 (27%)
Total number of respondents (percentage)	*44 (100%)*	*44 (100%)*	*44 (100%)*	*44 (100%)*	*44 (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 13 for an intuitive graphical representation.

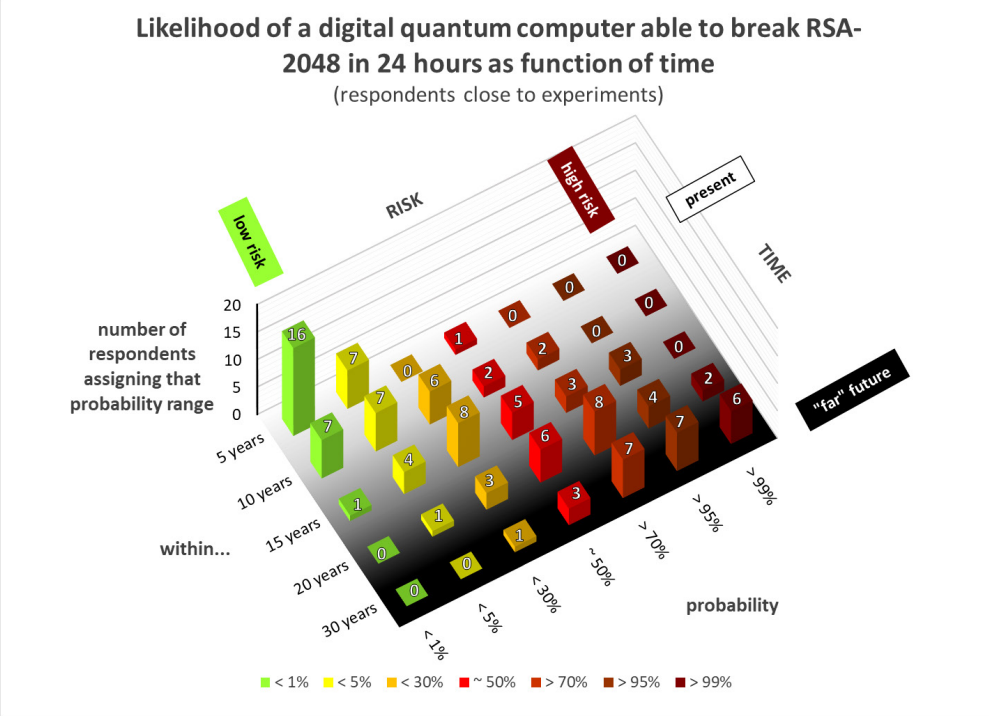

Figure 14: Number of respondents who are close to experiments and have indicated a certain likelihood that a quantum computer able to factorize a 2048-bit number in less than 24 hours will be built within a certain period of time. See also Table 2 for data and informal 'likelihood' wording.

	Period of time				
How likely	***5 years***	***10 years***	***15 years***	***20 years***	***30 years***
Extremely unlikely (< 1% chance)	16 (67%)	7 (29%)	1 (4%)	0 (0%)	0 (0%)
Very unlikely (< 5% chance)	7 (29%)	7 (29%)	4 (17%)	1 (4%)	0 (0%)
Unlikely (< 30 % chance)	0 (0%)	6 (25%)	8 (33%)	3 (13%)	1 (4%)
Neither likely nor unlikely (about 50% chance)	1 (4%)	2 (8%)	5 (21%)	6 (25%)	3 (13%)
Likely (> 70 % chance)	0 (0%)	2 (8%)	3 (13%)	8 (33%)	7 (29%)
Very likely (> 95% chance)	0 (0%)	0 (0%)	3 (13%)	4 (17%)	7 (29%)
Extremely likely (> 99% chance)	0 (0%)	0 (0%)	0 (0%)	2 (8%)	6 (25%)
Total number of respondents (percentage)	*24 (100%)*	*24 (100%)*	*24 (100%)*	*24 (100%)*	*24 (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 14 for an intuitive graphical representation

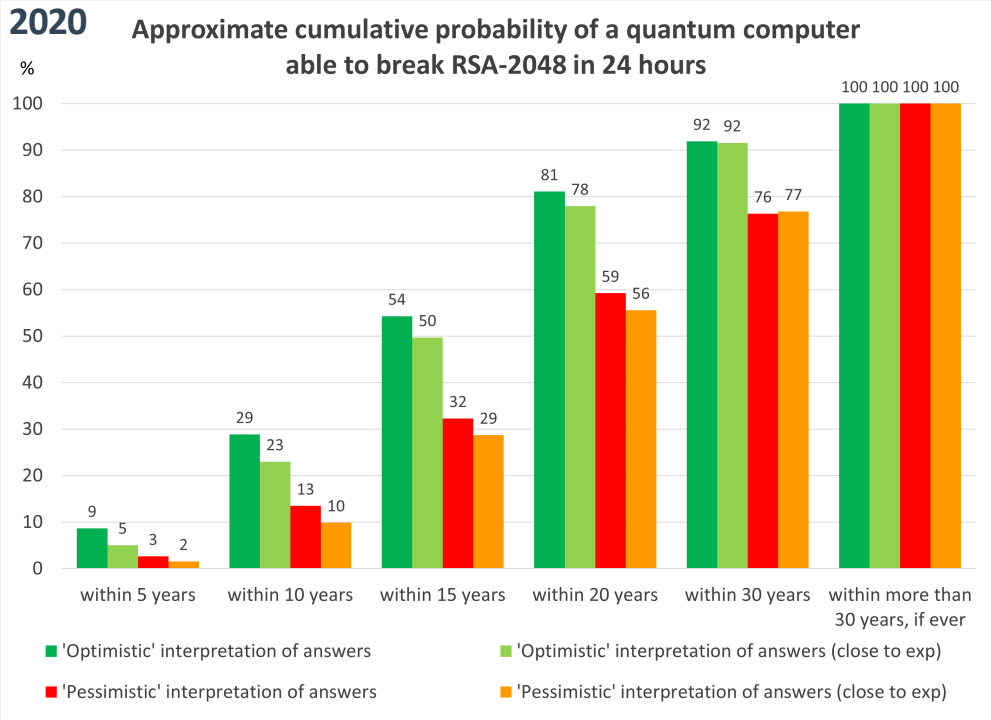

Figure 15: Cumulative probability for the creation of a quantum computer able to break RSA-2048 in 24 hours in a certain number of years in the future, based on the data of Table 1 and Table 2. The series correspond to 'optimistic' and 'pessimistic' interpretations of the responses, given the probability/ likelihood intervals given. We have considered, separately, the answers of those respondents closer to experiments. See Appendix for more details.

Note of correction, January 2022
A previous version of the 2020 report contained a calculation error that affected Figure 15, where it comes to the average probability associated to an optimistic interpretation of the answers provided by the respondents. This has now been corrected.

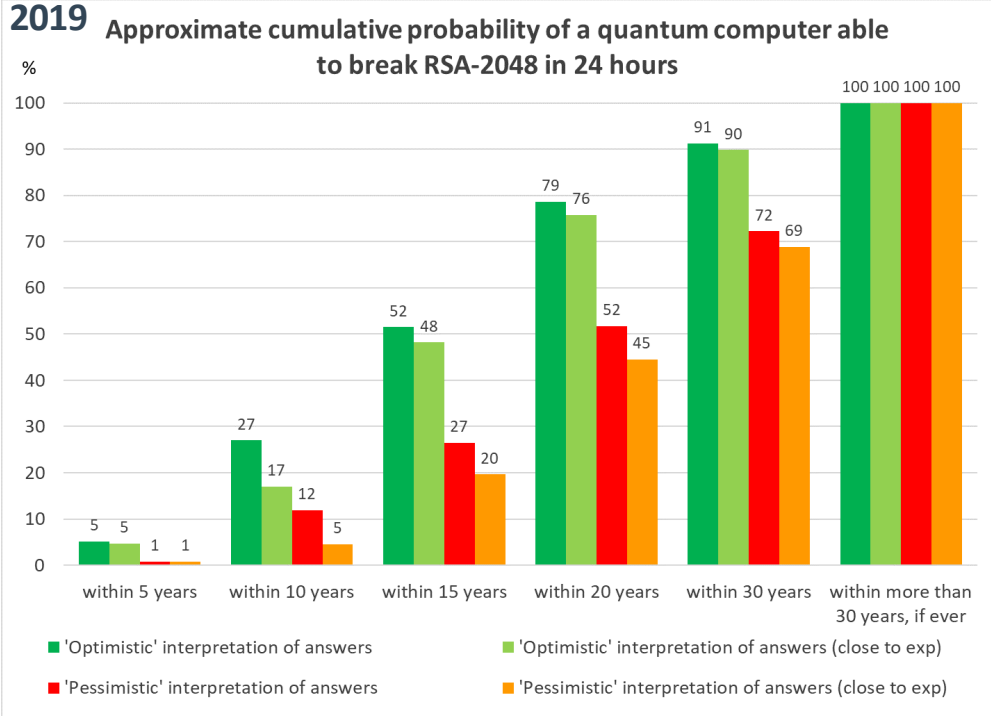

Figure 16: The same kind of plot as in Figure 15, as it appeared in the 2019 report.

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 16 in the Appendix.

Our Figure 15 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 13% in the next 10 years, and growing steadily in the years that follow; the estimated probability is more than 30% by the 15-year mark, and almost 60% by the 20-year mark, even in the pessimistic interpretation. In Figure 15 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).

4.1.2.1. Comparison with Last Year

It is interesting to check how opinions about the likelihood of the quantum threat in time have changed since last year. In Figure 16 we reproduce the same kind of plot as in Figure 15 but from the 2019 report. It is immediately evident that opinions overall have become more optimistic about the timeframe in which a cryptographically relevant computer could be built.

As mentioned in the previous section, and as described in the Appendix, the way in which the cumulative probability distributions are calculated is relatively rough (Figure 15). It is therefore useful to compare the change in responses themselves. This comparison is presented in Figure 17 and Figure 18, respectively for all respondents (44 this year and 22 last year) and only for the 21 respondents who took part in our survey both in 2019 and 2020. Note that the comparison is based on the percentage, rather than the number, of respondents who chose a particular answer, so that it is possible to compare the opinions of all the 2020 respondents with all of those 2019 respondents.

Such comparisons reveal a general shift towards increased likelihood (from the perspective of this report, increased probability of the quantum threat), particularly in the medium and long term. On the other hand, in the short term the experts seem to now lean towards a decreased likelihood. These somewhat conflicting tendencies could stem from better understanding the hurdles towards building a quantum computer and approaches for circumventing them. This better understanding might increase confidence in its eventual realization and reduce concerns about fundamental new hurdles or show-stoppers potentially emerging. At the same time, greater understanding of the challenges might also reduce the expectation of a short-term breakthrough. Another factor that may affect the short-term predictions is the effect of the present ongoing pandemic, which the experts estimate will slow down progress in the short term.

We note that the quite high optimistic probability for "within 5 years" of Figure 15 could be an 'artifact' of the 'optimistic' weight (70%) assigned to the responses "about 50% chance".

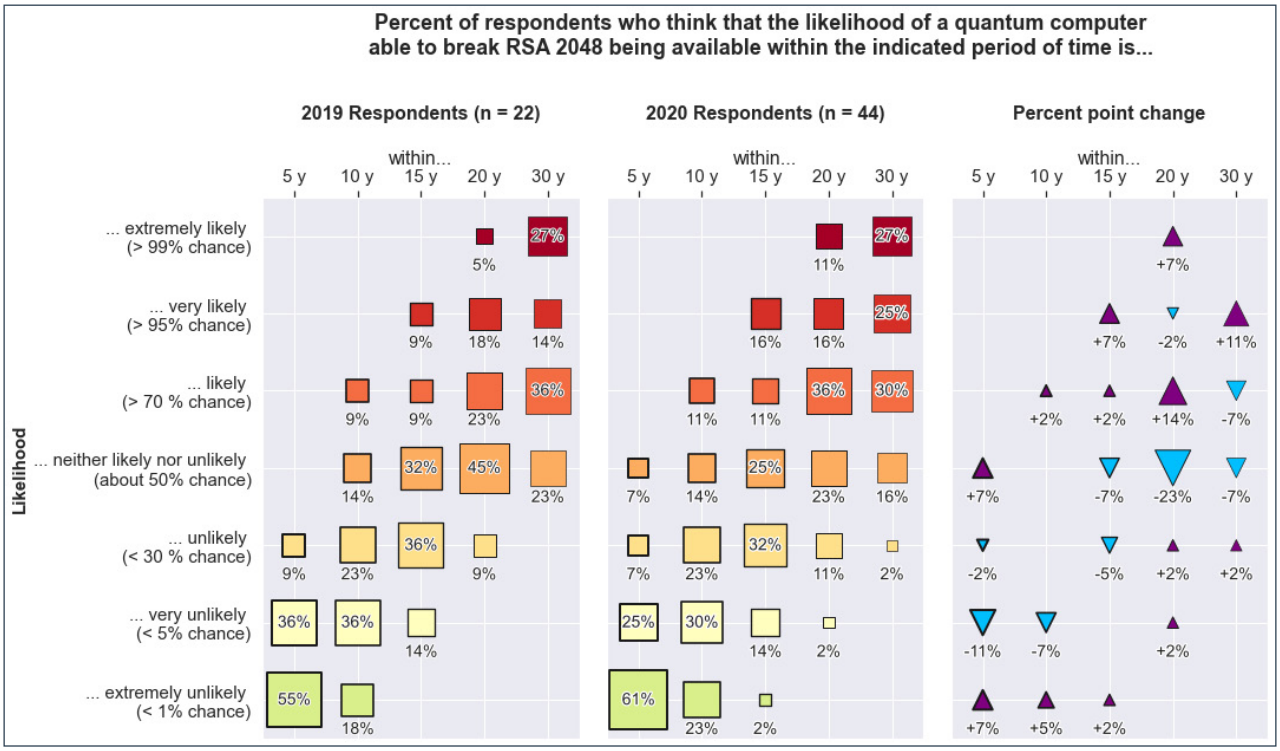

Figure 17: Change in opinion for the set of all respondents of the 2020 report with respect to all respondents in 2019.The left panel represents the distribution of choices for the all the 22 respondents of the 2019 report. The central panel represents the same, but for the 2020 respondents. The rightmost panel depicts the percent point change.

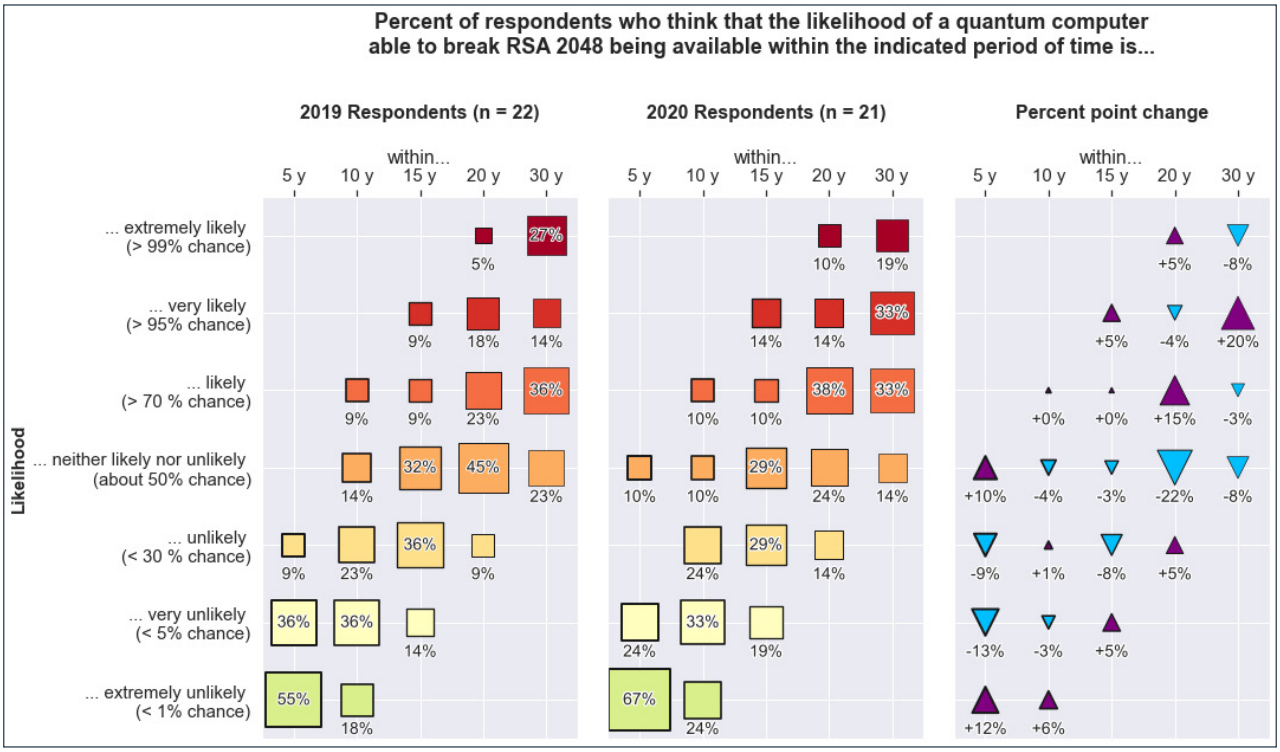

Figure 18: Change in opinion for the set of respondents that took part in both the 2019 and 2020 surveys. See caption of Figure 17 for details.

4.1.3. Logical Qubits and Fault-Tolerant Schemes

Arguably, the next major milestone towards building a fault-tolerant quantum computer is the realization of a fully controllable single logical qubit (see also Section 4.4). This would entail the ability to prepare, store, and manipulate a single logical qubit for an arbitrary number of operations, at least in principle, encoding the qubit in a sufficient number of physical qubits through some fault-tolerant encoding scheme.

The respondents largely indicate that such a milestone is close, with few even indicating that this will be demonstrated with high probability within one year, and a majority of respondents (30/44) indicating that this will happen with "about 50%" or more probability within three years (see Figure 19).[4]

Dave Bacon is optimistic about fautlt-tolerant proofs-of-principle but is conscious that, despite being a considerable achievement, they are just a step in the right direction:

Fault-tolerant demonstrations are just around the corner. However, the idea that you just need to get one fault-tolerant demo working and then it will be easy to scale up seems naive. [...] Getting to fault-tolerant quantum computation is critical to long term marketplace success against other forms of [computation] and it will be a long road to get there.

How to get there? According to many respondents, the surface code (with lattice surgery) is still the front-runner

as fault-tolerant schemes, especially when it comes to superconducting qubits. Nonetheless, other codes, progresses in/with other codes, and general efforts to improve fault-tolerant schemes were mentioned. Several respondents think that the best error-correcting mechanism will ultimately depend on the underlying physical platform, in several ways.

Dave Bacon writes:

The surface code is the most promising fault-tolerant architecture for systems that need to respect geometric locality in two dimensions. However [depending on the noise], there is some chance that other schemes will outperform the surface code.

Joe Fitzsimons highlights the importance of using the right code, and casts doubts on the choice of the surface code:

I think code choice is still the biggest open problem for scalable quantum computing. I am not entirely convinced that the focus on the surface code is entirely warranted at this point.

Bill Coish sees instead the surface code as being implementable in many systems:

I think there are several architectures that have potential to realize [the surface code] (ions, superconducting qubits, and spins). It's still not clear which will be the best [..].

4 A pre-print (not yet peer-reviewed) appeared after the end of our survey and during the write-up to this report, where a step towards a fault-tolerant logical qubit is implemented in an ion trap (Egan, 2020).

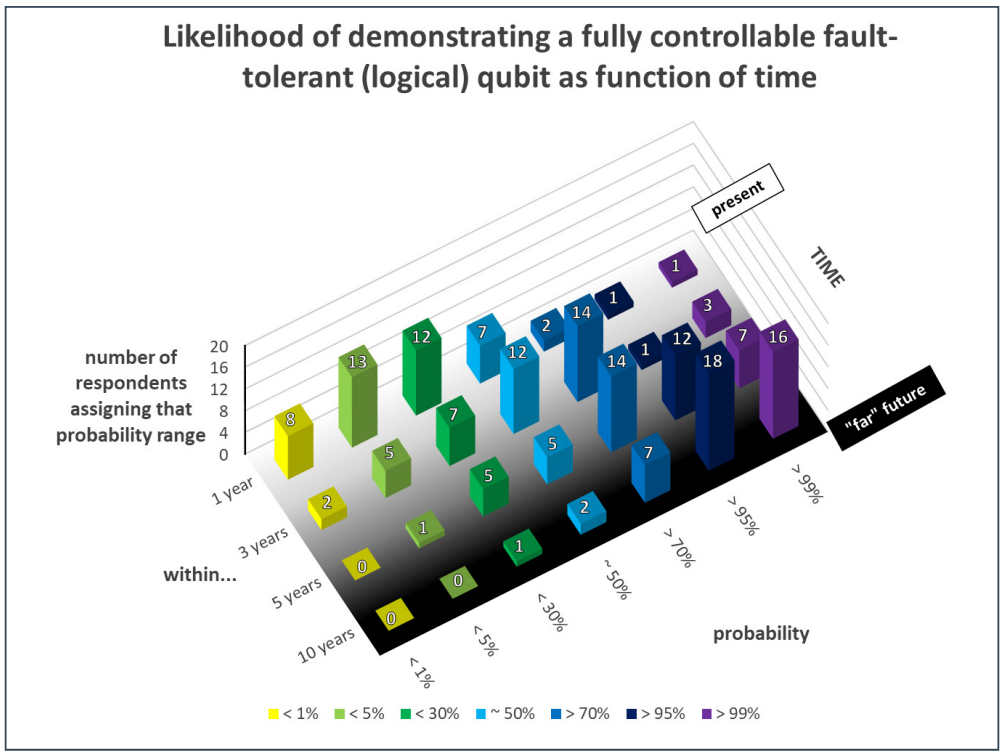

Figure 19: We asked our respondent to assign a likelihood to the demonstration of a fault-tolerant logical qubit. The consensus appears to be that this will happen relatively soon.

Daniel Gottesman has high hopes for so-called Low-Density Parity Check (LDPC) codes:

Based on current results, the surface code remains the leading candidate. However, I am extremely hopeful that further developments of fault tolerance based on LDPC codes will reach a point where they can be considered seriously. LDPC codes have the potential to greatly reduce the overhead needed for fault tolerance, but either require long-range gates (plausible perhaps in photonic systems and perhaps ion traps), or some further theoretical progress.

Ashley Montanaro also thinks that there is room for improvement:

I feel that further progress on fault-tolerance schemes is needed and expected before we reach the stage of fault-tolerant quantum computation.

One respondent puts it very simply:

I think the most promising scheme has not yet been discovered.

Speaking of alternatives to the surface code and its variants, Andre Morello writes of so-called bosonic codes, intended for bosonic systems with infinitely many energy levels:

As a solid-state experimentalist, I have long ignored bosonic codes [..], and failed to realize the potential they have. Only recently I began to understand them, and I have become more optimistic about the realization of a fault-tolerant system using such codes.

Another experimentalist comments in general about the importance of connectivity (see also Daniel Gottesman's comment above):

The choice of architecture/implementation will depend strongly on whether experimental platforms are able to introduce new connectivities between qubits [..] or reduce the error associated with certain connectivities.

Developments in the last year have not changed Stephanie Simmons' opinion about the necessity to operate at different levels when it comes to real systems in the lab:

I still expect a layered approach incorporating a combination of decoupling, refocusing, error-detection, and error-correction techniques to be the most promising.

4.1.4 Level of Funding for Quantum Computing Research

As mentioned in Section 1.3.4, the present level of investment in quantum technologies and in particular in quantum computing is at an historical high. This is an important fact, because high and sustained levels of investment are needed to deliver the promise of a full fault-tolerant quantum computer.

As world leaders in the field, involved in national and international projects and collaborations, working/consulting for industry and leading start-ups, our respondents have a significant vantage point to estimate the evolution of funding. We asked them to forecast what is likely to happen in the next two years.

A large majority of the respondents (28/43) expect investments towards quantum computing to increase in the next two years, with 10/43 expecting a significant increase (see Figure 20). Only 2/43 respondents foresee a decrease, and 12/43 see, in any case, a level of investment that stays approximately the same. Within the period considered, the global pandemic raises some uncertainty.

Reasons provided to expect continued or growing funding comprise:

- an ongoing quantum race, both at the level of companies and of countries,
- quantum computing being part of stimulus packages in the wake of COVID-19,
- partial but significant transition from fundamental science/a mostly scientific-driven field to an engineering-driven field, and
- venture capitalists' interest in looking for the "next big thing".

Respondents emphasize how there is currently "momentum", driven also by the field being strategic:

There is currently an important momentum in the field of quantum computing, and I believe that it will still increase in the next few years, with the ambitious plans of some countries boosting similar plans in others, etc. Quantum computing is perceived as a strategic field that attracts attention from both governments and industries and I think this trend will continue, especially if important milestones are reached regularly.

A 'bullish view' is expressed by Andrea Morello:

I see a great enthusiasm towards funding and investing in quantum technologies. There is a lot of capital looking for the next big thing, and quantum is an obvious one.

Elham Kashefi's tone is also very positive, and points to the fact that the arena got bigger and more crowded:

We are in a global race, milestones both on hardware and applications have been achieved and the quantum ecosystem all over the world has been growing and expanding so I can only envision further growth for the field now that a wide range of stakeholders are entering the field.

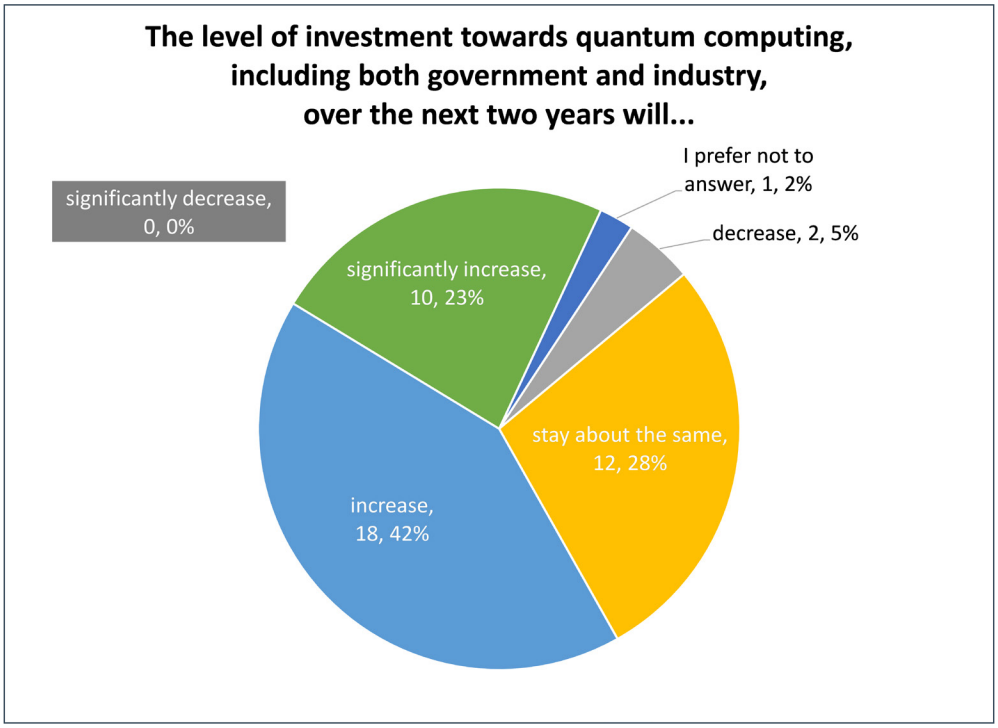

Figure 20: Expected change in the level of investment towards quantum computing in the next two year. While the option "significantly decrease" was given, none of the respondents chose it, and only two out of 43 respondents indicated they expect a decrease.

With respect to governmental support, Stephanie Simmons writes:

Governments move slowly enough that I imagine that current quantum hardware investments from government will continue, but that no large additional giant funds will be allocated.

Besides the role of governments, Jungsang Kim stresses that:

the private sector investment could increase substantially if the early quantum computer prototypes generate some traction with the R&D community in both academia and industry.

Ashley Montanaro predicts that:

private investment will continue to increase, and government investment will stay approximately the same (following recent very large investments in this area).

The importance of reaching some milestones—and the dangers of not doing it—are stressed by Nick Menicucci:

The major increase already happened. [..] I see investment increasing at a moderate pace in the near future (next <5 years), but beyond that it could go in any direction. For instance, investors could eventually grow weary if roadblocks are not overcome, leading to a "quantum winter" until a breakthrough occurs that reshapes the landscape once again.

Another respondent writes that:

in the near future the investment is likely to peak due to the amount of over-promise in the field and especially among start-ups.

Stephanie Simmons points out a distinction between venture-capital-funded and corporate initiatives:

Industry will go through dramatic shifts in quantum investment. Some existing quantum start-ups may survive, however patient capital will be become increasingly rare over the next year. [..] Corporate efforts and corporate venture arms will continue to think longer-term, and I imagine these kinds of investments will stay relatively stable.

A respondent highlights some potential dangers as well as some benefits of "consolidating" the field:

My feeling is that there is so much invested at this point is will be difficult for the momentum to dissipate in two years. However, I do suspect that the industry could collapse if any of the quantum start-ups go bankrupt. Alternatively, if a consolidation of some of the

quantum firms occurs, it will limit frivolous claims and solidify the path for the quantum technology industry.

A respondent stresses that the major way to avoid a reduction in funding is by delivering results:

It is important for the scientific community to deliver impactful quantum technology and applications in the next few years to keep the public interest engaged in this field. It is important for quantum science and technology to develop broadly including the development of fault-tolerant computing.

The COVID-19 pandemic may conceivably lead to different trends in funding (see also Section 4.5).

Bill Coish thinks that:

we are currently in (and will continue to be in) a 'recovery' mode through the pandemic, where resources are made available by governments, but I expect a shift away from perceived discretionary spending over the longer term. Unfortunately, I think 'discretionary' will include quantum computing unless there is a practical application (e.g. breaking RSA) in a very short time frame [..].

Yvonne Gao puts it in terms of a generic problem of resource allocation:

With COVID-19, I think many public-sector research funding will be shifted towards relief measures, vaccine developments, digital transformation, etc. in the near future, Quantum computing will remain an active and thriving field but will likely experience a slight reduction in funding due to the other more urgent priorities.

Nonetheless, one respondent highlights how:

[q]uantum technology has been cited by a number of governments as an area of investment for [post-COVID-19] stimulus.

4.2. Significance of Having Achieved Quantum Supremacy

One of last year's questions regarded the timeline for achieving quantum supremacy. Given Google's claim to have succeeded in accomplishing this (see Section 1.3.2 and (F. Arute et al., 2019)), we asked the experts to weigh in on the significance of its work.

Generally, the consensus is that such an experiment gave rise to an extraordinary result, from both a scientific and an engineering point of view, independently of any debate about its effective achievement of quantum supremacy.

Another point made by several respondents is that quantum supremacy itself is still extremely far from the holy grail of building and running a fault-tolerant quantum computer. One respondent writes:

I think it is an amazing experiment in that it is the first [one] to probe quantum physics of such a large number of particles, [exploring a] truly high-complexity regime, and many-body entanglement. In some sense, [I think it is] the strongest evidence we have [..] that quantum many-body [systems are] indeed described very accurately by good old quantum mechanics. In other words, it gives strong evidence that if there was no noise in the system, quantum mechanical systems could indeed [act as quantum computers].

Sergio Boixo, who was involved in (F. Arute et al., 2019), says of the experiment:

[It is the] demonstration that quantum computers offer a fundamentally different computational resource from all other computing platforms. It shows significant progress towards achieving a fault-tolerant quantum computer, although many hard problems remain.

Alexandre Blais comments:

Although the claim that quantum supremacy has been achieved can be challenged, this paper is in my opinion an impressive demonstration of the state of the art in the field. However, it is not clear how it impacts progress towards fault-tolerant quantum computation. The biggest impact is probably to give confidence that progress has happened and will continue to happen.

Some respondents pointed to the significance for fault-tolerance as an essential step was not shown in the paper. An expert writes:

An essential aspect of surface code error correction which was not demonstrated here was repeated measurement and post-processing / feedforward needed for error correction.

Nonetheless, they continue:

several aspects of the experiment are relevant for fault-tolerant quantum computation, namely the low error rates, and the repeated layers of single and two qubit gates in a lattice architecture (which is similar to that needed for surface code computation).

Another expert emphasizes the importance of the kind of noise the experiment dealt with:

the experiment does not differentiate between noise which is absolutely detrimental for quantum computation, for which no fault tolerance can help, versus noise which is in fact local, and allows for quantum fault tolerance,

and concludes that for such a reason:

[the experiment] says almost nothing [..] about the power and realizability of large-scale quantum computers in the presence of noise. Likewise, it says almost nothing about prospects for fault-tolerant quantum computation.

The role of noise is emphasized by another respondent, who points out how the noise may affect the claim to have achieved quantum supremacy, as there are analyses that show that:

sufficiently noisy quantum computers are easy to simulate compared to low-noise or error-corrected quantum computers.

The same respondent, as others, stresses another issue with the notion of quantum supremacy itself, rather than necessarily with the experiment itself:

the problem solved in the Google experiment is of no immediate practical relevance; rather, it was constructed for the sole purpose of demonstrating quantum supremacy. It was designed [..] for the demonstration to go ahead using a small number of noisy qubits, and without necessitating error correction.

Despite these limitations, a respondent describes the work as a source of inspiration:

This is a tour-de-force experiment (and analysis) that provides inspiration for all future developments in this area.

Artur Ekert points out that the well-publicized milestone has contributed also with:

lots of "constructive" hype that helps funding similar research elsewhere.

Many of the experts point out that further progress can go in at least two directions from the "line in the sand" that the experiment created:

- design/utilize NISQ devices for useful tasks (going beyond the mere lack of classical simulatability) that could lead to commercial applications;
- continue along the path towards fault-tolerance and a full-fledged quantum computer.

From the perspective of this series of reports, the latter is more significant, but the former kind of development would make it possible to tap into funding / sources of income that could sustain the work and the development of technology needed for the creation of a cryptographically-relevant quantum computer.

4.3. Recent Developments

When asked to name important recent developments (approximately since September 2019), many respondents refer to the achievement of quantum supremacy (see Section 4.2). Interestingly, Dave Bacon, one of the researchers involved in that result, writes that a major consequence of that result is that:

Google no longer has to focus on quantum computational supremacy.

We interpret this as the indication that, having achieved such a milestone, Google and other companies/groups can now proceed towards NISQ applications and the development of a real fault-tolerant computer.

It is unsurprising that several respondents indicate that there has not been any progress quite of the same significance as quantum supremacy. It appears that the last year has been devoted to steady and needed improvements, and to fostering the flourishing quantum landscape (see Section 1.3.4).

One respondent writes:

No single event comes to mind; the goals have been refined slightly with new code variants that adapt well to noise and new ideas [..], but broadly the field knows what it needs to do. The most interesting development is really the transition from (many of) the leading groups being university-based through to commercial, often by groups launching start-ups. This has brought in money but also "started the clock" on the need to deliver real value to avoid loss of investor confidence and a " quantum winter" effect.

Another respondent comments

I think a lot of the progress that needs to happen right now, and is happening, slowly, is on materials/engineering issues that are not going to make headlines in the short term.

Bill Coish highlights that

widespread efforts to develop control software and interfaces to 'real' quantum computers [..] have allowed for a democratization of quantum computing that could lead to big leaps forward. Most of these efforts have only become public in the last year, so we have yet to see the real impact of these initiatives.

On the theoretical side, several respondents point to developments in the design of error-correcting codes and fault-tolerant schemes that lower the demands on experimental control/precision. On the experimental side, there has been progress in the implementation of such kind of codes/schemes.

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 2021—many of our respondents mentioned the experimental implementation of some form of error correction or even the potential implementation of a fault-tolerant qubit (see Section 4.1.3).

One respondent set two broad goals that may facilitate this:

First, we need to establish what the real fault-tolerant thresholds are, including realistic noise models, and not consider asymptotic limits.

Second, [we need to improve] classical processing for the feed forward operations that are going to need to be done.

One respondent points out how often big steps need to be preceded by smaller steps:

I believe that real breakthroughs will be required both at the theoretical (fault tolerance) and experimental (system control) levels. So, I am more interested in possibly seemingly 'small' results that may open new avenues in the longer term. A surprise result coming from silicon or optical qubits could be very interesting.

Kae Nemoto sets platform-dependent goals:

For superconducting qubits, a better and fast measurement, and a better uniformity of qubits. For [hybrid architectures], a realization of entanglement between light and spin-qubits in a scalable device. For silicon, a scalable architecture design which can be experimentally demonstrated in a few years.

4.5. How COVID-19 is Affecting Quantum Computing Research

The respondents expressed various degrees of concern about how the COVID-19 pandemic is affecting quantum computing research.

A stark contrast is highlighted, overall, between experimental and theoretical work. Experimental work has been delayed most, with estimates going from a few months to one year and more. This is due to several factors, from restrictions on accessing equipment on campus to a partially disrupted supply chain. "Experiments have ground to a halt", said Bill Coish, a theoretician who often collaborates with experimentalists. Another respondent wrote:

For experimentalists typically two+ months of lab time were lost completely, and current arrangements are still at a reduced level impacting rate of progress. Overall, perhaps 2020 is an eight-month year in terms of progress.

Not all experimentalists have been affected equally, much depending on policies of universities as well as the availability of setups to operate experiments remotely. One writes:

In the short term, my academic research group was able to return to our work in the laboratory after about two months, and while we are still adapting to new policies, our research is currently progressing as it was before the pandemic.

Andrea Morello is even more positive:

My research has been so far largely unimpacted by COVID-19. Although my university did go into shut-down, we got permission to have sporadic attendance in the laboratories, whereby we were able to conduct the few operations that need the physical presence of a researcher (connect instruments, fill cryogenic liquids, etc.). For the rest, the experiments can be controlled via remote desktop from home. The clean rooms—arguably one of the safest places to be during a pandemic—were open at reduced capacity, and we managed to continue fabricating devices.

Frank Wilhelm, who also collaborates with experimentalists, points out that, even if device fabrication continues, it was nonetheless slowed down:

Laboratory remote control is quite advanced, but the fabrication activities for quantum computing were slowed down a fair bit.

That different experimentalists report different degrees of impact is emphasized by Alexandre Blais:

The negative impact is probably less for well-established groups which can, for example, run experiments remotely. The situation can be much more challenging for new faculty members.

In general, the activities of theoreticians have certainly been impacted less, but researchers still report notable differences in—and perceptions of—the impact, likely depending on also on the amount of remote teaching and bureaucracy. One respondent wrote:

People (including me) have more time to think, work, finish projects, as well as go deeper into ideas we always wanted to explore. [Teleconferencing] allows reasonable interactions, and so existing collaborations continue.

On the other hand, Bill Coish reported that:

[while] there may [..] be some benefits for particular individuals, [e.g.] due to lighter administrative requirements some people may have been able to spend more time on research, [...] this is likely true only in rare cases. For me, the added overhead associated with remote work has slowed the progress of my group substantially.

Many respondents, both experimentalists and theoreticians, expressed how a major impact of the pandemic has been the inability to interact in person, which hinders the training of students, collaboration, networking, and the serendipitous generation of ideas though discussions.

Yvonne Gao emphasizes that

the immediate impact is the restriction on the flow of talents in quantum computing.

Many experts express some concern about investments (see also Section 4.1.4). Frank Wilhelm-Mauch worries that

some companies, who saw quantum computing as a blue-sky luxury activity seem to be [reconsidering].

An optimistic respondent writes:

The most serious effect might come from funding redirected to other sectors, more relevant to the pandemic, but for the moment I feel that quantum technologies in general and quantum computing in particular are rather seen as a means for governments to showcase their commitment to innovation and sovereignty that are considered crucial in times of crisis.

Will there be, in general, a long-lasting impact that significantly slows down the development of a fault-tolerant quantum computer? Not necessarily, given the timescales involved. The same respondent writes:

I would say that given that the timeline of a fault-tolerant quantum computer is quite long-term the effects of the pandemic will hopefully be erased at that scale assuming the pandemic is under good control globally within one or two years.

4.6. 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[5]. 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,
- the dangers of hype and of high (and potentially, too short-term) expectations from funders, government and the public, and
- the difficulty of making predictions about the rate of development in the field.

Respondent:

It is an absolutely remarkable time! I think the biggest question is whether NISQ algorithms can be made in any way truly useful (for solving problems of true interest). [If this is not the case], the field might be in danger, since it needs huge money investments to keep it alive and going, so that it can evolve to the point where the threshold for fault tolerance is achieved. We need good milestones on the way, and if NISQ computations [do] not provide this motivation, we will need another intermediate goal - which I am not sure what it will be. The danger is that the time scale for the next truly exciting development - if NISQ algorithms fail that might be fault tolerant single qubit memory, for example - is too long for the short breath of the industry. But hopefully some other intermediate goal is found on the way...

Sergio Boixo:

In experimental quantum algorithms it remains important to study scalability and a path to surpass the performance of classical algorithms.

Sergio Boixo:

In the NISQ era there is an opportunity to achieve practical applications with improved quantum algorithms and error-mitigation techniques. Advancements here will also be used with early fault tolerant quantum processors.

Frank Wilhelm-Mauch:

We are addressing the optimization of given hardware by making gates that are fast (hence avoiding incoherent error) and precise (avoiding coherent error). A key obstacle is that theoretical work is reliant on a good and precise model, which we need to extract at the same time.

Respondent:

[We] need an overwhelming case to demonstrate the advantages of quantum computing, i.e., [we] must demonstrate some task that is considered practical by a wider community [beyond the] quantum computing community.

Daniel Gottesman:

I remain extremely hopeful that we will continue to make good progress towards a large quantum computer. Eventually there will be a crash of sorts, with a winnowing out of many start-up companies and some large companies cutting back or eliminating quantum computing research. Hopefully, there will be survivors that will continue making progress (and most likely the survivors will be those who have made the most progress). I am concerned that the COVID recession could bring this crash forward to the next year or two.

Artur Ekert:

Expect [the] unexpected

Respondent:

For all platform, independent of what they are, there will be a plateau at around 300 physical qubits. Beyond that qualitatively new techniques must be developed.

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

Yvonne Gao:

It's an extremely dynamic field and we are witnessing relentless progress globally in quantum computing but also other technologies that stem from understanding we gain from the current quantum devices. Going forward, I think this field will become increasingly inter-disciplinary and will lead to many exciting new ideas and innovations.

Jungsang Kim:

I personally do not believe that the cryptography application will drive the development of near-term quantum computer technology, but there will be sufficient applications that are currently unknown that will stimulate the continued technology and market development.

Jungsang Kim:

We are [ramping up] our research effort to build modular quantum computers that connect 2-4 such systems together, which will lead to ~100+ qubit systems, and show a path towards quantum computers with thousands (if not hundreds of thousands) of qubits. The major bottleneck in the progress, beyond demonstrating key protocols and their integration into our systems, is the technology to build these systems with high enough reliability (manufacturability) and low enough cost to demonstrate scale.

Ashley Montanaro:

This is an extremely exciting time for quantum computing — quantum computers have outperformed our best classical supercomputers, and it is plausible that within a 2-5 year period, NISQ devices could solve problems of real practical importance beyond the capacity of existing methods. Large-scale fault-tolerant quantum computing is further distant but seems achievable on a 10-15 year timescale. A key question remains to determine the best applications for quantum computers within both of these eras.

Andrea Morello:

There are at least two tangential reasons to be optimistic about progress in quantum computing. One is that a large ecosystem of classical control devices and system, backed both by start-ups and by established companies, is entering the arena. These companies are stepping up the game in producing instruments to help the research scientists working on quantum computer scale-up. Cryogenics, optics, electronics, all are doing their part to help out. It's great. The other one is that the universities are taking notice of the growth in demand for quantum-trained engineers, and starting to create degrees targeted to creating a pipeline of skilled workers. This will take a few years to show its effect, but eventually the new quantum-trained talent will make a huge difference.

Bill Coish:

There's a mismatch between the digital, discrete formalism most often used in quantum computer science and the reality of continuous-time processes leading to analog outcomes. There are (sometimes very large) advantages that can be found by adapting quantum protocols to the analog experimental ('real') world, rather than forcing the experiment to approximate a digital process.

Kae Nemoto:

NISQ machines are not really computers, but the biggest quantum coherent system the humans manage to control. It is important to explore what it means to us, our physics, as well as to find potentially interesting and important problems to run on it in a better way. The requirements for fault-tolerant quantum computer are different, and the development also needs to be based on its own requirements.

Respondent:

[C]urrently we still have little understanding about how to use NISQ devices properly to get a quantum advantage.

Respondent:

How to make quantum software (running on NISQ devices or simulators) widely accessible is of crucial importance to the wider IT community.

Respondent:

I'm worried about excessive hype. Though I am optimistic about the long-term impact of quantum computing, I feel that there are widespread unrealistic expectations regarding the time scale for reaching useful applications.

Lieven Vandersypen:

These are exciting times! Compared to five years ago, my belief that large-scale fault-tolerant quantum computers are feasible, has significantly grown. This renewed optimism is based on the rapid progress we have achieved as a community with semiconductor spin qubits, as well as on the rapid progress in adjacent experimental fields.

SUMMARY AND OUTLOOK

A fully working quantum computer can be seen as the 'holy grail' of quantum technologies, but also as a major threat for cybersecurity. For example, it is a threat for cryptosystems based on the difficulty of mathematical problems that are solvable by a quantum computer large and reliable enough to run the appropriate quantum algorithms.

The quest to build such a fully scalable and fault-tolerant quantum computer 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. Such a competition has substantially heated up in recent years and months, with the entry of new major private players, large grants from governments, and the birth and growth of many start-ups fuelled by venture capital. It has also been described as a marathon, rather than a sprint race. Nonetheless, there could be sudden accelerations. One survey respondent puts it nicely:

I think the path leading to a fault-tolerant quantum computer is fascinating and will be full of obstacles and surprises. It is extremely hard to make predictions in the field [..]. Very beautiful science will happen to lead to such a computer, and it will also need to be combined with excellent engineering.

Another respondent expresses how even an expert may be surprised at times by the relatively fast progress of quantum computing:

It is not always the case [..] but I find that my predictions are often more pessimistic than what actually happens. I take this as a sign that the research is accelerating.

The large group of experts we have polled judged that the development of a quantum computer that could break a scheme like RSA-2048 within the next 10 years is relatively unlikely, but its likelihood is far from negligible. Furthermore, the experts indicated that the chance of a quantum threat emerging within 15 years is significant. Comparing this year's opinions to the results of the survey we conducted one year ago, it appears that the experts have become a bit more optimistic about the timeline for quantum computing, suggesting the quantum 'threat' is an increasing concern for cyber-security.

Whenever one deals with opinions rather than hard facts, it is appropriate to consider how reliable or partisan such opinions might be. After all, our respondents are generally devoting their careers to quantum information science and quantum computing. One could therefore wonder whether they are necessarily biased toward believing in the possibility of realizing a fault-tolerant computer, potentially relatively fast. We cannot exclude this, but we are confident that our respondents have tried to provide the best possible estimates, based on their expertise. Working in a field that pushes the limit of what humans are capable of (as quantum computing corresponds to changing the paradigm of computation) requires some optimism. It also requires a deep critical capacity to be able to identify and overcome roadblocks.

The logical possibility that consequential quantum cryptanalysis is, for some reason, infeasible or impossible is captured in the small but non-negligible likelihood that quantumly breaking RSA-2048 will take over 30 years. Cyber-risk managers are naturally more concerned about the possibility that the quantum threat materializes early. It is up to each institution, company and manager to decide what risk they are ready to accept.

While building a cryptographically relevant quantum computer is a formidable task, it is important for people managing cyber-risk to understand that there is nothing close to a scientifically convincing or established argument for why efforts currently underway are likely to fail in the medium to long term. Progress last year, including the demonstration of quantum supremacy as well as the significant momentum of the field—in terms of activities, results, and resources—should probably trigger cautious reactions, directed to develop crypto-agility and resilience against quantum attacks.

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 necessarily be only one winner. Presently, according to the experts' opinions,

superconducting circuits and ion traps seem to have an edge over the competition, but surprises could come from other implementations and from the combinations of different technologies.

The demonstration of quantum supremacy was a milestone on the path towards the creation of a fault-tolerant quantum computer. While it did not say much *directly* about fault-tolerance, it demonstrated that the field has already developed the ability to operate on many qubit components while preserving their (joint) quantum behaviour. This is highly non-trivial and has stimulated further effort and commitment towards the next big steps, which will be to create and exploit Noisy Intermediate-Scale Quantum (NISQ) devices for concretely useful tasks, and, most importantly, to demonstrate error suppression via error-correction and fault-tolerance schemes to prolong the storage and manipulation of logical qubits.

From last year, we have doubled the number of experts who took part in our survey. The opinions we have collected 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 can estimate the time it has at disposal to implement post-quantum cryptographic solutions. A significant number of the experts felt the likelihood that quantum computers capable of breaking current asymmetric schemes materialize in 5 years is non-negligible and becomes substantial in the next 10-15 years. A respondent explains clearly how this is a call to action:

With respect to protecting confidentiality in the long term, it is important to be conservative and to assume such a scenario: Mitigating actions should be taken now by standardizing and rolling out post-quantum secure asymmetric 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, A Methodology for Quantum Risk Assessment, 2017).

The Global Risk Institute and evolutionQ Inc. will provide an update of this survey in approximately one year. This will allow us to 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.

Egan, L. e. (2020). Fault-Tolerant Operation of a Quantum Error-Correction Code. arXiv:2009.11482. Retrieved from https://arxiv.org/abs/2009.11482

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

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.

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

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

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).

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

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

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

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

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

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

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

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.

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

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

APPENDIX

List of Respondents

The respondents that took part also in the 2019 survey are listed at the top of this table (light grey background).

Name	Institution	Country
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 Physicsand Quantum Benchmark Inc	CAN
Jungsang Kim	IonQ, Inc. and Duke University	USA
Ashley Montanaro	PhaseCraft and 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	Jülich Research Center and Saarland University	GER
Shengyu Zhang	Tencent	CHN
Sergio Boixo	Google	USA
Dan Browne	University College London	GBR
Fernando Brandão	Caltech	USA
Michael Brett	Independent consultant	USA
Eleni Diamanti	CNRS and Sorbonne University	FRA
Joe Fitzsimons	Horizon Quantum Computing	SGP
Yvonne Gao	Centre for Quantum Technologies, National University of Singapore	SGP
Winfried Hensinger	Sussex Centre for Quantum Technologies, University of Sussex	GBR
Sir Peter Knight	Imperial College London	GBR
Elham Kashefi	School of Informatics, University of Edinburgh CNRS, LIP6, Sorbonne University	GBR
Yi-Kai Liu	National Institute of Standards and Technology (NIST)	USA
Klaus Moelmer	University of Aarhus	DNK
Bill Munro	NTT Basic Research Laboratories	JPN

Name	Institution	Country
Nicolas Menicucci	RMIT University	AUS
Kae Nemoto	National Institute of Informatics	JPN
John Preskill	California Institute of Technology	USA
Simone Severini	Amazon Web Services	USA
Lieven Vandersypen	QuTech and Kavli Institute of Nanoscience, TU Delft	NLD
David Wineland	University of Oregon	USA
James Daniel Whitfield	Dartmouth College	USA
Gregor Weihs	University of Innsbruck	AUT
Jun Ye	JILA, NIST and University of Colorado	USA
Peter Zoller	University of Innsbruck, and IQOQI Innsbruck	AUT

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 (with Zoller) the theory at the basis of trapped-ion quantum computation. He devised new methods to efficiently study quantum systems with classical computers, and to use controllable quantum systems (like cold atoms) as quantum simulators.

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

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

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

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

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

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.

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

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

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

Michael Brett
Presently an independent consultant, he was the CEO of the data analysis and quantum computing software company QxBranch, and later Senior Vice-President of Applications at Rigetti Computing.

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

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

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

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

Sir Peter Knight
He is a pioneer in the field of quantum optics. He is Chair of the UK National Quantum Technology Programme Strategy Advisory Board, chairs the Quantum Metrology Institute at the National Physical Laboratory, and is a science advisor to the UK government. His research centres on quantum optics and quantum technologies.

Elham Kashefi
A leading quantum cryptography researcher, renowned for her work on blind quantum computing. She is a professor at the University of Edinburgh, associate director of the Networked Quantum Information Technologies, a CNRS researcher at the Sorbonne University, and on the executive team of the Quantum Internet Alliance.

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

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

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

Nicolas Menicucci
A leading researcher who contributed key results for the development of continuous-variable cluster states and towards fault-tolerance in quantum optics approaches.

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

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

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

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

Gregor Weihs
He is Professor of Photonics at the Institute for Experimental Physics at the University of Innsbruck, where he leads the Photonics group. His research in quantum optics and quantum information focuses on semiconductor nanostructures and on the foundations of quantum physics.

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

James Daniel Whitfield
A leading expert in quantum algorithms, based in the Department of Physics and Astronomy of Dartmouth College. His research focuses on understanding the potential and the limitations of new and existing quantum computers to perform physical simulations.

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

Peter Zoller
A leading theoretical physics who works on quantum optics and quantum information. He is best known for his pioneering research on quantum computing and quantum communication and for bridging quantum optics and solid-state physics. He established (with Cirac) the theory at the basis of trapped-ion quantum computation.

QUESTIONS

Questions 1-7 involved identification of the respondent and gauging their familiarity with different subfields of quantum computing research as well as implementations.

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 (including integrated photonics), Quantum spin systems in Silicon, Quantum spin systems not in Silicon, Topological Systems and Cold Atoms.

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

QUESTIONS 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 (including integrated photonics), Quantum spin systems in Silicon, Quantum spin systems not in Silicon, Topological Systems, Cold Atoms, Other

QUESTION 11

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.

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 single fully controllable fault-tolerant (logical) qubit will be demonstrated 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

Please comment on Google's recent demonstration of so-called "quantum supremacy" [Arute, F., Arya, K., Babbush, R. et al. Quantum supremacy using a programmable superconducting processor. Nature 574, 505–510 (2019). https://doi.org/10.1038/s41586-019-1666-5] and how it impacts the progress towards achieving a fault-tolerant quantum computer.

QUESTION 15

How do you assess the impact that the COVID-19 pandemic will have on the progress being made towards achieving a fault-tolerant quantum computer? Feel free to elaborate on how your own work has been affected and the impact you expect it to have on academia and industry.

QUESTION 16-17

You think that, over the next 2 years, the level of investment (both by government and by industry) towards quantum computing will…

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

Question 17 asked for free-form comments on the issue.

QUESTION 18

What has been the most significant recent (approximately since September 2019) achievement in the progress towards building a fault-tolerant quantum digital computer?

QUESTION 19

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 20-21

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

QUESTION 22

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 3.2 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 4.1.2, we assigned the following cumulative probabilities to each response, which are the largest and smallest ones compatible with the ranges among which the respondents could choose:

Optimistic Assignment:

Extremely Likely (> 99% chance)	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 15, 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.