

# State of the art in fault-tolerant quantum computing

Questions and issues

# Academy report



Académie des technologies Le Ponant - Bâtiment A 19, rue Leblanc 75015 PARIS

+33(0)1 53 85 44 44

secretariat@academie-technologies.fr www.academie-technologies.fr

© Académie des technologies ISBN: 979-10-97579-63-0

Cover design: EnigmaEasel - Adobe stock



# STATE OF THE ART IN FAULT-TOLERANT QUANTUM COMPUTING

QUESTIONS AND ISSUES

Academy report

This report was written under the direction of Catherine LAMBERT, Thierry BONHOMME and Gérard ROUCAIROL.

It is the product of a working group, the full composition of which is specified in the appendix. It follows a rigorous process within the Academy of technologies, guaranteeing its independence and objectivity; it was approved and voted on at the General Assembly of the Academy.

May 2025

## CONTENTS

Cor	ntext	6
Executive summary		7
Résumé exécutif		11
Introduction		15
Chapte The	er 1 reality of quantum advantage and its needs	20
1.5. 1.6. 1.7. Chapt	Quantum algorithm Universal computer Fault-tolerant computing Theoretical complexity Quantum acceleration Heuristics and empirical approaches Use cases ter 2 r-correcting codes	20 22 23 25 27 30 31
2.1. 2.2. 2.3. 2.4. 2.5. 2.6.	Fault tolerance, the only route to deep quantum algorithms Error correction and fault tolerance The principle of quantum error correction The golden age of surface code New approaches requiring less hardware The challenges ahead	43 45 46 47 51 60

Chapter 3 Qubit technologies: challenges and roadmaps		
<ol> <li>3.1.</li> <li>3.2.</li> <li>3.3.</li> <li>3.4.</li> <li>3.5.</li> </ol>	Neutral atoms Superconducting Qubits Photons Spins in silicon Trapped ions	70 83 98 114 130
Chapter 4 Scaling up quantum computers		142
4.1. 4.2. 4.3.	Scalability through modularity Photonic interconnection Scalability and connectivity of superconducting qubits	143 150 163
Chapter 5 Additional information		174
5.1. 5.2. 5.3. 5.4. 5.5. 5.6. 5.7. 5.8.	Introduction Technical and economic analysis of the ecosystem - Historical comparisons A few questions Development of competing technologies on the horizon of the first FTQC? Application-oriented <i>benchmarks</i> Links with high-performance computing Human capital Financing strategies	174 174 177 180 183 189 195 196
Conclusion		198
Appendices		199
	<ul><li>A. Glossary</li><li>B. Members of the working group</li><li>C. People interviewed</li></ul>	200 216 218

### Context

This report was created by a committee set up by the National Academy of Technologies of France (NATF) with the aim of producing a reference document, to be completed and updated annually, assessing the feasibility and impact of *Fault Tolerant Quantum Computing (FTQC)*. In particular, it aims to identify the scientific, technological and economic challenges, to specify the use cases, the potential for value creation, the links with High Performance *Computing* (HPC) and the time horizons for the buildup of FTQC.

This first edition of the report focuses on the current state of the main qubit technologies developed in the world and in France, the state of the art in error correction demonstrations, and approaches for scaling up fault-tolerant quantum computing technologies. As such, it follows-on from a study carried out for the "Secrétariat Général pour l'Investissement" by the Académie des Technologies. It does not address the question of the potential strategic and economic interest in building such computers, a matter that will be addressed in a future report.

In 2019, the US *National Academies of Sciences, Engineering, and Medicine* published a report<sup>1</sup> on the state of the art of quantum computing. Like the American report, this report by the National Academy of Technologies of France (NATF) represents a rare initiative to provide a comprehensive overview of the difficulty to scale quantum physics experiments to build fault-tolerant quantum computers.

<sup>1 &</sup>quot;Quantum Computing: Progress and Prospects (2019) - <u>https://nap.nationalacademies.org/</u> catalog/25196/quantum-computing-progress-and-prospects

## **Executive Summary**

This report presents the results of a collective study carried out by a Working Group set up in June 2023 by the french Academy of Technologies to assess the feasibility and impact of Fault Tolerant Quantum Computing (FTQC). It follows on from a study carried out for the "Secrétariat général pour l'investissement" (SGPI).

The report begins with a review of the concepts of quantum algorithms and quantum computational advantage, and describes how this advantage varies according to algorithm classes and use cases. It highlights the issues and challenges involved in fault-tolerant quantum computing applications.

- There are strong theoretical arguments showing that a fault-tolerant quantum processor would solve certain classes of problems drastically faster than a classical computer, such as factoring large numbers and simulating quantum systems.
- For other problem classes, there are theoretical polynomial time speedups compared with the best classical algorithms, but it is less drastic. Practical usefulness will then depend on implementation details, notably the execution time of quantum gates.
- Initial resource estimates for algorithms dedicated to FTQC computers capable of solving industrial-scale problems are of the order of several thousand logic qubits and billions of logical operations, often with very long computation times. This illustrates the need to optimize the execution speed of quantum circuits and error correction.
- Industry potential use cases for FTQC are, as for now, little or poorly identified in France, probably because industry players have focused on exploring

applications for the small noisy quantum processors currently available or in the near future (known as NISQs) or analog quantum computers.

The report then highlights the fundamental role played by the implementation of quantum error-correcting codes in the realization of fault tolerant quantum computers and describes the latest advances in this field.

 Research into quantum error correction has progressed significantly over the last five years, with important theoretical advances and experimental demonstrations.

- Recent developments in bosonic codes and nonlocal LDPC codes have reduced the cost of error correction by at least an order of magnitude compared with the traditional surface code previously favored.
- Several major experimental projects were announced in 2023 and 2024. These proofs of concept demonstrated the experimental viability of error correction on quantum memories, through the realization of logical qubits that are more reliable than the physical qubits they rely on.
- Despite the progresses made in quantum error-correcting codes, the transition to applications that irrefutably demonstrate a quantum advantage requires the assembly of a number of logical qubits and the support of a number of logical operations several orders of magnitude greater than what is currently possible.

The report then reviews the most mature qubit technologies being developed in the world and in France. For each of these technologies, we describe the advances they have made, the challenges they face, their capacity to evolve and the resulting roadmap.

- There are five main groups of competing qubit technologies, with many variants, which have different strengths and weaknesses regarding manufacturing and scaling. The current state of knowledge does not allow any of them to be singled out as clearly more advantageous.
- The technological roadmaps of the major manufacturers of quantum computing technologies have, overall, been well respected to date.

The report then returns to the issue of scaling up existing solutions - a major global scientific and technical challenge if we are to reap the practical benefits of quantum computing in a wide range of applications - and examines the ability to ensure this scaling up in a modular way through the appropriate interconnection of various components.

- Monolithic integration of qubits into a single physical system has its limits, notably because of the noise and errors generated by increasing the number of qubits on a single chip or device and the connections required to address each qubit.
- Scaling up will often require a modular approach and interconnectivity via quantum communication links, with microwave and/or optical photons.
- The fidelity of distributed computational operations across multiple processors is still largely unexplored.

Finally, the report examines the question of comparisons, with other largescale technological programs developed over the last few decades, between FTQC and more conventional technologies (AI, HPC), as well as benchmarking quantum computers and their applications. It is difficult to assess the economic impact of FTQC a priori, and comparisons with other major programs, such as space research, can be enlightening.

- To compare advances and assess progress objectively and reliably, it will be essential to develop FTQC-specific benchmarks common to all platforms. This is made difficult by their diversity, their specific physical characteristics and applications, their low level of maturity, and the rapid evolution of technologies.
- Evaluation of FTQC's practical performance should be based on benchmarks that are close to real-life applications and that make sense for industry end users.
- Quantum computing is not intended to replace classical computing, but to complement it. quantum advantages are to be sought in computationally intensive rather than data-intensive applications.

The system architecture of a FTQC processor is a subject that has received little attention in France, and would benefit from the development of specific expertise and projects. These projects should include an understanding of the conditions for coupling high-performance computing (HPC) and FTQC, and the fit between technologies and applications, as well as the question of a possible energy advantage for FTQC, particularly for applications where FTQC and HPC could be in competition. In this respect, sharing and jointly analyzing feedback from the operation of early quantum processor demonstrators will be useful.

After this interim report, three areas need to be explored in greater depth to complete the study. From a technological point of view, the state of advancement of algorithms and applications, as well as compilation environments, and the question of enabling technologies, remain to be examined. Analysis of the economic outlook will involve studying the operating costs of the technologies described here, and the potential for value creation associated with the use cases. Finally, it will be necessary to assess the sector's training and skill requirements. In addition, it will be important to place the results of this study in a European and international context.

The report concludes that further scientific and technical progress is still needed to establish realistic industrial strategies. It recommends continued collective reflection to complete and regularly update the information gathered, in order to ensure a complete and up-to-date understanding of FTQC-related issues.

# Résumé exécutif

Ce rapport présente les résultats d'une étude collective menée par un Groupe de Travail créé en juin 2023 par l'Académie des technologies pour évaluer la faisabilité et l'impact du calcul quantique tolérant aux fautes (*Fault Tolerant Quantum Computing - FTQC*). Il fait suite à une étude réalisée pour le Secrétariat général pour l'Investissement (SGPI).

Le rapport revient d'abord sur les notions d'algorithme quantique, d'avantage quantique calculatoire, et décrit comment cet avantage varie selon les algorithmes et les cas d'usage. Il met en exergue les enjeux et défis concernant les applications du calcul quantique tolérant aux fautes.

- Il existe des arguments théoriques solides montrant qu'un processeur quantique tolérant aux fautes (FTQC) résoudrait certaines classes de problèmes drastiquement plus rapidement qu'un ordinateur classique, tels que la factorisation des grands nombres et la simulation de systèmes quantiques.
- Pour d'autres classes de problèmes, il existe un gain de temps théorique par rapport aux meilleurs algorithmes classiques, mais il est moins drastique. L'utilité en pratique dépendra alors des détails d'implémentation (notamment le temps d'exécution des portes quantiques).
- Les premières estimations de ressources pour des algorithmes dédiés aux calculateurs FTQC à même de résoudre des problèmes de portée industrielle sont de l'ordre de plusieurs milliers de qubits logiques, avec souvent un très long temps de calcul. Cela illustre le besoin d'optimiser la vitesse d'exécution des circuits quantiques et de la correction d'erreurs.
- Les cas d'usage industriels du FTQC sont, pour l'instant, peu ou mal identifiés en France, probablement car les industriels se sont concentrés sur

l'exploration des applications des petits processeurs quantiques bruités disponibles actuellement ou à court terme (dits NISQ) ou d'ordinateurs quantiques analogiques.

Le rapport souligne ensuite le rôle fondamental que joue la mise en œuvre de codes correcteurs d'erreurs quantiques pour réaliser des ordinateurs quantiques universels et exacts (c'est-à-dire tolérants aux fautes), et décrit les dernières avancées de ce domaine.

- La recherche en correction d'erreurs quantiques a beaucoup progressé au cours des cinq dernières années, avec des avancées théoriques et des démonstrations expérimentales significatives.
- Les développements récents concernant les codes bosoniques et les codes dits LDPC non locaux permettent de réduire le surcoût de la correction d'erreurs d'au moins un ordre de grandeur par rapport au traditionnel code de surface qui était privilégié auparavant.
- Plusieurs réalisations expérimentales importantes ont été annoncées en 2023 et 2024. Il s'agit de preuves de concept qui démontrent la faisabilité expérimentale de la correction d'erreur, par la réalisation de qubits logiques plus fiables que les qubits physiques qui les composent.
- Malgré les progrès réalisés en matière de codes correcteurs d'erreur quantique, le passage à des applications qui mettent en évidence un avantage quantique de manière irréfutable nécessite d'assembler un nombre de qubits logiques de plusieurs ordres de grandeur supérieur à ce qui est possible actuellement.

Puis le rapport dresse l'état des lieux des technologies de qubits les plus matures et qui sont notamment développées en France. Pour chacune de ces technologies sont alors précisées les avancées dont elles ont fait preuve ainsi que les défis auxquels leur maîtrise est confrontée, leur capacité d'évolution et la feuille de route qui en résulte.

 Il existe au moins cinq technologies de qubits concurrentes, avec de nombreuses variantes, qui ont des points forts et des points faibles différents concernant la fabrication et le passage à l'échelle.

- L'état actuel des connaissances ne permet pas d'en distinguer une comme étant clairement plus avantageuse.
- Les feuilles de route technologiques des grands fabricants de technologies de calcul quantique ont dans l'ensemble été bien respectées jusqu'à présent.

Le rapport revient alors sur le passage à l'échelle des solutions existantes, enjeu scientifique et technique mondial majeur pour pouvoir bénéficier concrètement des avantages quantiques calculatoires dans de nombreuses applications, et étudie la capacité à assurer ce passage à l'échelle de façon modulaire par l'interconnexion appropriée de divers composants.

- L'intégration monolithique de qubits dans un seul système physique a des limites, notamment en raison du bruit et des erreurs engendrées par l'augmentation du nombre de qubits sur une même puce et des connexions requises pour adresser chaque qubit.
- Le passage à l'échelle nécessitera une approche modulaire et de l'interconnectivité par des liens de communication quantique, avec des photons micro-ondes et/ou optiques.
- La fidélité des opérations de calcul réparties sur plusieurs processeurs est encore peu étudiée.

Enfin, le rapport examine la question des comparaisons, avec d'autres programmes technologiques de grande envergure développés au cours des dernières décennies, entre le FTQC et des technologies plus classiques (IA, HPC) ainsi que le *benchmarking* d'ordinateurs quantiques et de leurs applications. Il est difficile d'évaluer l'impact économique du FTQC *a priori*, et des comparaisons avec d'autres grands programmes, comme la recherche spatiale, peuvent être éclairantes.

 Pour comparer les avancées et évaluer les progrès avec objectivité et fiabilité, il sera indispensable de développer des *benchmarks* spécifiques pour le FTQC communs à toutes les plateformes. Cela est rendu difficile par leur diversité, leurs spécificités en matière de caractéristiques physiques et d'applications, leur maturité basse, ainsi que par l'évolution rapide des technologies.

- L'évaluation des performances pratiques du FTQC doit être envisagée à travers des *benchmarks* proches d'applications réelles et qui ont du sens pour les utilisateurs finaux industriels.
- Le calcul quantique n'est pas voué à remplacer le calcul classique, mais à le compléter ; les avantages quantiques sont à rechercher du côté des applications intensives en calcul plutôt que des applications intensives en données; dans le futur les technologies FTQC auront à se positionner par rapport à celles fondées sur du Silicium 3D.
- L'architecture système d'un processeur FTQC est un sujet peu abordé en France et mériterait le développement d'expertise et de projets spécifiques. Ces projets devraient inclure la compréhension des conditions de couplage entre le calcul haute performance (HPC) et le FTQC et les adéquations entre technologies et applications, ainsi que la question d'un éventuel avantage énergétique du FTQC, notamment pour les applications ou FTQC et HPC pourraient être en concurrence. À cet égard, partager et analyser conjointement les retours d'expérience d'exploitation de démonstrateurs de processeurs quantiques précoces sera utile.

Après ce rapport d'étape, trois axes devront être explorés plus en profondeur pour compléter l'étude. Du point de vue technologique, l'état d'avancement des algorithmes et des applications, ainsi que des environnements de compilation, et la question des technologies habilitantes, restent à examiner. L'analyse des perspectives économiques nécessitera d'étudier les coûts de fonctionnement des technologies décrites ici et le potentiel de création de valeur associés aux cas d'usage. Enfin, il conviendra d'évaluer la formation et les besoins en compétences du secteur. En outre, il sera important de placer les résultats de cette étude dans le contexte européen et international.

Le rapport conclut que des progrès scientifiques et techniques sont encore nécessaires pour fonder des stratégies industrielles réalistes. Il recommande de poursuivre la réflexion collective pour compléter et mettre à jour régulièrement les informations collectées, afin de garantir une compréhension complète et actualisée des enjeux liés au FTQC.

# Introduction

The American physicist Richard Feynman is credited with the idea, put forward in 1981, of creating quantum computers to study quantum phenomena by simulating them. Subsequently, numerous researchers demonstrated that if we had computers whose calculation principles were based on certain properties characteristic of quantum mechanics, then it was possible to design algorithms that showed significant accelerations, in particular exponential accelerations, compared with their equivalents (Bernstein-Vazirani, Deutsch-Jozsa, Shor...). This has led to many organisations, be they public or private, large or small, to embark on the actual development of a quantum computer in most industrialized countries.

The aim of this report is to inventory these developments. It is based on the international expertise of the members of the working group set up by the Académie des technologies and on a detailed analysis of various research programs and experiments carried out in France and presented by their leaders. However, the use of the principles of quantum mechanics in the field of information technology is not limited to accelerating computing, but can also involve the development of sensors providing new types of data, and the secure transmission of information within telecommunication networks. These other areas are not covered here.

Two main routes have been followed so far to build a quantum computer. An "analogue" approach, which proposes physical devices that are not very or not at all programmable, adapted to the simulation of quantum phenomena or even to implement certain optimization heuristics. A "discrete" or "digital" approach, proposes building universal quantum computers that can be programmed using a sequence of operations similar to the way in which a conventional computer is used.

In this case, however, the unit of information processed is no longer the "bit", but the quantum bit known as the "qubit". The properties of superposition of states and entanglement that characterize quantum mechanics make it possible to perform operations on a register of several gubits that have no classical equivalent, notably by exhibiting various forms of parallelism. This is what contributes to the potential speedup of a quantum computer and provides the elements for what is known as the computational quantum advantage. However, obtaining a quantum advantage to solve a given problem is difficult, and this is the subject of research in quantum algorithms. Furthermore, the execution of a quantum algorithm encounters a formidable obstacle that none of the world's current experiments has yet managed to overcome in its entirety. These are the errors that accumulate during any operation on the qubits: their initialization, the quantum gates that act on them, and finally gubits readout. One of the effects of these errors is to generate the decoherence that progressively damage entanglement and superposition, which distorts both the operations on the gubits and their readout. There are many sources of error. In particular, they arise from interactions between gubits and their external environment, as well as from imperfections in the tools used to control gubits using electronic or laser signals.

Trying to work around quantum noise and still benefiting some quantum advantage despite this noise, a first generation of solutions has emerged. Called the NISQ generation (*Noisy Intermediate Scale Quantum*), it applies to most of the developments currently underway around the world. It consists of exploiting a relatively small number of qubits, between 50 and 150, using shallow algorithms with a constrained number of operations that tolerate a limited amount of noise. Applications include computing the ground state of many-body quantum systems for chemistry (*Variational Quantum Eigensolver*) and solving combinatorial optimization problems (*Quantum Approximate Optimization Algorithm*). However, the small number of qubits and gate operations considered considerably limits the potential concrete applications. The NISQ generation now exploits the so-called error mitigation technique, which uses quantum computation results post-processing can reduce the errors generated using statistical methods.

Recently, a new generation has emerged. Although theorized a long time ago, it has remained experimentally out of reach for a long time.

This generation aims to actively correct quantum noise during computation. It is FTQC generation (*Fault Tolerant Quantum Computing*), which is the topic of this report. In an FTQC computer, basic physical qubits are assembled to form logical qubits, which are fewer in number than the physical qubits and have lower error rates. The redundant physical qubits are exploited with error correction codes to detect the impact of noise and correct it as it occurs during the circuit execution. Recent error-correcting code techniques were created, which require fewer physical qubits. They have opened up new perspectives. They make it more realistic to build reliable quantum computers that can be used for a wide range of applications.

The report is structured as follows.

The first chapter introduces the basic concepts of quantum computing and the practical reality of a quantum advantage. This quantum advantage is not monolithic. It depends on the nature of the algorithms to be implemented and therefore has a different impact depending on the field of use. Furthermore, when this advantage exists in theory, its practical exploitation depends on the quantity and quality of the logical qubits available to the computer that is supposed to execute the algorithm in question. This chapter concludes by comparing the state of the art in terms of the physical realization of logical qubits with several applications algorithms requirements. This comparison makes it possible to estimate how far scale-up is required.

The use of error- correcting codes is currently the only way devised to enable fault-tolerant quantum computers. This is covered in the second chapter. Until recently, the main error-correcting code studied were variants of the "surface code". Depending on the various physical implementations of qubits, the redundancy required for this type of code is about 100 to 10,000 physical qubits to create logical qubits in the context, for example, of Peter Shor's integer factoring algorithm. However, recent discoveries have made it possible to build qubits that are either intrinsically protected

against certain types of error (*bit-flip*), or to improve the local connectivity of qubits, which until now was limited to two dimensions in the nearest-neighbor fashion. This leads to significant reductions in the number of physical qubits per logical qubit. These approaches have opened up a vast field for exploring new solutions that require less hardware and are very helpful for the practical implementation of fault-tolerant quantum computers.

There are many physical qubit technologies being studied and tested around the world to implement physical qubits and the operations associated with them. In the third chapter, the technologies with the highest level of technology on the market are reviewed: neutral atoms, superconducting qubits, photons, spins in silicon and trapped ions. For each of these technologies, the advances they have made are described, along with the challenges they face, their capacity to evolve and the resulting roadmap.

Scaling up the number of physical qubits with which a quantum computer can be built is therefore the first major challenge that needs to be overcome. After noting the limits of monolithic qubit integration, the fourth chapter focuses on increasing the number of available physical qubits by interconnecting several quantum processing modules, and tackles the various problems associated with this interconnection. In particular, photonic interconnection and the coupling of quantum processors are explored. The scaling and connectivity of superconducting qubits is also analyzed in greater depth.

The fifth chapter outlines the first elements of a technical and financial analysis for the creation of a quantum computing industry. This investigation will be extended and completed in a second report by the Academy. The three points addressed in this chapter are as follows:

- The scale of the hopes raised by the existence of quantum computers and the scale of the public and private funding devoted to them around the world mean that research into quantum computers can be compared with other global programs such as nuclear fusion, space research, the sequencing of the human genome, and CERN's LHC (*Large Hadron Collider*) particle accelerator. From this point of view, FTQC research is closer to space research from a geostrategic point of view.

- In terms of performance, FTQC technology is competing in the short and medium term with other computing technologies which may reduce the size of its potentiall market over this period. 3D silicon circuit technology and the specialization of microprocessors dedicated to matrix computing (GPUs, *Graphics Processing Units, TPUs, NPUs*) mean that quantum processors are likely to coexist with traditional processors in supercomputer architectures, rather than replacing them. In the field of artificial intelligence (AI) and machine learning, in addition to the use of GPUs, the performance of quantum applications depends on the speed of of inputs/outputs operations with very large data sets. In this case, quantum acceleration might be seriously handicaped.
- Being able to compare architectures on the basis of their performance requires appropriate *benchmarking* methods. We need skilled personnel, and we need to be able to rely on essential enabling industries (cryogenics, etc.), lasers, semi-conductors, quantum software...) and to strengthen training courses are just a few of the questions that have been sketched out here, and which will be added to in the next section.

### Chapter 1

# THE REALITY OF QUANTUM ADVANTAGE AND ITS NEEDS

#### 1.1. QUANTUM ALGORITHM

A quantum algorithm is an algorithm in the usual sense. It operates on a quantum memory. It applies a series of quantum gate operations on this quantum memory, measure part of it, or alternate between the two. This sequence is described by the quantum algorithm. The results of the measurements can be used to control the sequence of operations to be performed. The description of an algorithm is therefore completely deterministic. Quantum programming tools have been available for a long time, and some of them are very advanced, with a rich set of programming languages, development environments and libraries.

Using this quantum memory, a quantum algorithm can, for example, simulate the evolution of a quantum system by decomposing the evolution into a sequence of elementary operations that can then be performed on its memory. This is a form of emulation, in that the quantum memory is then in a physical state that approximates the state of the quantum system under study.

The main difference with a classical memory, i.e. one governed by classical physics, lies in the existence of phenomena, or paradoxes, in quantum physics that cannot be explained or realized by classical physics. It is, of course, possible to describe them, and therefore to simulate them numerically on a classical computer, but, as we shall see, the computing power of a quantum computer capable of compiling and executing a quantum algorithm is theoretically infinite.

The computing power of these computers is far greater than that of our conventional computers, or that of the computers of the future, because it is potentially exponentially greater. This computational superiority only applies to certain specific tasks, which the scientific community has been identifying and mapping for over thirty years.

A quantum bit or quantum system has two levels which are represented in the form  $|\psi\rangle = \alpha |0\rangle + \beta |1\rangle$ , where  $\alpha$  and  $\beta$  are two complex numbers, called amplitudes.  $|0\rangle$  and  $|1\rangle$  correspond to the basic states of the qubit, as for a classical bit, but a qubit can be found in any quantum superposition of these states. Note that this is not a statistical mixture of the two basic states: a statistical mixture describes the situation where there is uncertainty about the state under consideration, whereas a superposition describes a coherent combination of the  $|0\rangle$  and  $|1\rangle$  states, and enables interference<sup>2</sup> to be generated.

The probabilistic correlation of several observations, and therefore of several states, is a well-known phenomenon in classical physics. The same is true in quantum physics, where it is referred to as entanglement. It is possible to entangle two quantum objects, such as two atoms or two electron spins, so that the result of observing one (by physical measurement) is correlated with that of the other. In quantum physics, on the other hand, experience shows that the result of a measurement of one of the two quantum objects, themselves in a superposition state, is not fixed until it is made, rather like a die that keeps spinning until it is looked at. It is then that a counter-intuitive phenomenon occurs, because this random result is correlated to that obtained by the measurement of the other bound quantum object, even if it is so far away that no information has time to be transmitted between the two objects. This may seem paradoxical, but it is the correlation between the two particles that is preserved, rather than the transmission of information from one to the other. Quantum entanglement also, and above all, occurs during computing, and before any measurement, thanks in particular to the two-gubit guantum gates that generate these correlations between gubits and that can be observed later by measurement. It makes it possible to generate

<sup>2</sup> In fact, it is also possible to consider statistical mixtures of superpositions, called mixed states and represented by density matrices.

interferences between several qubits and then amplify certain amplitudes, for example that of the quantum state corresponding to the result of the algorithm, so that it is revealed during the final measurement. These interferences are sometimes referred to as quantum parallelism.

This kind of phenomenon has now been verified experimentally over great distances, including by satellite. It is now accepted that quantum physics cannot be efficiently simulated, in all generality, by a classical computer, even using probabilistic techniques. On the other hand, it is quite an art to identify that fringe of problems where quantum computers could provide an incomparable computational advantage.

#### 1.2. UNIVERSAL COMPUTER

Our current non-quantum computers are universal, i.e. non-specialized, and can execute on demand anything that can be calculated by another computer, leaving aside the constraints of time and memory. This universal character is in fact linked to the mathematical concept of the Turing machine. On the one hand, any machine can be described by a "code" or program, and on the other, there is a "universal" machine that can simulate any other machine by reading its code. In quantum computing, there is also a quantum variant of this universality. We can define a Turing machine that formally models quantum computation. Some of these quantum machines are also universal, so any other quantum machine can be programmed into them. This point is less intuitive than in classical computing, because here it is physical transformations - operations on quantum states - that are configured according to the task at hand. However, these operations are also universal, because they can be described by unitary matrices with real or complex coefficients, rather than by logic gates based on 0s and 1s.

Quantum programming tools have already been developed. They are ready for the machine that will be able to control a sufficiently large quantum memory. These programming languages usually describe quantum circuits, which are then realized on platforms that manipulate quantum particles. As with logic, just a few types of gate are enough to create all possible circuits: gates acting on one qubit, plus a gate of a unique type acting on two qubits (the *control-not*, for example). This is also a notion of universality. A set of quantum gates is universal if it can perform any unitary transformation by combination. There are also finite sets, at the cost of an approximate universality, but sufficient for calculation. For example, a universal set of gates can be made up of the Hadamard and T gates acting on a single qubit plus the *control-not* gate acting on two qubits.

Since a quantum computer is a more complex system than a classical computer, we might think that it will be reserved for certain tasks, and that there is therefore no point in trying to run all the classical algorithms on quantum machines. The fact remains, however, that classical logic routines, such as arithmetic, will also have to be possible on quantum states. Since transformations on quantum states have to be reversible, one might think that this would pose a problem when it comes to performing a non-reversible logical operation, such as adding two bits or qubits<sup>3</sup>. Fortunately, it was known even before the start of quantum computing that any classical calculation can be made reversible with an additional cost in memory and calculation time. What's more, it turns out that a reversible classical calculation is a special case of a unitary transformation, so in principle it can be converted directly into a quantum circuit at no extra cost. Of course, on small computers of a few thousand bits or more, the (controlled) extra cost of switching to reversible (and then quantum) computing will undoubtedly be an obstacle at the beginning.

#### 1.3. FAULT-TOLERANT COMPUTING

On the other hand, unlike today's computers, the error rate of quantum processors, or even just quantum memory, remains high, and so cannot be ignored. This paradigm was studied for a time when the first (classical) computers were being developed. This is the problem of fault-tolerant computing, for which solutions exist in

<sup>3</sup> This operation is non-reversible, since it is impossible to recover the state of two bits a and b from their sum a+b.

classical computing. These have been exported to quantum computing and validated theoretically. Initial experimental demonstrations have also been carried out (see chapter 2).

The errors that occur during a quantum circuit execution are different from those that occur during a classical calculation. In a classical bit, there may be a substitution of a 0 by a 1, or the opposite. In a quantum bit, there can be a whole continuum of modifications, because the state of a quantum bit is represented by two complex numbers (its coordinates) which can undergo a non-discrete transformation, known as a continuous transformation. These continuous errors come from different sources of noise that affect the physical objects that support the qubits. They result from imperfections in the physical systems that act on the qubits to generate quantum gates, as well as from the environment.

Surprisingly, however, it is possible to correct these continuous errors using quantum error correction codes. The great discovery was to realize that it is sufficient to take into account two types of error (bit inversion and phase inversion) to correct all the errors, thanks to the linearity of quantum transformations. Inspired by classical corrector codes, but mathematically much more difficult to construct, quantum error correction codes protect the quantum memory against sources of noise. They can also protect the circuit execution, when the error rate is low enough. Of course, this comes at an additional cost in terms of the additional number of qubits required and computing time. From a theoretical point of view, this extra cost is considered reasonable. From a practical point of view, where, in 2025, the thousand physical qubits will barely have been reached, it's a different story.

Note that (classical) corrector codes are used every day in telecommunications, for example. There is therefore a great deal of know-how that has led, among other things, to the discovery of quantum corrector codes that are very well suited to quantum computing. Let's not minimize the theoretical power of these quantum corrector codes, which, by detecting a discrete error, can correct an *a priori* continuous error. In the 1990s, the very existence of such codes was called into question.

These error correction codes lead to two levels of measurement of the number of qubits used: logical qubits and physical qubits. The physical qubits correspond to the quantum systems actually produced on the processor. They are noisy by construction, but can be used to create a logical qubit with less "logical" noise. A logical qubit then becomes an abstraction used by the algorithm, which requires a sufficiently low error that it does not disturb the response of the algorithm (or only to a negligible extent).

The number of physical qubits required to produce a logical qubit depends on the physical error rate, hardware constraints (e.g. connectivity) which influence the choice of a suitable quantum error correcting code, and the "logical" error rate depending on the intended use. The link between these two notions of qubits is developed in Chapter 2.

#### 1.4. THEORETICAL COMPLEXITY

It is not uncommon to see, even today without a quantum computer, fast quantum algorithms being discovered, followed a few months later by classical algorithms that are almost as good. There is in fact a real emulation between communities, a bit like in cryptography when an encryption, signature or other application is developed, standardized, then attacked, before being upgraded cyclically until the next attack is unveiled. It is virtually impossible to prove that a problem is difficult. Complexity theory allows us to distinguish between the two, and therefore to better understand the potential of quantum computing.

This theory classifies problems according to their proven or conjectured difficulty. The most difficult are the undecidable problems. You might as well move on here, as they are mathematically proven to be unsolvable, even with all the time in the world. Computational problems that are not necessarily easily solvable (i.e. in polynomial time), but whose solution is easily verifiable, are known as NP problems. The class of NP-complete problems includes all problems that are as difficult to solve as the other NP problems. If we could find an efficient algorithm for an NP-complete problem, we could efficiently solve all NP problems. Today, it is conjectured that the

Not all NP problems can be solved efficiently (it relates to the famous "P=NP?" question), and solving a NP-complete problem is a challenge that currently takes exponentially longer depending on the size of the problem. Complexity theory, a branch of computer science, has developed a solid theory of their presupposed difficulty. This theory is fairly stable, and there is little chance that these problems will ever admit of a rapid algorithmic solution. On the other hand, for certain instances, used every day in the industrial world, there are heuristics that provide acceptable solutions. This is one of the most strategic areas of research in operations research.

There are other intermediate problems, for which it is difficult to construct a convincing theory of their difficulties. Yet no efficient algorithm is known. One example is the integer factorization problem. This involves expressing a (non-prime) number as a product of prime numbers. Verifying a solution is easy, as it involves simple multiplication. This problem is therefore NP, but there are arguments to suggest that it cannot be NP-complete. Factorization and its variants form the basis of contemporary public key cryptography, which is standardized and used in most secure transactions and operations. However, this is one of the problems that a quantum computer could solve quickly (i.e. in polynomial time linked to the number of digits in the number to be factorized).

We have just seen that there are undecidable, NP-complete, NP, but not complete problems. Efficiently solvable problems are said to be in P, for solution in polynomial time (linear, quadratic...).

We can also consider the BPP class (Bounded-error Probabilistic

*Polynomial*), problems that can be solved by a probabilistic algorithm in polynomial time with (negligible) probabilistic error. The quantum equivalent of this class is BQP. The latter is therefore *a priori* broader than BPP and P (*polynomial*). But, alas, it is not reasonable to think that BQP contains an NP-complete problem, at least not any more than BPP or P would. However, theoretical arguments suggest that there would be a whole series of difficult problems (for classical computers) in BQP, but that none of them would be complete. So, you have to look for them one by one! For example, factoring is an *a priori* non-complete NP problem that is not *a priori* in BPP, but in BQP, because the factorization algorithm

is performed in polynomial time. The advantage of having such an NP problem is that the solution is easily verifiable. In addition, some problems, such as boson sampling, used in photonics, are in BQP, but not in NP, because they are not easily verifiable.

Note here that a quantum computer cannot compute what is undecidable. It has also been proved that any problem solved by a quantum computer (for example in BQP) can also be solved by a classical computer, with a polynomial memory, but with a potentially exponential time (PSPACE class). From a computing point of view, the advantage to be sought is therefore in computation time (alongside other potential advantages not linked to algorithmic complexity, such as energy cost or data confidentiality), and not in computability. In other words, quantum computing does not call into question the Church-Turing thesis, but it does call into question its extended version.

#### 1.5. QUANTUM ACCELERATION

Which quantum computer for which use? This question is being asked today in an environment where the first computers were of modest size compared with the scale of contemporary computing. We are talking about a few hundred bits in a quantum state, i.e. physical qubits, compared with several gigabytes on our smartphones. Worse still, these qubits are currently noisy, i.e. with too high an error rate to carry out a calculation. Yet a reliable quantum memory containing just a few thousand quantum bits would be enough to provide quantum computing power that could not be matched by the most powerful conventional computers.

There is an *a priori* rich source of problems for which a quantum computer would provide an exponential advantage in computing time. Factoring is one such problem, as is the simulation (or emulation) of a quantum system (quantum chemistry, for example). It is also possible to deal with linear algebra problems, such as solving linear systems, which are intractable when the number of variables is too large. These problems are universal for quantum computation.

There is also a whole series of algorithms whose gain is not exponential, but is still significant. A quadratic acceleration is even almost

on a whole series of quantum optimization algorithms, particularly in operations research.

A final category also emerged, between exponential and quadratic speedups, around machine learning and solving partial differential equations, where ultimately large linear systems have to be inverted. For a while, an exponential gain was hoped for, but in the end, it was only polynomial.

There are three types of quantum acceleration:

- exponential acceleration, for example Shor's algorithm for factoring large numbers or simulating quantum systems;
- polynomial accelerations, including algorithms using HHL (*Harrov-Hassidim-Lloyd*) for e.g. approximate matrix inversion, as well as the recommendation algorithm it should be noted that, generally speaking, even a polynomial acceleration can bring significant gains in the industry;
- quadratic acceleration (a special case of the previous one), such as Grover's algorithm for searching an unordered database, and amplitude estimation (useful for applications such as Monte-Carlo simulation) although a quadratic factor is important in industry, this acceleration will probably be less obvious in practice, since, given the speed of a quantum processor, its asymptotic advantage will only become apparent at disproportionately long computing times, currently estimated in centuries

These accelerations have all been mathematically proven, as it is not possible to achieve these algorithms today. Although they are very interesting from a research point of view, it is important to consider the study of practical cases rather than theoretical worst-case scenarios. Factors such as the clock cycle and error correction overheads make the actual acceleration more uncertain, particularly for quadratic and polynomial accelerations.

Quantum acceleration on a real input compared with existing classical algorithms is therefore not guaranteed in practice. It is also not clear when this asymptotic acceleration will manifest itself, given all the internal machinery and its additional costs, which depend on the technologies used and their progress. Efforts are mainly focused on reducing the costs associated with error correction, a real barrier that seems to be gradually lowering, and on transposing quantum algorithms onto the chosen architecture, such as decomposing operations into elementary quantum gates depending on the connectivity of the quantum bits used. These steps must also be carried out for data loading and for the classical algorithms used to process classical tasks in superposition.

Some research groups are working on estimating the resources that would be needed to solve a useful problem. Here are a few journal articles that review existing algorithms and use cases, estimating the resources required, whether in terms of quantum logic qubits or physical qubits, taking error correction into account:

- "Assessing requirements to scale to practical quantum advantage4".
- "Quantum algorithms: A survey of applications and end-to-end complexities<sup>5</sup>".

Resource estimates are regularly carried out, attempting to take all factors into account. Tools are even developed for this purpose for written programmes such as :

 "Using Azure Quantum Resource Estimator for Assessing Performance of Fault Tolerant Quantum Computation<sup>6</sup>".

<sup>4</sup> M. E. Beverland, P. Murali, M. Troyer et al. https://arxiv.org/abs/2211.07629 (2022)

<sup>5</sup> A. Dalzell, S. McArdle, M. Berta et al. https://arxiv.org/abs/2310.03011 (2023)

<sup>6</sup> W. van Dam, M. Mykhailova, M. Soeken, https://arxiv.org/abs/2311.05801 (2023)

#### 1.6. HEURISTICS AND EMPIRICAL APPROACHES

Until recently, the development of quantum algorithms has been based on a theoretical approach, since none of the quantum algorithms with established superiority has yet been implemented on full-scale data on a quantum computer. Quite rightly, a heuristic, even empirical, approach is now being widely deployed in order to take advantage of the first existing or future computers. Algorithms such as *Variational Quantum Eigensolver* (VQE) and *Quantum Approximate Optimization Algorithm* (QAOA) implement such heuristics, however, with no guarantee of acceleration.

This approach is traditional in computer science, and particularly in industry, for dealing with infeasible problems such as NP-complete problems (SAT, MAXCUT, etc.), for which there are basically solutions suitable in practice. The middle way is to look for accelerations for certain practical cases, even if theory tells us that in the worst case no acceleration is expected.

These new heuristic approaches, combining theory, experimentation (on the NISQ platform or simulated on HPC) and conjecture, have therefore been actively developed. Inspired by their theoretical analogues, they greatly simplify the algorithms previously proposed by taking into account hypotheses about the data, and sometimes also by dispensing with theoretical validation. We can see here a certain approach to deep learning that has proved its worth, even if quantum learning is still a long way off, and we will no doubt have to start by going beyond NISQ.

However, to date, no advantages have been demonstrated. Worse still, the excitement of the first results is often later erased by (classical) algorithmic advances that match the calculations performed on the first quantum circuits.

Basically, this is just another incentive to create a calculator that is large enough, fault-tolerant and as universal as possible.

#### 1.7. Use cases

The precise identification of use cases for fault-tolerant quantum computing is difficult, probably partly because it is hard to demonstrate usefulness before achieving feasibility. Assessing the economic impact of solving these use cases is an even more complex problem, and it is very difficult to provide a precise answer today.

Some industry sectors are more involved than others in research into faulttolerant quantum computing applications, and these sectors vary from country to country. In France, the energy, financial services and defense sectors seem to be well positioned. There is also the field of computational chemistry for drug discovery.

However, there is some consensus that fault-tolerant quantum computing will be more beneficial for *compute-intensive* applications than for *data-intensive* applications. Data-intensive applications, such as machine learning, spend considerable execution time manipulating inputs and outputs, whereas quantum acceleration is more about the execution of the computation itself, and loading data into a quantum computer is a physically slow process.

The primary motivation for the creation of quantum computers is to enable the resolution of problems whose resolution time on classical computers is *a priori* exponential as a function of the number of parameters. In practice, this quest has broadened somewhat, as a number of quantum algorithms have been overtaken by their classical counterparts, in particular because the latter have been able to undergo developments stimulated by the potential quantum advantage.

So the notion of a quantum advantage now incorporates a number of different parameters: an acceleration in computing time, which is generally polynomial or exponential, better quality results, the ability to generate good results with less training data in the case of machine learning and, *ultimately*, an economic advantage with a solution that costs less in absolute terms, particularly in terms of energy costs, which are very high in high-power computing centers.

#### 1.7.1. CLASSIFICATION BY TYPE OF PROBLEM

Taking a closer look at sections 1.5 and 1.6, we can classify the current problems that could potentially benefit from the contribution of quantum computing according to the following issues, bearing in mind that these benefits remain to be proven from a theoretical and practical point of view. The relevance of these problems depends on the sectors in which companies operate. The ability of fault-tolerant quantum computers to solve them is still open to question, for various reasons mentioned in this report. The variety of algorithms, use cases and hardware solutions envisaged is so great that no general conclusions can be drawn from current knowledge in the field.

Simulations of quantum physical systems, which concern solid state physics as well as quantum chemistry. The aim is generally to determine the rest energy configuration of these systems (*ground-state*) with greater precision than with conventional methods. By extension, other quantum algorithms aim to determine the spectrography of molecules, their vibrational or rotational structure, excited states, how they interact with each other and how various chemical reactions work. This would make it possible, for example, to create *ab initio* molecules. Applications include the creation of new solid materials, such as batteries that are more efficient in terms of energy density or charging speed, or new chemical compounds in the energy sector, such as for carbon capture, or the identification of new, more energy- efficient catalysts for the production of fertilizers or various therapeutic solutions<sup>7, 8, 9, 10</sup>. The existence of accelerations

<sup>7</sup> Quantum-centric Supercomputing for Materials Science: A Perspective on Challenges and Future Directions\_by Yuri Alexeev, Liang Jiang *et al*, DoE, CERN, RIKEN, University of Toronto, University of Maryland, Infleqtion, Dell, IBM, BMW, Boeing, Bosch, Algorithmiq, ExxonMobil, *et al*, *arXiv*, December 2023 (60 pages).

<sup>8</sup> Quantum Chemistry in the Age of Quantum Computing by Yudong Cao, Alan Aspuru-Guzik *et al*, 2018 (194 pages). © 2019 American Chemical Society

<sup>9</sup> Quantum chemistry, classical heuristics, and quantum advantage by Garnet Kin-Lic Chan, arXiv, July 2024 (45 pages).

<sup>10</sup> Quantum Chemistry on Quantum Computers by Libor Veis, 2021 (31 pages).

for these types of calculation is not generically guaranteed<sup>11</sup>. Furthermore, as we will see below, current estimates of resources in terms of the number of logical and physical qubits and computing time are within ranges that make these use cases highly uncertain given the current state of our knowledge and technological developments.

- Decision and optimization problems, which can be found in a wide range of professions. These include, for example, optimizing the routing and filling of vehicles for deliveries, or allocating resources in factories with a large number of machines, incoming and outgoing objects, and technical staff to allocate to these tools. There are also many optimization problems in financial services and insurance<sup>12</sup>. Known optimization algorithms tend to have polynomial rather than exponential accelerations. The scientific literature is less abundant on how to solve these problems with fault-tolerant computers.
- Machine learning for both training and model inference <sup>13,14</sup>. These problems are currently solved fairly effectively using conventional algorithms based on hardware architectures with specialized processors, which optimise massive matrix calculation. There are a large number of quantum machine learning algorithms that could achieve better results or at lower cost in a significant number of areas (generative networks, image recognition, language processing), or even speed them up. One of the current motivations for research in this field is the enormous energy costs involved.

<sup>11 &</sup>lt;u>Is there evidence for exponential quantum advantage in quantum chemistry?</u> by Seunghoon Lee, Ryan Babbush, John Preskill *et al*, *arXiv* August 2022 (81 pages).

<sup>12 &</sup>lt;u>A Review on Quantum Approximate Optimization Algorithm and its Variants</u> by Kostas Blekos *et al*, *Physics Reports* Volume 1068, 2 June 2024 (85 pages).

<sup>13</sup> Quantum Artificial Intelligence: A Brief Survey by Matthias Klusch, Jörg Lässig, Daniel Müssig, Antonio Macaluso, and Frank K. Wilhelm, DFKI, arXiv, August 2024 (21 pages).

<sup>14 &</sup>lt;u>A comprehensive review of Quantum Machine Learning: from NISQ to Fault Tolerance</u> by Yunfei Wang, and Junyu Liu, The University of Chicago and Brandeis University, *arXiv*, January-March 2024 (53 pages).

large language models. However, obtaining a quantum advantage depends on optimizing the loading of training data into these algorithms. The unsuitability of quantum computing for ingesting large volumes of data remains a definite handicap. Furthermore, a generic quantum advantage has not yet been demonstrated for quantum *machine learning* algorithms<sup>15</sup>.

- Engineering problems requiring, in particular, the solution of complex partial differential equations<sup>16</sup> with vast areas of application in the design of various systems in aerospace, engines and fluid mechanics<sup>17</sup>. Note that these use cases do not generate guaranteed exponential acceleration.
- Cryptanalysis, which includes methods for breaking public encryption or signature keys used on the Internet, such as those based on the RSA encryption and decryption algorithm. The best-known quantum algorithm is Peter Shor's integer factorization algorithm, but it is not the only one. The same Peter Shor's discrete logarithm algorithm and Grover and Simon's algorithms could be used to break various public key cryptographic systems, albeit requiring very long computation times. The existence of this future threat could, for example, affect data intercepted today that would have value in the distant future (*Store Now Decrypt Later*). This has led to the creation of countermeasures such as so-called post-quantum encryption algorithms, whose standardization by the NIST in the USA was finalized in the summer of 2024 for three out of four selected protocols, and whose deployment is beginning to be planned.

<sup>15</sup> Is quantum advantage the right goal for quantum machine learning? by Maria Schuld and Nathan Killoran, arXiv, March 2022-February 2023 (13 pages) and <u>Why measuring performance is our biggest blind spot in quantum machine learning</u> by Maria Schuld, Xanadu, *PennyLane blog* March 2022.

<sup>16</sup> Investigation on a quantum algorithm for linear differential equations by Xiaojing Dong, Yizhe Peng, Qili Tang, Yin Yang, and Yue Yu, arXiv, August 2024 (22 pages).

<sup>17 &</sup>lt;u>Quantum computing for simulation of fluid dynamics</u> by Claudio Sanavio, and Sauro Succi, *arXiv*, January 2024 (11 pages).

The practical realization of quantum computing solutions in all these areas is still hypothetical. It depends on numerous scientific and technological factors and uncertainties, which will be discussed in this report. These relate, of course, to the ability to scale up hardware and error correction, but also to practical considerations in terms of actual computing time and acceleration, on a case-by-case basis. The notion of a quantum advantage cannot yet be asserted across the board for these different use cases.

The viability of these solutions depends on many factors and unknowns: the creation of more efficient quantum algorithms, the evolution of methods for preparing and encoding data in quantum registers, improvements in the quality of qubits, scaling, the execution speed of quantum gates, the creation of even more efficient error correction codes, the ability to interconnect quantum processors efficiently and quantum-fashion, and the ability to parallelize quantum circuits on a large number of affordable machines to reduce computing times. The diversity of these challenges, of the solutions envisaged by researchers and companies in the sector, and of the scientific, technological and economic uncertainties associated with them, make any predictions hazardous. One thing is almost certain: given the development and experimentation cycles, even if we are optimistic, it will take a long time to solve all these problems.

#### 1.7.2. CLASSIFICATION BY INDUSTRY

At a very high level, as shown in Figure 1 below, there are two main categories of industrial applications:

 The first is applications that facilitate the work of researchers in academia and industry, particularly in the discovery of new materials, molecules and chemical processes. This most often involves creating digital twins that exploit the laws of quantum physics. In the case of applied research, products designed using quantum computing will then be manufactured and distributed using traditional processes; - the second is linked to the day-to-day operations of companies in many sectors, which aim to make them more efficient, by improving the services they provide or reducing their costs. There are many use cases for optimization in financial services, transport, logistics and distribution, as well as in all manufacturing industries. Telecommunications operators are also concerned with optimizing their networks (base station placement, frequency plans, traffic routing, etc.).

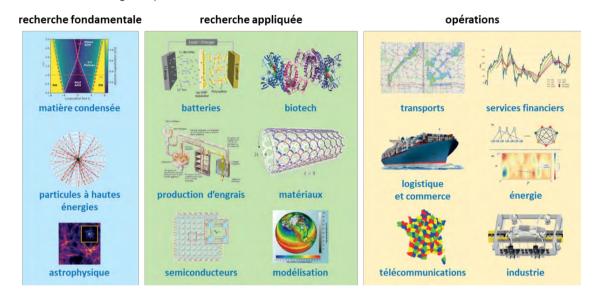


Figure 1: mapping of use cases, separating the fields of fundamental research, applied research and business operations. (cc) Olivier Ezratty, January 2025.

The working group interviewed a number of French industrial players, including start-ups, about their potential uses for quantum computing. Here is a summary of the use cases identified.

#### 1721 APPLICATIONS IN CHEMISTRY FOR DRUG DISCOVERY

Computational chemistry aims to model and predict the behavior of molecules in order to design medicines atom by atom. It is a field at the frontier between chemistry and quantum mechanics. It is crucial to reconsider the scales of time

and size when we try to predict the behavior of molecules: a change of nature occurs when we reach sufficiently small scales, during the transition from classical mechanics to quantum mechanics.

The main objective is to discover drugs by exploiting the increased precision offered by quantum computing. The algorithms used include VQE (*Variational Quantum Eigensolver*) and QPE (*Quantum Phase Estimation*). VQE, a hybrid algorithm, combines classical optimization techniques with quantum routines. It is a heuristic, but this is very often the case in theoretical chemistry, unlike theoretical physics. VQE is compatible with NISQ (*Noisy Intermediate-Scale Quantum*) and FTQC, while QPE represents an opportunity to go even further in terms of precision, but requires FTQC.

Simulating a water molecule with a quantum computer requires 400 to 600 logic qubits. With 1,000 logical qubits, we could solve interesting problems with this type of molecule. For larger molecules, which we would need to design vaccines for example, we would need even more logical qubits.

#### 1722 DEFENCE APPLICATIONS

The ambitions are many: to solve problems that cannot be solved by HPC and to solve solvable problems more quickly or at lower cost (energy and economic).

The use cases envisaged are the resolution of partial differential equations for applications in antenna and radar optimization, optimization in satellite control and trajectories, machine learning for detecting anomalies in images, Monte-Carlo simulations for locating and simulating materials.

We should also mention, orthogonally, cryptography (in particular the development of post-quantum cryptography, which requires the testing of quantum attacks).

#### 1723. ENERGY APPLICATIONS

Various players in the field of power generation are supplementing their expertise in high-performance computing with quantum computing, in the belief that in a few years' time they will be able to have appropriately sized quantum computers at their disposal. These teams are working on materials simulation, safety studies and partial differential equations. They are also interested in optimizing the charging of electric vehicles.

Until now, studies have focused almost exclusively on emulation, with very few tests on real machines. With 100 to 1,000 logic qubits, it would be possible to solve some interesting use cases.

Another area concerns the search for fossil fuels, and players are looking at the uses of quantum computing in quantum chemistry, combinatorial optimization and the resolution of partial differential equations.

Interesting use cases in decarbonizing energy production include simulation for  $CO_{(2)}$  capture and wind farm optimization. With as few as 300 logic qubits, it seems possible to solve some interesting problems.

#### 1.7.3. CLASSIFICATION BY RESOURCE

We conclude this section with a graphical representation of these use cases in terms of the resources required. Figure 2 lists case studies documented with precise resource estimates and in scientific papers. The case studies in the figure are fairly representative of what is currently expected as a function of the number of fault-tolerant qubits available and the number of quantum gates to be executed. Figure 3 lists a number of classes of algorithms with their logic resources: number of logic gates (generic, T-gates or Toffoli gates depending on the case, bearing in mind that T-gates and Toffoli gates are often used as benchmarks because they are very expensive to correct compared with other quantum gates such as X, Z, H and CNOT), fidelity of physical or logic qubits and sub-problems. In

terms of chemical simulation for the creation of new processes, such as the search for new catalysts for ammonia production or the creation of corrosion-resistant materials, several simulations are required with compounds and data of different complexities.

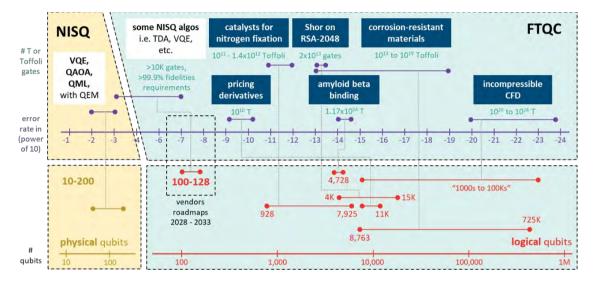


Figure 2: Mapping of use cases according to resource requirements. It illustrates the fact that the most value-generating applications in chemical and materials simulation, financial computing and cryptanalysis require a very large number of logic gates and logic qubits. The first viable applications of fault-tolerant computing could be those designed to Noisy quantum computers (NISQ), but require qubits of greater quantity and quality than are currently available. (cc) Olivier Ezratty, January 2025. These include<sup>18</sup>:

- simulation of quantum physical systems: VQE (in NISQ and FTQC), catalysts for nitrogen fixation (for more efficient production of ammonia via the Haber-Bosch process), corrosion-resistant materials, beta-amyloid bonding (in connection with the treatment of Alzheimer's disease);
- decision and optimization problems: QAOA (solving combinatorial problems in NISQ and FTQC), valuation of financial derivatives (in financial services);
- machine learning: QML (in NISQ and FTQC), TDA (FTQC), for topological data analysis, an automatic clustering method based on machine learning;
- engineering problems: incompressible CFD, fluid mechanics;
- cryptanalysis: Shor algorithm for RSA 2048.

There has been a sort of shift in usage scenarios. Many of the current NISQ case studies in practice require fault-tolerant quantum computers, including for variational algorithms whose circuits are too deep to be executed with current processors, including with the use of error mitigation techniques. It is therefore likely that the first use cases for fault-tolerant computers will be for NISQ

Sources: TDA: Towards Quantum Advantage on Noisy Quantum Computers by Ismail Yunus Akhalwaya et al, September 2022-March 2024, Feasibility of accelerating homogeneous catalyst discovery with fault-tolerant quantum computers by Nicole Bellonzi et al, June 2024, Real Option Pricing using Quantum Computers by Alberto Manzano et al. March 2023, How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits by Craig Gidney and Martin Ekerå, 2019, Quantum Resources Required for Binding Affinity Calculations of Amyloid beta by Matthew Otten et al, June 2024, Quantum computers by Nicole Bellonzi et al, June 2024, Quantum computing for corrosion-resistant materials and anti-corrosive coatings design by Nam Nguyen et al, June 2024, Feasibility of accelerating homogeneous catalyst discovery with fault-tolerant quantum computers by Nicole Bellonzi et al, June 2024, Feasibility of accelerating incompressible computational fluid dynamics simulations with fault-tolerant quantum computers by John Penuel et al, June 2024.

algorithms that have already been designed and need qubits with fidelities in excess of 99.9%. This will correspond to the first generations of fault-tolerant computers with a hundred or so logic qubits with error rates ranging from  $10^4$  to  $10^{-8}$ .

The use cases currently identified for fault-tolerant computers require physical resources that go beyond these characteristics, with thousands of logical qubits and more than 10<sup>12</sup> operations, the TeraQops mentioned later in the document.

The analyst firm GQI created a database of 174 case studies in August 2024<sup>19</sup>. It then used the Microsoft Azure *Resource Estimator* tool to normalize the resource estimates shown in the figure below. She shows that, for the most interesting chemical simulation cases (QPE, *Quantum Phase Estimation*), the calculation times are prohibitive, given that the points positioned between the middle and the right of the surfaces are the most credible in terms of physical qubit fidelity and quantum gate speed. These simulations could be carried out by running their circuit in parallel on several independent quantum processors. This will only be possible if the price and operating cost of each processor are reasonable. On the other hand, simulations of quantum dynamics seem to be achievable in reasonable times (left), but still requiring a very large number of physical qubits.

<sup>19</sup> The GQI Quantum Resource Estimator Playbook - Quantum Computing Report by Doug Finke, Quantum Computing Report, August 2024.

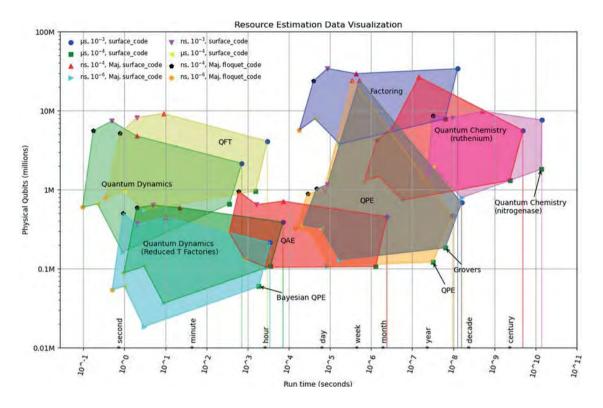


Figure 3: Classes of algorithms and their physical resources in number of qubits and computation time. ns = quantum gates realised in about 100 ns. μs = quantum gates realised in about 100 μs. Powers of 10 correspond to the error rates of two-qubit physical gates. Maj corresponds to qubits based on Microsoft Majorana fermions (which do not yet exist). The others correspond to superconducting qubits (for ns gates) or trapped ions (for μs gates).

## Chapter 2

## **ERROR-CORRECTING CODES**

## 2.1. Fault tolerance, the only route to deep quantum algorithms $^{\scriptscriptstyle 20}$

The acceleration of experimental progress in the control of quantum systems led John Preskill in 2018 to identify the era of the NISQ (*Noisy Intermediate-Scale Quantum*), i.e. that moment in history when quantum processors of a few hundred qubits become available and can tolerate thousands of operations (or gates) before an error occurs. While these numbers remain far too small to implement the best-known quantum algorithms, they have raised hopes of finding interesting applications for these processors, which are at the limit of what can be emulated with conventional supercomputers, for example by implementing optimization algorithms. It seems that a quantum computer will only become useful if it can perform a very large number of logical operations, at least 10<sup>10</sup>. The typical order of magnitude is that of TeraQuop, corresponding to 1,000 billion error-free logical operations. For example, Shor's factorisation algorithm for breaking RSA 2028 requires around ten TeraQuops (Ekera, Gidney 2021). Chemical simulation problems appear to be just as demanding in terms of logical operations.

Achieving this TeraQuop regime means that each elementary operation on one or two qubits can only tolerate an error rate of the order of 10<sup>-12</sup>. Such values are easily accessible for classical transistors, but much more difficult to obtain in the quantum world.

<sup>20</sup> An algorithm with a very large number of quantum gate cycles is called "deep".

The experimental state of the art is currently around 10<sup>-3</sup>, and even if we can hope for an improvement of one or two orders of magnitude in the future, it seems unlikely that we will be able to gain the needed 9 orders of magnitude simply by improving the experimental devices that are the physical qubits and their control systems.

Fortunately, the solution has been well known since the seminal work of Peter Shor and is based on quantum error correction and fault tolerance. Once thought impossible because of the continuous nature of quantum states and errors, error correction is now based on solid theoretical foundations. The idea here is to encode the logical information, which we wish to protect, in a delocalized way across a large number of physical systems. These systems are typically disrupted by local phenomena that have little long-distance correlation and will not be able to access the protected information. In particular, regular parity measurements will enable any errors that have occurred to be identified and corrected. Unless we discover completely new error mechanisms where nature conspires to disrupt physical systems in such a way as to systematically prohibit quantum computation, these quantum error correcting codes will make it possible to protect logical information at the level required to perform the algorithms we want.

A crucial concept in this context is that of the logical qubit, which corresponds to a quantum bit defined in a higher-dimensional physical space. This space can consist of a large number of two-level physical systems, i.e. physical qubits, or of one or more bosonic modes, which are infinite-dimensional spaces. To be useful, these logical qubits must possess two important properties: their quantum state must be capable of being manipulated by simple physical operations performed on the physical systems that make them up, and these operations must be capable of being performed with error levels lower than the intrinsic physical error level. This is achieved through quantum error correction.

What fault-tolerance theory shows is that for realistic noise models, there is a threshold value for the physical error rate, such that if the true error rate is below this threshold, then arbitrarily long quantum computations can in principle be performed by encoding logical information into a larger number of physical qubits. The precise values of this threshold and above all of the redundancy required for fault tolerance are not yet well known, but it is anticipated that the threshold is around 0.001, and that a logical qubit could be encoded in a thousand physical qubits, possibly much less. This cost is certainly very high even though not prohibitive: e.g. for the factorization, we switch from a computation involving a few thousand perfect qubits for example to a physical processor comprising around ten million noisy physical qubits. Furthermore, these values are probably

very pessimistic and improvements of several orders of magnitude are conceivable with the development of better error correction techniques.

#### 2.2. ERROR CORRECTION AND FAULT TOLERANCE

We need to distinguish between two guite different objectives. Error correction involves protecting a memory against a very specific type of error. This is the case, for example, with conventional communications: information is encoded redundantly before being transmitted on an imperfect channel subject to errors, and the receiver seeks to correct the errors in order to obtain the sent message. A memory is similar: in this case, the imperfect channel corresponds to the interactions between the quantum memory and its environment during the time the information remains stored, and the aim is for the information to remain intact for as long as possible. This scenario crucially assumes that errors are confined to the information transmission stage, but that we have access to ideal encoding and decoding techniques, unperturbed by noise. This scenario only imperfectly models the relevant task in the quantum world, where every manipulation of the information is inevitably noisy: we cannot therefore assume perfect encoding or decoding. On the contrary, every physical operation designed to correct the noise is itself a source of noise! In this case, it's a question of correcting the errors as quickly as they appear. In fact, managing to keep the number of physical errors in the system below the correction threshold throughout the calculation is sufficient, because the very last stage of a quantum computation involves measuring the state of the computer, which returns classical information, the result of the calculation. Since this information is classical in nature, it is easy to protect using classical coding techniques, and any residual errors can be corrected without difficulty.

In addition to having to deal with physical operations that are all imperfect, the other challenge facing *fault tolerance* is that we don't simply want to protect logical information, but on the contrary to be able to manipulate it. We therefore find ourselves torn between two contradictory objectives: to isolate the information as well as possible from its environment in order to protect it, while at the same time making it possible to manipulate it in order to perform the desired calculation. These two ambitions are relevant both at the level of the physical system and at the level of the logical information. At the physical level, we need to be able to apply physical operations on the qubits, which requires the existence of control lines for each system, and this can be a source of interference or problematic inaccuracies. At the logical level, we want to be able to carry out operations that are simple to implement so as not to allow the inevitable errors to propagate too much throughout the system, while remaining protected in a quantum code. We now know that the most naïve approaches to achieving this objective are doomed to failure: for example, the Eastin-Knill theorem shows that an encoding in which each logical gate is obtained by applying the same physical gate to each of the underlying physical qubits cannot both correct errors and perform a calculation that is difficult for a

classical computer. Several solutions for getting round this problem are known today, in particular the distillation of magic states, which is a fault-tolerant protocol for producing very pure quantum states, known as magic states, from noisy copies: these states, which cannot be obtained directly in an error-correcting code of the stabilizer type, are needed to implement non-Clifford gates, such as the T gate, and therefore make quantum computation universal. Distillation is based on circuits composed solely of Clifford operations and measurements, and plays a central role in the architecture of fault- tolerant quantum computers.

All the solutions for getting round the problem of errors illustrate the difficulty of carrying out quantum calculations in a world where experimental devices are necessarily imperfect.

The task of fault tolerance is therefore as follows: given a quantum algorithm that we wish to implement, and a realistic noise model for the physical system encoding the information, we need to find a way of protecting and manipulating the logical information that will enable us to emulate the considered theoretical calculation. This always has a cost in terms of redundancy; how many physical systems are needed per logical qubit? The aim is to understand how to reduce this factor as much as possible. The potential sources of improvement are diverse: improving the quality of the physical systems and their control, finding better quantum error correcting codes, developing more efficient techniques for manipulating the encoded information at the logic level. The most widely studied theoretical solutions are based on the use of the surface code, which aims to correct two types of error (bit-flip and phase-flip). These solutions imply a cost of around 1,000 to 10,000 physical qubits per logical qubit for algorithms such as factorization. More recent solutions rely, for example on new, more efficient error correcting codes or on so-called bosonic qubits equipped with error selfcorrection mechanisms that intrinsically protect the gubits against one type of error cause (bit-flip). These solutions could save several orders of magnitude, and will be discussed in section 2.5.

#### 2.3. The principle of quantum error correction

The state of a quantum bit is a vector  $|\psi\rangle = \alpha |0\rangle + \beta |1\rangle$  in a complex space of dimension 2. Here,  $\alpha$  and  $\beta$  are two complex numbers, called amplitudes, such that  $|\alpha|^2 + |\beta|^2 = 1$ , and the two vectors  $|0\rangle$  and  $|1\rangle$  correspond to states 0 and 1. The fact that the amplitudes take a continuum of values is reminiscent of analogue computers, where errors cannot be corrected. In the quantum case, there are also an infinite number of possible errors that can disrupt the state of the system. The crucial difference between the two scenarios lies in the linearity of quantum mechanics: this guarantees that if we know how to correct the errors corresponding to a basis of possible errors, then we can correct any error. In particular, it is enough to know how to correct bit-flip errors (which exchange the |0> and |1> states), phase errors (which add a phase -1 to the |1> state), and the combination of these two errors to be able to correct any error. In other words, we're back to a discrete error model, where we have to consider twice as many errors as in classical error correction: the bit-flip errors are the same as in the classical case, and the phase errors are new and specific to quantum computation.

In particular, the work of Daniel Gottesman, who defined the stabilizer codes, and that of Calderbank, Shor and Steane has made it possible to import many classical error correction techniques into the quantum world. The general idea is to encode a certain number of logical qubits in a larger number of physical qubits, and to measure parities in order to project the true error onto a bit or phase error, then to use a classical decoding algorithm to identify the error in question and correct it.

#### 2.4. THE GOLDEN AGE OF SURFACE CODE

Soon after the pioneering work mentioned above, researcher Alexei Kitaev discovered a remarkable guantum error correcting code: the surface code. This code encodes a logical qubit into a planar arrangement of a set of L x L physical qubits on a square grid, with connectivity between neighbouring qubits, perfectly suited to architectures based on superconducting gubits. Below the threshold, the logical error rate plog decreases very rapidly with the physical error rate poly. This code, and its generalizations known as topological codes, whose properties are intimately linked to those of the topological manifold (2D surface or code defined in 3 dimensions) that defines them, have been the subject of particular study for nearly 20 years. In particular, we know how to correct errors efficiently and implement logic gates on these codes. For example, the general idea behind quantum computing with these codes is to encode each logical gubit in a surface code. Logical operations are then carried out using so-called lattice surgery techniques aimed at merging and then splitting two surface codes, as well as with the distillation of magic states. By taking the value of L sufficiently large, typically around 30, we can ensure that all logical operations have the desired error level, and therefore with an overhead of around 1,000 physical gubits per logical gubit for the TeraQuop regime, bearing in mind that this number depends closely on the fidelity of the physical quantum gates, which tends to improve gradually.

For a long time, researchers thought they had the ultimate code for quantum computing. The limitation to local architectures in two (or possibly three) dimensions was dictated by the physical constraints of superconducting circuits, and we now know that in two dimensions the surface code is essentially optimal. In addition, the topological properties of the code are used to develop logical gates that do not propagate errors during calculation. For these reasons, research into alternative approaches remained fairly confidential during this period.

# Experimental state of the art: superconducting platform

With the advances made over the last fifteen years in superconducting qubit design, nanofabrication techniques and microwave environment filtering methods, this platform has reached the maturity required to produce errorcorrecting codes with physical error rates below the fault tolerance threshold. Among superconducting qubits, the transmon, thanks to its simple design and very good coherence times, is currently the number one choice of several industrial and academic players for scaling up to a fault-tolerant quantum computer.

The constraint of 2D connectivity on a chip largely limits the choice of errorcorrecting codes. The dominant approach, followed by Google and IBM, for example, is to create a surface code based on transmons. In a recent experiment published in Nature in 2023, on a chip made up of 72 transmons, the Google team has demonstrated the creation of a surface code of distances 3 and then 5, corresponding respectively to sizes of 3× 3 and 5× 5 physical data qubits. In this experiment, all the gubits benefit from gate fidelity for one-qubit gates of the order of 99.9%-99.95% (i.e. error rates of 0.0005 to 0.001) and for two-qubit gates of the order of 99%-99.8%. However, gubit reading is less faithful (96% - 99%) and, above all, suffers from a very long duration compared with the other operations (500 ns compared with 25 ns for 1-qubit gates and 34 ns for 2-qubit gates). A reset time of 160 ns must be added to this. This slowness and the infidelity of the measurement process, as well as the 2- gubit gate errors, are the main reasons for an error rate of around 2% to 4% per error correction cycle. The Google team has nevertheless managed to demonstrate a slight improvement in the error rate from 3.028± 0.023% (for a distance 3 surface code) to 2.914± 0.016% per error correction cycle (for a distance 5 surface code). This indicates that the mean values of the physical components in this experiment are at the fault tolerance threshold. This experiment represents the most complete implementation of a quantum error correction protocol to date of any existing platform, with several repeated error correction cycles and scaling from distance 3 to distance 5. A rapid improvement in the performance of this realization is conceivable in the coming years. Indeed, the major physical error component per error correction cycle corresponds to the duration and fidelity of transmon reading. However, the performance achieved in the Google experiment is not at the level of the state of the art. In particular, with the integration of the latest filtering techniques, academic groups

(ETH Zurich, Riken) have demonstrated much faster and more faithful readings (typically a reading time of 40 ns and 99.6% fidelity). While this report was being written, new advances by the Google team appeared in the form of a pre-publication on arXiv (arXiv: 2408.13687). In this prepublication, the Google team announces the completion of two new chips at 72 and 105 qubits with improved quantum memory fidelities compared with the 2023 experiment. This improvement is attributed to the higher T1 relaxation time and T2 coherence time (by a factor of 3). These advances enable the team to demonstrate the operation of a surface code of distances 3, 5 and 7 below the fault tolerance threshold (by a factor of 2.14). Notably, the distance 5 and 7 codes also achieve the break-even corresponding to a lower logical error rate per error correction cycle than the most coherent physical component of the system. Finally, the 72-qubit chip is equipped with a real-time decoder with an average latency of 63 microseconds at distance 5 (compared with the error correction cycle time of 1.1 microseconds).

It seems plausible that this approach could in the short term reach a level of maturity where the physical error rate would be up to an order of magnitude below the fault tolerance threshold, on chips with several hundred gubits. However, even with this level of performance, we will need a thousand transmons per logical gubit to achieve error rates of the order of 10<sup>-10</sup> per error correction cycle. Furthermore, to ensure effective error correction by the surface code, it is necessary to eliminate the sources of error that act in a correlated manner on the physical qubits of the surface code. In recent years, two major sources of correlation have been identified and studied: the leakage of the transmon population to excited states beyond the computational space, which can then propagate in the circuit to create problematic error correlations in time and space; and the impact of high-energy events leading to the creation of phonons propagating in the substrate and thus generating an overpopulation of quasiparticles throughout the sample. These two correlated sources of error have been studied in recent years and solutions have been proposed. Google's recent experiment shows that by applying these solutions, logical error rates of 10<sup>-10</sup> are potentially achievable, but that going below these rates would require the correction of correlated errors of different natures.

With the significant extra cost of error correction (around 1,000 transmons per logical qubit), scaling beyond one logical qubit would certainly come with complications. Larger systems, for example to demonstrate fault-tolerant logical gates or to distil magic states, will certainly require advances in materials science, for example to better understand and control the creation of parasitic 2-level systems that disrupt the operation of this multi-qubit system. In parallel with research into this type of problem, enabling the maturation of the transmon-based platform, academic and industrial players

are also working on alternative ways of reducing, sometimes significantly, the additional cost of error correction. For instance, here are two dominant approaches in the context of superconducting circuits:

- Replacing transmon qubits with other superconducting qubits: several candidates (Fluxonium, heavy Fluxonium, 0-π qubit, cos 2φ qubit, etc.) with different levels of maturity are currently being studied. These qubits can provide a level of intrinsic protection against the dominant sources of noise (charge or magnetic flux noise), and make it easier to achieve error rates below (and sometimes well below) the fault tolerance threshold. This would imply that a smaller surface code would be sufficient to achieve a given logical error rate.
- Dual-rail encoding replacing the errors taken into account by surface codes (bit-flip, phase-flip) with erasure errors: in this approach inspired by the photonics platform (pursued, for example, by the start-up QCI and by AWS), the information of a qubit is encoded on the single-excitation states of two superconducting modes. Energy relaxation, which is a major source of noise in qubits, means that the information is erased with a transition to the zero-excitation state, but this can be observed by an efficient reading of this state. This results in a first-order replacement of the dominant Pauli errors by erasure errors that are more easily corrected by a surface code. In the absence of Pauli errors, this simplification results in an error correction threshold rising from 0.5% to around 4%.

#### 2.5. NEW APPROACHES REQUIRING LESS HARDWARE

While surface code was long considered to be the best road to fault-tolerant quantum computing, new ideas have emerged in recent years that have largely reshuffled the deck, to such an extent that it is now very difficult to identify which architecture will enable fault-tolerant quantum computing to be achieved first. Two basic assumptions have been called into question, namely that physical information must necessarily be encoded on two-level systems and that physical systems are constrained to local connectivity in two dimensions. Questioning the first hypothesis led to the development of bosonic codes, while questioning local connectivity led to the advent of quantum LDPC or qLDPC (*low-density parity-check*) codes. These two approaches offer remarkable prospects in terms of reducing the number of physical systems per logical qubit, typically by an order of magnitude.

In both cases, the theoretical ideas are not completely new, since bosonic codes had already been studied in the late 1990s in the context of photonic codes and quantum LDPC codes considered in the early 2000s. But the renewed interest in these exotic approaches is based on recent experimental considerations. For example, the use of strong non-linearities provided by Josephson junctions in superconducting circuits makes it possible to achieve with microwave photons what was not possible with optical photons. On the other hand, the advent of platforms such as reconfingurable array of neutral atoms makes it possible to couple qubits arbitrarily by moving the corresponding atoms, thus freeing ourselves from the constraints of local codes in two dimensions, at the cost, however, of two-qubit quantum gates which are currently very slow (1 ms).

From a computer science and error-correcting code point of view, these two approaches illustrate perfectly the importance of creativity in solving problems such as fault tolerance. The playing field is immense, and it is almost certain that the solutions being considered today will not be the ones that will eventually be deployed, with too many avenues remaining unexplored to date.

What the bosonic codes challenge is the idea that logical gubits need to be encoded on two-level physical systems. In fact, for a large number of physical systems, there are not necessarily two very natural levels for encoding a physical aubit. For example, the state of the microwave field confined by a superconducting circuit is described by an infinite-dimensional Hilbert space, called a Fock space, and the traditional approach is to consider the two lowest energy levels to define a physical gubit. To do this, we add a Josephson junction to the circuit, which generates an anharmonicity that allows the first two levels to be decoupled from the higher energy levels. Getting rid of the other levels in this way has an experimental cost, and one interesting idea is to understand whether, on the contrary, we could not make these other levels an asset for error correction. After all, the very principle of error correction is to exploit physical redundancy to protect logical information. The idea of bosonic codes is therefore to define a better qubit within this infinite-dimensional space. For example, the space generated by two coherent states  $\alpha$  and  $\beta$  and  $\beta$ , of opposite amplitudes  $\alpha$ and  $-\alpha$ , is the choice put forward by the cat qubits. The theoretical interest of this new physical qubit is that it is already partially protected thanks to the redundancy offered by the bosonic code. In particular, by choosing a sufficiently large value of  $\alpha$ , the qubit is intrinsically protected against bit-flip errors, and logical encoding just needs to provide additional protection against phase errors. This can be implemented with simple classical codes, dramatically reducing the cost of a logical qubit. Of course, new complications arise with this scheme: in particular, it is crucial that the physical circuit never transforms the inevitable phase errors into bit-flip errors, which are assumed to have been handled by the bosonic code.

This cat qubit is just one of many possibilities. For example, a potentially more interesting encoding is the GKP code by Gottesman, Kitaev and Preskill, which provides simultaneous (but partial) protection against both types of error. This is achieved at the cost of a more complex stabilization scheme for the physical qubit, but could be a very interesting solution in the longer term.

## Experimental state of the art: superconductors, ions and photonics

The strong, dissipation-free non-linearity provided by Josephson junctions makes it possible to generate, manipulate and stabilise non-classical states of the microwave field. Cat gubits can thus be stabilised by a two-photon exchange mechanism. A recent experiment by the Alice&Bob start-up and the Quantic team (ENS, Mines, Inria, CNRS) shows that gubits stabilised in this way are exponentially protected against bit errors, achieving bit-flip times of the order of ten seconds (an improvement of 6 to 7 orders of magnitude compared with the unprotected case). This experiment also demonstrates the implementation of a one-qubit gate (rotations around the Z axis) maintaining the same protection. To correct the other type of error, phase errors, it is then sufficient to use conventional error-correcting codes, which correct only one type of error as in the case of classical communications, and which can now be very close to a classical code (with a few specific features linked to the fact that we want to be able to carry out logical operations that do not introduce bit errors either). This leads to several simplifications and shortcuts in terms of hardware overhead: for example, it is possible to concatenate such an encoding with a classical LDPC code in a 2D architecture and achieve very interesting encoding rates. A demonstration of gates (in particular the two-gubit CNOT gate) preserving the intrinsic protection against *bit-flips*, as well as scaling up to tens of cat qubits to demonstrate the correction of phase flips, are the next steps in demonstrating the viability of this approach. This is the approach followed by the Alice&Bob start-up and by AWS. In addition, GKP encoding (pursued, for example, by the Quebec start-up Nord-Quantique) provides more symmetrical low-level protection, eliminating both bit errors and phase errors. A demonstration of this fault-tolerant low-level correction (which eliminates the propagation of errors from the syndrome measurement system to the uncorrectable errors of the harmonic oscillator), protectionpreserving logical gates and concatenation with a small surface code, are necessary steps to demonstrate the viability of

this approach. In parallel with these advances on the superconductor platform, similar experiments are being carried out at ETH Zurich on a trapped ion platform. Here, the information is encoded in GKP states of the vibrational mode of an ion. The non-linearity required for state generation, syndrome measurement and error correction is provided by coupling this vibrational mode to the internal state of the ion. The state of progress of these experiments is quite similar to that of the GKP experiments with superconducting circuits.

With photonic qubits, the preferred computing paradigm is *Measurement*-Based Quantum Computing (MBQC)<sup>21</sup>, which is based on a network of entangled gubits and measurements by successive layers of gubits. This approach is also universal (see Chapter 1) and therefore has the same computational power as the most widespread model of quantum circuits, and it can also be made fault-tolerant<sup>22, 23</sup>. In this case, the error-correcting codes described above can be translated into an MBQC version by the foliation process, and there are also correcting codes specifically built for the MBQC model which have no circuit counterpart<sup>24</sup>. To carry out an MBQC calculation, a large network of entangled gubits is required. This can be created as the calculation progresses, but a certain number of qubits and entanglement links must still exist simultaneously, and producing such a network of entangled qubits is not easy. One way of meeting this challenge is to create small entangled states, known as resource states, and to entangle these states with each other using so-called fusion operations<sup>25</sup>, which correspond to a Bell measurement, a type of two-qubit measurement that is particularly well suited to

<sup>21</sup> R. Raussendorf, H. J. Briegel, "A one-way quantum computer", Phys. Rev. Lett. 86 5188-5191 (2001).

<sup>22</sup> R. Raussendorf, J. Harrington, K. Goyal. "A fault-tolerant one-way quantum computer", *Annals of physics*, 321(9):2242-2270 (2006).

H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, M. Van den Nest. "Measurement-based quantum computation", *Nature Physics*, 5(1):19-26 (2009).

<sup>24</sup> N. Nickerson, H. Bombín, "Measurement-based fault tolerance beyond foliation", arXiv:1810.09621 (2018).

<sup>25</sup> D. E. Browne, T. Rudolph, "Resource-Efficient Linear Optical Quantum Computation", *Phys. Rev. Lett.* 95 010501 (2005).

photonics. However, fusion operations are probabilistic, and building a network of entangled qubits in this way is based on the theory of percolation: more qubits than necessary for the desired state are fused, and if the probability of success of each fusion is sufficiently high, the targeted state can be achieved. A more recent approach, known as *Fusion-Based Quantum Computing* (FBQC)<sup>26</sup> takes this fusion process directly into account, making it possible to incorporate the probabilistic nature of fusion into the error correction protocol. A FBQC computing architecture then corresponds to the specification of a type of resource state and a fusion network, and, for a given fusion network, a calculation is specified by the choice of measurement bases for the different fusions.

With photons, demonstrations of small-scale error correction are more difficult than with stationary gubits, because the photons are consumed as the gates are executed. The threshold effect characterising the passage into a region where error correction is advantageous is therefore more marked, and it is difficult to carry out experiments where the logical errors are of the same order of magnitude as the physical errors, as is the case with the experiments described above for superconducting gubits, trapped atoms or ions. Demonstrations of error correction do exist, however, such as the creation of a Shor code encoding a logical qubit on nine qubits<sup>27</sup> and are useful for the development of computing as well as for quantum communications. But some players find it more convincing to demonstrate that the basic components are on the way to achieving the values needed to pass the threshold and to concentrate on the manufacturing capacity that will enable the transition to scale<sup>28</sup>, rather than carrying out small-scale proofs of concept. Finally, when several degrees of freedom are available, for example with quantum emitters with controllable spin,

<sup>26</sup> S. Bartolucci, P. Birchall, H. Bombín, et al. "Fusion-based quantum computation", Nat. Commun 14 912 (2023).

<sup>27</sup> R. Zhang *et al.* "Loss-tolerant all-photonic quantum repeater with generalized Shor code", *Optica* 9 152 (2022).

<sup>28</sup> K. Alexander et al. "A manufacturable platform for photonic quantum computing", arXiv:2404.17570 (2024).

Hybrid architectures can be considered, allowing a high degree of modularity in the implementation of error correction and fault tolerance  $protocols^{29,30}$ 

Even if we confine ourselves to physical gubits defined on two-level systems, it is possible to do much better than the surface code. The idea is to allow connections between gubits that are not neighbors on a two-dimensional grid. In this case, if the parity checks required for error correction involve only a small number of qubits, we are dealing with a so- called quantum LDPC code. The advantage of these codes is that they offer much better performance than the surface code. Whereas the surface code encodes k logical qubits in k L<sup>2</sup> physical qubits while being able to correct errors affecting (L-1)/2 qubits, quantum LDPC codes can encode k logical gubits in ck physical gubits and correct errors affecting c'k qubits, for constants c, c'>0. While these existence results are theoretical for the moment, small quantum LDPC codes have already been optimized, notably by IBM, and provide a gain of a factor of 10 in terms of redundancy compared with the surface code, to encode 12 logic qubits for one million error correction cycles [Bravyi et al, Nature 2024]. Here again, this improvement is accompanied by new difficulties, such as how to efficiently manipulate the logical information encoded by these codes. The theoretical approach of the IBM researchers targets the superconducting circuit platform and proposes a biplanar architecture with long-distance couplers. While a biplanar architecture seems to be the right way forward with the development of flip-chip technology, the effectiveness of long-distance couplers has yet to be proven by experiments. Other platforms, notably those based on atoms, lend themselves particularly well to encoding beyond 2D connectivity.

<sup>29</sup> S. Simmons, "Scalable Fault-Tolerant Quantum Technologies with Silicon Color Centers," PRX Quantum 5 010102 (2024).

<sup>30</sup> G. de Gliniasty et al. "A Spin-Optical Quantum Computing Architecture", Quantum 8 1423 (2024).

# Experimental state of the art: neutral atom platform

Neutral atoms trapped in optical tweezers have seen a significant leap in their degree of controllability in recent years. These major advances have promoted this platform, initially developed as an ideal candidate for the analogue simulation of quantum many body systems, to an excellent candidate for fault-tolerant digital quantum computing. Encoding the quantum information in the hyperfine clock states of the trapped atom (Rubidium in the case of start-ups Pasqal and QuEra) ensures robustness (long coherence times). Then, excitation towards the Rydberg levels ensures strong interaction between the atoms to implement high-fidelity 2-qubit (or multi-qubit) gates.

More specifically, one-qubit gates are implemented by Raman processes and today achieve fidelities of the order of 99.97%, with durations of the order of 500 ns. CZ's two-qubit gates are based on the Rydberg blockade phenomenon: when two atoms approach each other, the passage of one atom through the Rydberg level prevents the passage of the other. It is therefore possible to prepare entangled states of two atoms in proximity and take advantage of this to implement a CZ gate. In recent experiments at Harvard and QuEra, fidelities of the order of 99.5% and gate times of 200 ns have been demonstrated. High-fidelity redout of these gubits is ensured by fluorescence imaging. While high fidelities of the order of 99.9% have been demonstrated, these measurements suffer from certain limitations. Non-destructive measurements are possible, but remain very slow compared with other operations (tens of milliseconds). A relatively fast measurement (10 microseconds) is possible, but is accompanied by the loss of the measured atoms. It is therefore necessary to permanently and efficiently replace these measured atoms with freshly prepared atoms.

In terms of computing architecture, this platform differs greatly from superconducting circuits. The superconducting qubits each have their own control and readout line, which makes it possible to implement complex logical circuits and ensures a decorrelation of control errors. However, this certainly comes with complexities related to scaling and multiplexing. In addition, this platform is limited by 2D or near-2D (e.g. biplanar) connectivity. Addressing neutral atoms, for single-qubit operations, can be done directly by combining global microwave excitation and focused laser beams to produce frequency light shifts on selected qubits. But for all other operations (multi-qubit gates and readout) and even for easier scaling with single-qubit operations, the most promising approach is to coherently

transport atoms using 2D acousto-optic deflectors. By transferring the atoms to appropriate zones, it is then possible to perform the necessary operation on several atoms simultaneously. This approach significantly reduces the number of control lines. In addition, this transport can make it possible to achieve more or less arbitrary connectivity, which can make it possible to implement LDPC codes with much better performance than the surface code. However, this parallelization of operations can be accompanied by limitations such as correlated errors. Typically, the laser noise that collectively excites the atoms can potentially give rise to correlated errors on these atoms. It is therefore necessary to demonstrate in future experiments that this correlation is not a hard limit to the effectiveness of the error correction.

In a recent experiment, published in *Nature* in 2024, the Harvard team and the start-up QuEra combine these capabilities to build a set of building blocks for a fault-tolerant computer on a device made up of 280 qubits. This includes preparing the logical states of a surface code of distances 3, 5 and 7 and creating a transverse CNOT gate between two logical qubits encoded in this way. In addition to the surface code, the authors have produced several logical qubits encoded in 2D colour codes encoding one logical qubit in 7 physical qubits, with distance 3, and subsequently used in a fault-tolerant preparation of a GHZ (Greenberger-Horne-Zellinger) entangled state between 4 logical qubits encoded in this way. They also demonstrate their ability to perform efficient mid-circuit readouts by implementing a *feed-forward* entanglement teleportation protocol.

Lastly, the richer connectivity between the qubits enables them to prepare the logic states of 3D error detection codes encoding 3 logical qubits in 8 physical qubits with a distance of 2, which in particular enables non-Clifford CCZ gates to be implemented transversely within a code block.

Despite this impressive set of demonstrations, the experiment is limited to the preparation and manipulation of code states and does not demonstrate error correction through repeated measurements of error syndromes. It therefore remains to be demonstrated that the implementation of error correction can indeed lead to an exponential improvement in error rates. The major obstacle in this context is due to the fact that the dominant errors in the operations of this platform are not the Pauli errors for which these codes were designed. In particular, two major types of error are the leakage of information outside the computation space and the loss of atoms in the optical tweezers. For the first type of error, it is possible to convert these leakage errors into erasures, similar to those discussed in alternative approach 2 for superconducting circuits. But this would certainly require rubidium to be replaced by other atomic species such as ytterbium or strontium. Indeed, the richer structure of the electronic energy levels of ytterbium and strontium makes it possible to observe population leakage and bring the atom back into computational space without disturbing normal operation in the absence of leakage. For the second type of error, it is also possible to treat them as erasure errors after efficient and rapid injection of the freshly prepared atoms. It therefore seems necessary to demonstrate all of the above capabilities with these modifications before being able to achieve exponential<sup>31</sup> error suppresion with error-correcting codes.

<sup>31</sup> The gain obtained with error correction is said to be exponential when the error rate of the logical qubits is reduced exponentially, while the number of physical qubits increases non-exponentially (e.g. only polynomially).

#### 2.6. The challenges ahead

As we can see, the field of quantum error correction and fault tolerance is booming. The first experiments demonstrating the performance of these methods are being carried out right now in several laboratories, and numerous theoretical avenues for improvement are being studied by the community.

While the theoretical foundations of the field are well understood, the experiments to come will be decisive, as they will enable us to explore completely new regimes, corresponding to logical error probabilities of less than  $10^{-5}$  or  $10^{-6}$  in the short term. These experiments will be crucial for understanding whether the error models studied today remain relevant, or whether new errors, correlated for example, come into play in these regimes. We already know that high-energy particles can create errors that are problematic for superconducting circuits. No doubt we will discover other similar phenomena. These correlated errors are eminently problematic for the quantum codes we are considering today, but are very probably correctable once identified and understood. All we need to do is develop new error correcting codes with the required properties.

Generally speaking, research into quantum error correcting codes and fault tolerance is still in its infancy. The last two decades have been largely devoted to surface codes, but it is now clear that many alternatives exist and that many avenues remain to be explored.

One difficulty is the small size of the research community, probably just a few dozen researchers in France and less than a thousand worldwide, faced with a fairly colossal task. One of the challenges in the years to come will be to train more students in these subjects, and perhaps also to enable part of the community working on NISQ to redirect their attention to error correction issues.

## Metrics for judging the success of experiments

To assess the success of the various experiments, the players in the field use different metrics, linked to the limitations of each platform. To ensure the viability of an approach towards a fault-tolerant processor, it is necessary to demonstrate the exponential suppression of errors by increasing the distance of the code. It is therefore necessary to repeat several error correction cycles with codes of different distances and achieve this exponential suppression below the fault tolerance threshold. These are the experiments carried out, for example, by the Google team with surface code of distances 3, 5 and 7. The recent Quantinuum/Microsoft experiment on a platform made up of 32 gubits encoded in the states of trapped ions also shows these repeated error correction cycles on a code [[12,2,4]] called "Carbon code". Here, with the distance fixed at 4, the authors demonstrate an improvement in the logical error rate below the physical error (we then speak of break-even). This metric justifies the success of correction in eliminating errors, but does not allow us to see whether, by scaling the distance, we can achieve very low logical error rates of the 10<sup>-10</sup> type. As each code has a higher capacity for error detection than error correction, experiments often use a post-selection of error-free events, which results in much lower logical error rates. Platforms that suffer from limitations in terms of mid-circuit reading of gubits, or non-Pauli errors (leakage or loss of atoms), also benefit from these post-selection experiments to anticipate the performance of their platform once these limitations have been lifted. Although this postselection can be used to mitigate errors in medium-sized calculations, it is the exponential suppression of the logical error rate in an error correction experiment that ensures viability towards a fault- tolerant processor capable of handling problems of arbitrary size.

### Chapter 3

## QUBIT TECHNOLOGIES: CHALLENGES AND ROADMAPS

In this part of the report, we examine the state of the art, the challenges and the roadmaps for creating fault-tolerant quantum computers. We focus on five main categories of qubits: superconducting qubits, neutral atoms, trapped ions, quantum box spins and photons. These are the categories with the highest level of technological maturity on the market<sup>32</sup>.

In all, nine major categories of qubit types can be identified, themselves usually grouped together in a simplified way into three groups: qubits using isolated atoms, those relying on electrons and their spin in variable structures (superconducting artificial atoms, electrons trapped in potential wells, electrons in gaps in crystalline structures or vacancies, topological qubits, etc.), and finally, flying qubits which are generally photons, but could also potentially be electrons.

<sup>32</sup> Other types of qubit, such as Microsoft's topological qubits or cavities in diamonds, are at a lower level of maturity at this stage.

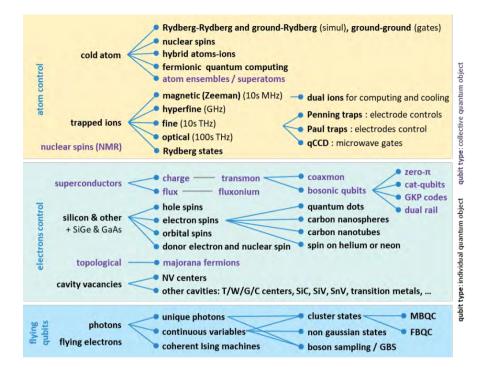


Figure 4: A map and genealogy of the main known types of qubit. The qubits in violet have a physical support based on several quantum objects. (cc) Olivier Ezratty, 2021-2024.

The vast majority of the types of qubit inventoried (Figure 4) are supported by industrial vendors, demonstrating the wide range of technological bets made by academic researchers, industry vendors, the private investors who support them and the governments who fund them directly or indirectly.

At the beginning of 2025, over 90 industry vendors had been identified (Figure 5). The level of technological maturity of the offerings varies widely. Some have not even demonstrated a qubit operation or even a two-qubit quantum gate. Others have exceeded a hundred functional qubits, notably with superconducting qubits and neutral atoms. Trapped ions stand out with very good quantum gate fidelity, but with greater difficulty in increasing their number.

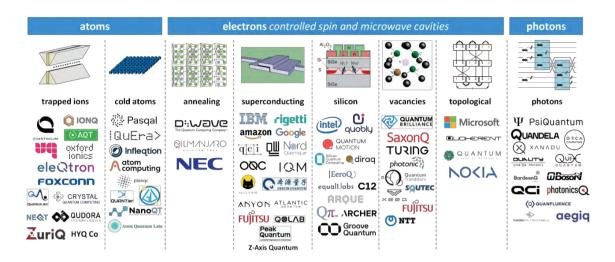


Figure 5: Overview of industrial players in the market (cc), Olivier Ezratty, March 2025.

	atoms		electrons superconducting & spins			photons
				Al 2Q2	8	***
	cold atoms	trapped ions	superconducting	silicon	NV centers	photons
qubit size	about 1 µm space between atoms	about 1 µm space between atoms	(100µ)²	(100nm) <sup>2</sup>	<(100nm) <sup>2</sup>	nanophotonics waveguides lengths, MZI, PBS, etc
best two qubits gates fidelities	99.5% (QuEra)	99.84% (56 q, Quantinuum)	99.68% (IBM Heron r2)	>99% (SiGe)	99.2%	99% (but probabilistic)
best readout fidelity	97%	99.99%	99.4%	99% (SiGe)	98%	89% (efficiency)
best gate time	≈1 ns	0.1 to 4 µs	20 ns to 300 ns	≈5 µs	10-700 ns	<1 ns
best T <sub>1</sub>	>1s	0,2s-10mn	100-400µs	20-120µs	2.4 ms	$^\infty$ & time of flight
qubits temperature	< 1mK 4K for vacuum pump	<1mK 4K cryostat	15mK dilution cryostat	100mK-1K dilution cryostat	4K to RT	1.8K-4K for photons gen. & det.
operational qubits	1,600 (Infleqtion)	56 (Quantinuum)	156 (IBM) 72 (China)	12 (Intel) in SiGe	5 (Quantum Brilliance)-10	216 modes GBS (Xanadu)
scalability	up to 10,000s per QPU	1000s, with multiple traps	<1,000 per chip	potentially millions	unknown	100s-1M in multiple circuits

Figure 6: Figures of merit for the main types of qubit. The gate times include the electronic control, but not the classical part of the processing. T<sub>1</sub>are the best relaxation times for qubits. The fidelity of two-qubit gates evaluates in percentages the probability that the operation will generate the expected result. These are the best figures of merit taken in isolation. This does not mean that a quantum computer gathers all the figures of merit in a single column!

Some data sources: <u>Neutral Atom Quantum Computing Hardware: Performance and End-User Perspective</u> by Karen Wintersperger *et al*, April-May 2023 (27 pages), trapped ions (<u>TrappedIonQuantumComputing:ProgressandChallenges</u>, 2019,<u>MaterialsChallengesfor</u> <u>Trapped Ion Quantum Computers</u>, 2020, <u>Infineon</u>, IonQ and Quantinuum, <u>High-Fidelity Bell-State Preparation with 40Ca+Optical Qubits</u> by Craig R. Clark *etal*, PRL, September 2021 (7 pages)), silicon (<u>Roadmap on quantum nanotechnologies</u>, 2020), superconducting (many IBM papers), NV centers (<u>Quantum computer based on color centers in diamond</u>, 2021). I list only the most demanding two qubit gates and readouts fidelities. Cold atoms systems are usually simulators, but data pertains to gate-based implementations. Olivier Ezratty, 2020-2024.

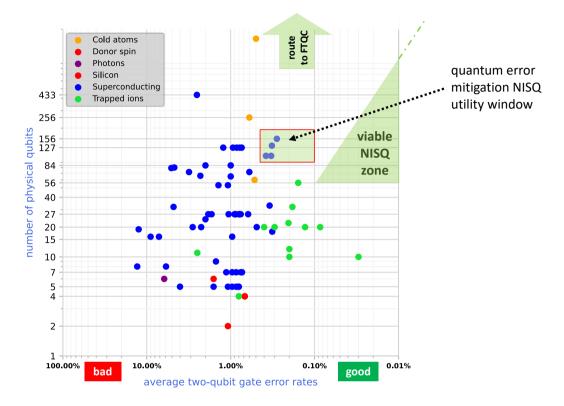
Comparing the technological maturity of these approaches is very tricky,

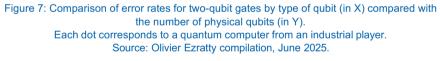
particularly when dealing with the prospects of building fault-tolerant quantum computers (figure 6). By analogy, we are more in the field of climatology than weather forecasting, in the sense that we can certainly compare today's figures of merit (weather), but their evolution over time depends on a very large number of parameters, interdependencies and uncertainties (climate).

Figure 7showcases the fidelities of two-qubit gates operations. It corresponds to the most relevant figure of merit to compare these platforms, at the very least, in a NISQ regime (). It illustrates the challenge of increasing both these gate fidelitie and the number of qubits involved. Ideally, in the NISQ regime, we should have fidelities greater than 99.9%, which increases with the number of qubits used.

This is also the threshold that needs to be reached for error correction to be viable on a large scale.

Figure 8 presents a summary of the challenges facing each type of qubit. For each of these challenges, many academic and industry groups are considering a wide variety of solutions that are still difficult to evaluate.





	superconducting	neutral atoms	trapped ions	silicon spins	photons
challenges	<ul> <li>noise and crosstalk ≯ with # of qubits.</li> <li>electronics energetic cost.</li> <li>scaling cabling, circulators.</li> <li>scaling cryostats.</li> </ul>	<ul> <li>atom controls beyond 1,000 qubits.</li> <li>harder to implement gate-based QC and QND.</li> </ul>	<ul> <li>slow gate speed.</li> <li>scaling &gt;50 qubits.</li> <li>ion heating.</li> <li>control signals variability.</li> </ul>	<ul> <li>controlled electrostatic potential.</li> <li>error correction.</li> <li>qubits entanglement.</li> <li>fab cycle time.</li> </ul>	<ul> <li>creating entangled photons cluster states.</li> <li>photon losses in circuits.</li> <li>photon detectors efficiency.</li> </ul>
solutions	<ul> <li>materials improvement.</li> <li>3D chipset stacking.</li> <li>cryo-CMOS or SFQ.</li> <li>microwave signals multiplexing.</li> <li>more powerful cryostats, JJ circulators.</li> <li>scale-out with photons.</li> </ul>	<ul> <li>movable atoms for 2-qubit gates and readout.</li> <li>multiple elements.</li> <li>more powerful lasers.</li> <li>atoms/photon conversion for interconnect.</li> </ul>	<ul> <li>QCCD and microwave qubit drive.</li> <li>ions shuttling.</li> <li>multiple elements.</li> <li>Rydberg states ions.</li> <li>QPU photonic interconnect.</li> </ul>	<ul> <li>material and interfaces improvement.</li> <li>integrated cryoelectronics.</li> <li>more powerful cryostats.</li> <li>more efficient fabs (GF).</li> </ul>	<ul> <li>bright &amp; deterministic photon sources.</li> <li>deterministic sources of cluster states.</li> <li>MBQC / FBQC models.</li> <li>hybrid integrated nanophotonics.</li> <li>photon routing.</li> </ul>
caveats	<ul> <li>photonic interconnect overhead and statistics.</li> <li>energetic cost of microwave multiplexing and control drive.</li> <li>SFQ backaction on qubits.</li> </ul>	<ul> <li>gate control precision.</li> <li>losing the atom while computing.</li> <li>potential applicability limited to mid-scale simulations.</li> </ul>	<ul> <li>photonic interconnect viability.</li> <li>photonic interconnect statistics and impact on speedups.</li> </ul>	<ul> <li>understanding noise sources.</li> <li>scalability potential is capital intensive.</li> <li>2-qubit gates fidelities improving slowly.</li> </ul>	<ul> <li>photon statistics.</li> <li>small cluster states.</li> <li>nanophotonic circuit losses.</li> <li>probabilities everywhere.</li> </ul>

Figure 7a: Technological challenges by type of qubit for scaling up quantum computers. Source: Olivier Ezratty, 2022-2024.

What follows is a compilation of the public roadmaps of a large number of manufacturers in the sector, which are described in the following five sections for each type of qubit. These roadmaps are generally documented with one to three indicators including: the number of logical qubits, the number of executable quantum gates (which condition the fidelity of these logical qubits) and sometimes also the number of associated physical qubits. The median time for achieving a hundred or so logical qubits is around 2029, while some players, such as IBM, expect to reach a thousand logical qubits by 2033.

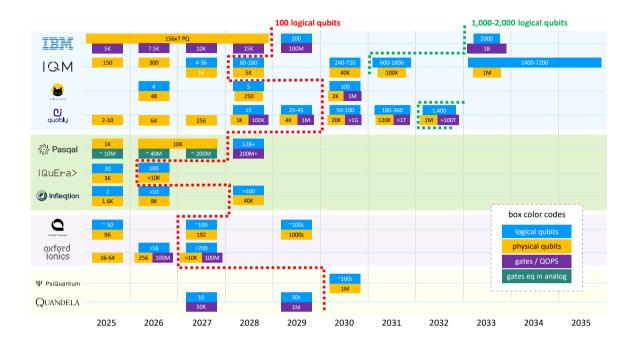


Figure 8: Compilation of the roadmaps of the main industrial vendors in quantum computing. (cc) Olivier Ezratty, May 2025.

At the end of the day, the big question is: how far can we go in controlling a large number of quantum entangled objects with good fidelity?

As far as energy consumption is concerned, researchers are currently investigating whether quantum computers could consume less energy than a conventional computer when performing useful computing tasks. This is an important question that could pave the way for many applications of quantum computing in a world looking to save energy.

However, the question remains entirely open, as the research community is only just beginning to develop ways of assessing and optimizing the energy consumption of small, noisy quantum computers (often referred to as NISQ computers), and the task is even more challenging for future fault-tolerant quantum computers.

The energy cost of quantum computers depends on the type of qubit, its quality and the cost of error correction. It is generally divided between qubit control electronics (microwave generators,

lasers, cables, amplifiers...), cryogenics and classical computing for control of the whole, both in terms of controlling the quantum computer with especially the classical part of error correction, but also with the classical part of quantum algorithms, which can be particularly costly. These costs increase significantly when we use logical qubits, which are sets of a potentially very large number of physical qubits.

There are two possible scenarios:

- Quantum computers solving problems which are also solvable with conventional computers with similar qualitative results.
- Fault-tolerant quantum computers, which will solve problems that are out of reach of existing and future conventional computers. They will then need to have a reasonable energy footprint, if possible not exceeding that of the most powerful supercomputers today (20 to 40 MW in power). Initial estimates in this area vary widely. 100 logic qubits at PsiQuantum (based on photons) could require power of the order of 10 to 20 MW. However, 4,000 logical qubits would be needed to perform useful calculations, for example in chemical simulation. At several hundred MW, the power required would become enormous and difficult to accept, especially as the corresponding calculations could take several months. As a comparison, CERN consumes a peak power of 200 MW, but only for a small part of its operating time. Other technologies could significantly reduce this bill, for example Alice&Bob's cat qubits, which could require far fewer resources to resist errors, Quandela thanks to more efficient sources of entangled photons, Quobly's spin qubits and, finally, certain types of trapped ions. Some people envisage supporting thousands of logic gubits with less than 1 MW, but all the hypotheses mentioned in this paragraph remain to be validated.

This work on estimating and optimizing, as well as standardization and *benchmarking of the energetic cost of quantum computers*, is addressed by the Quantum Energy Initiative<sup>33</sup>, an international community of academic and industry scientists working on the issue, which was launched in 2022.

#### 3.1. NEUTRAL ATOMS

Until recently, qubits based on neutral atoms were reserved for the specific field of quantum simulations. Since 2023, a number of scientific and technological advances have opened up the field of fault tolerance for neutral atoms.

Neutral atoms are usually trapped in a vacuum by lasers applying the Doppler effect and the Sisyphus effect<sup>34</sup>, the latter having been discovered by Claude Cohen-Tannoudji and Jean Dalibard at the Ecole Normale Supérieur Laboratoire Kastler Brossel (LKB) in 1987. Neutral atoms are prepared in several stages. A cloud of atoms is first trapped and cooled in a magneto-optical trap using coils that create a magnetic field and six lasers beams in three directions. Other lasers are then combined with light-shaping tools (SLMs, spatial light modulators that control the phase of the photons emitted by the lasers, and AODs, acousto-optical light deflectors that modify the orientation of the light beams) to create optical tweezers used to move and trap atoms at will in vacuum, and generally in two-dimensional geometric structures called lattices.

The two quantum states of atoms used to create qubits correspond to different energy levels, where the transitions are controlled by a variable mixture of laser beams and microwaves, in particular so-called high-energy Rydberg states. These states can be used as qubit states and/or to couple qubits together with

<sup>33</sup> Quantum Energy Initiative - Creating an interdisciplinary research line for quantum technologies energetics (quantum-energy-initiative.org)

<sup>34</sup> https://www.college-de-france. fr/sites/default/f iles/documents/serge-haroche/ UPL313866821288387325 Aperc u\_de\_Cours2015.II.B.20\_03.pdf

two-qubit gates or to configure a Hamiltonian for quantum simulations.

Neutral atom qubits can in fact be used in two ways: with computation based on one- and two-qubit quantum gates and for quantum simulations, using a prepared state of interconnected qubits that converge to a minimum energy level, helping to find a solution to N-body physics and various optimization problems. Quantum simulations can, for example, simulate Fermi-Hubbard models that model highly correlated electronic materials such as condensed matter and high-temperature superconducting materials. They can also be used to solve combinatorial problems by emnbedding QUBO (*Quadratic Unconstrained Binary Optimisation*) problems.

## neutral atoms qubits key takeaways

highlights         • operational systems with 100-300 atoms in simulation mode and not far from a provable quantum advantage.         • long qubit coherence time and fast gates.         • Identical atoms, that are controlled with the same laser and microwave frequencies.         • first logical qubits thanks to movable atoms.         • works in both simulation and gate-based paradigms.         • no need for specific integrated circuits.         • uses standard apparatus: lasers, optics, controls, cryogenics.         • low energy consumption.	challenges  crosstalk between qubits: can be mitigated with two-elements atom architectures.  slow operations: 1 Hz simulation cycle.  harder to implement with gate-based model: need a stable universal gate set with support for error correction.  need to move atoms around: to enable many-to-many two-qubit gates and reducing QEC overhead, enabling qLDPC, but with risk of losing the atoms on the way.  implement full QND measurement: needed for error correction.  losing atoms during computing while disconnecting the tweezers.
variations  dual species qubits: to improve QND measurement.  udeus spin qubits: with longer coherence times, in seconds.  circular Rydberg atoms: more stable.  fermionic computing: better for chemical simulations.  bosonic codes (cat-codes, GKP) with atom qubits: better QEC.  hybrid gate-based and analog: for specific case studies.  atoms trapped on nanophotonic circuits: enabling interactions.  hybrid neutral atoms and ions: getting the best of both worlds.  atom ensembles: studied particularly in China.	<ul> <li>path to scalability</li> <li>powerful lasers: needed to control &gt;300 atoms, with stabilized power, fiber lasers.</li> <li>atoms positioning: large scale SLM+AOD to trap up to 10K atoms.</li> <li>QND measurement: using dual species and/or atoms shuttling.</li> <li>faster operations and duty cycle: better and more powerful lasers, fast classical controls (FPGA, ASIC).</li> <li>implement full universal gate set : T gate &amp; magic state distillation.</li> <li>QPU interconnect: photon based, including for memories, can also use nanophotonic circuit traps.</li> </ul>

Figure 9: Strengths, challenges, varieties and scalability options of cold atom qubits.

#### 3.1.1. ADVANCES

In 2017, Mikhail Lukin from Harvard and a team from MIT controlled 51 rubidium atoms<sup>35</sup>. This record was raised to 256 atoms in July 2021<sup>36</sup>. In 2023, the same team operated two-qubit quantum gates (CZ) with 60 atoms with a fidelity of 99.5%. In France, Antoine Browaeys' team at the Institut d'Optique and Pasqal controlled 72 cold atoms in a 3D structure in 2018, 196 in 2020<sup>37</sup>, 500 in 2021 and828 in 2024<sup>38</sup>.

Scalability can be improved mainly by increasing the lifetime of atomic interactions. This also requires continuous improvement in the fidelity of one- and two-qubit gates, with a target of 99.9% fidelity for two-qubit gates. A record was set in 2022 by the University of Calgary, with fidelities of 99.85% for CZ gates and cesium atoms<sup>39</sup>.

One approach, also tested with trapped ions, involves mixing two species of neutral atoms. It is being evaluated by the University of Chicago, with 512 atoms in a lattice using an equivalent proportion of cesium and rubidium atoms. These species of atoms have different gate drive laser wavelengths<sub>40,41</sub>

Probing many-body dynamics on a 51-atom quantum simulator by Hannes Bernien, Mikhail Lukin et al., 2017 (24 pages).

<sup>36</sup> Quantum phases of matter on a 256-atom programmable quantum simulator by Sepehr Ebadi, Dolev Bluvstein, Vladan Vuletić, Mikhail D. Lukin et al, Nature, July 2021 (20 pages).

<sup>37 &</sup>lt;u>Synthetic three-dimensional atomic structures assembled atom by atom by Daniel Barredo,</u> Antoine Browaeys *et al*, 2018 (4 pages).

Rearrangement of single atoms in a 2000-site optical tweezers array at cryogenic temperatures by Grégoire Pichard, Adrien Signoles, Antoine Browaeys, Thierry Lahaye, Davide Dreon *et al, arXiv*, May 2024 (6 pages).

Interpretation of the second secon

<sup>40 &</sup>lt;u>Dual-Element, Two-Dimensional Atom Array with Continuous-Mode Operation</u> by Kevin Singh *et al*, University of Chicago, February 2022 (11 pages).

<sup>41 &</sup>lt;u>A dual-species Rydberg array</u> by Shraddha Anand, Conor E. Bradley, Ryan White, Vikram Ramesh, Kevin Singh, and Hannes Bernien, University of Chicago, *arXiv*, January 2024 (24 pages).

A great deal of research is underway to improve the capabilities and scalability of quantum simulations, by estimating the fidelity of quantum simulators based on neutral atoms<sup>42</sup>, measuring high-dimensional entanglement<sup>43</sup>, understanding the effect of noise and enhancing resilience to errors<sup>44, 45</sup>. Other work is investigating the ability to organize atoms in 3D<sup>46</sup>, with up to 30,000 atoms<sup>47</sup>.

Other advances deal with using Rydberg atoms with circular orbitals (CRA)<sup>48, 49</sup>.

Rydberg atoms with low angular momentum have large elliptical orbitals, creating a dipole. Their lifetime is limited to a few milliseconds and 1% of the atoms are lost during that time<sup>50</sup>.

Other researchers use ytterbium instead of rubidium, which is more stable, and strontium is also considered because of its longer stability times<sup>51</sup>.

<sup>42</sup> Preparing random states and benchmarking with many-body quantum chaos by Joonhee Choi and Manuel Endres, Caltech, MIT, Harvard, Berkeley, University of Illinois at Urbana-Champaign, University of Innsbruck, Nature, January 2023 (25 pages).

<sup>43 &</sup>lt;u>Detecting high-dimensional entanglement in cold-atom quantum simulators</u> by Niklas Euler and Martin Gärttner, Universität Heidelberg, May 2023 (26 pages).

<sup>44 &</sup>lt;u>Quantum advantage and stability to errors in analogue quantum simulators</u> by Rahul Trivedi, Adrian Franco Rubio and J. Ignacio Cirac, December 2022 (23 pages).

<sup>45</sup> Optimization of Algorithmic Errors in Analog Quantum Simulations by Nikita A. Zemlevskiy et al, August 2023 (18 pages).

<sup>46 &</sup>lt;u>Scalable multilayer architecture of assembled single-atom qubit arrays in a three-dimensional Talbot</u> <u>tweezer lattice</u> by Malte Schlosser, February 2019-March 2023 (13 pages).

<sup>47 &</sup>lt;u>Picosecond-Scale Ultrafast Many-Body Dynamics in an Ultracold Rydberg-Excited Atomic Mott Insulator</u> by V. Bharti *et al*, *PRL*, September 2023 (7 pages).

<sup>48 &</sup>lt;u>Array of Individual Circular Rydberg Atoms Trapped in Optical Tweezers</u> by Brice Ravon, Jean-Michel Raimond, Michel Brune, Clément Sayrin *et al*, April 2023 (9 pages).

<sup>49 &</sup>lt;u>Quantum Computing with Circular Rydberg Atoms</u> by Sam R. Cohen and Jeff D. Thompson, Princeton, *PRX Quantum*, March-August 2021 (26 pages).

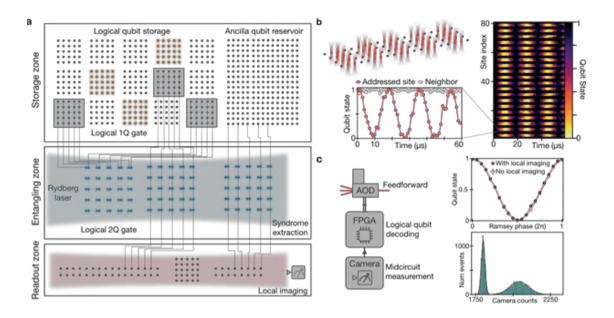
<sup>50</sup> Long-Lived Circular Rydberg Qubits of Alkaline-Earth Atoms in Optical Tweezers by C. Hölzl, A. Götzelmann, E. Pultinevicius, M. Wirth, and F. Meinert, Universität Stuttgart, *PRX*, May 2024 (11 pages).

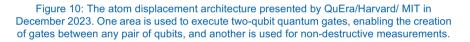
<sup>51 &</sup>lt;u>Cavity-enhanced optical lattices for scaling neutral atom quantum technologies to higher qubit numbers</u> by A. J. Park *et al*, November 2022 (18 pages).

In December 2023, QuEra, Harvard and MIT announced that they had created 48 logic qubits using 280 atoms. In their device, the atoms are shuttled in space for implementing two-qubit gates between any pair of qubits with Rydberg interactions and mid-circuit measurements in cavities<sup>52</sup>. Their two-qubit logic gate had an error of 7%, while the physical gates fidelities was 99.5%. They obtained good fidelity with 4-qubit GHZ entangled states but with using post-selection, which is a trick that cannot be used for scalable quantum computating. It's more an error detection technique than real error correction. They had not yet set up a complete set of universal gates. In the same *Nature* article, the researchers were experimenting with their own "quantum advantage" using an IQP (*Instantaneous Quantum Polynomial-Time*) algorithm using their 48-qubit logic qubits. It was demonstrated by IBM researchers that this could easily be simulated in the classical way, up to 96 logic qubits and even 192 logic qubits using Google TPUs<sup>53</sup>. Still, since the 2023 experiment, Harvard, MIT and Quera have made advances in the development of FTQC approach with their neutral atoms.

<sup>52</sup> Logical quantum processor based on reconfigurable atom arrays by Dolev Bluvstein, Mikhail D. Lukin *et al*, *Nature*, December 2023 (42 pages).

<sup>53</sup> Fast classical simulation of Harvard/QuEra IQP circuits by Dmitri Maslov, Sergey Bravyi, Felix Tripier, Andrii Maksymov, and Joe Latone, IBM and IonQ, arXiv, January 2024 (9 pages).





Neutral atoms have a number of advantages that are worth mentioning. They do not require specific and expensive semiconductor circuits, unlike all the other qubit modalities. They use relatively standard enabling technologies: lasers, discrete optical devices and even affordable cryogenics. The latest systems require affordable helium 4 based 4K cryogenics to cool both the ultra-high vacuum ion pump and the chamber containing the atoms in the magneto-optical trap (MOT). Lastly, these systems appear *to* be relatively energy-efficient. Several thousand atoms can *a priori* be controlled for a power budget of less than 20 kW.

#### 3.1.2. CHALLENGES

Here are some of the challenges to continue to develop these qubits and enable them to be implemented in fault-tolerant architectures.

Operation speed: each movement of atoms to create two-qubit gates takes 10 microseconds. Two-qubit quantum gates are therefore very slow, although this is compensated for by the ability to link all the qubits in the system in this way without using expensive SWAP gates. In analogue mode, a simulation cycle currently takes around one second. There are technological prospects for gaining one or two orders of magnitude in this cycle time.

Crosstalk between qubits: this can be reduced with two-element atomic architectures.

Complexity of the gate model: it is more difficult to implement than analogue quantum simulations. It requires a set of stable universal gates with support for error correction, including T-gates and magic states distillation.

Atoms shuttling: this is the technique currently chosen for the implementation of two-qubit gates with maximum connectivity, to reduce error correction overhead and allow the use of qLDPC correction codes, but with the risk of losing the atoms along the way and slower operations.

Atoms losses: during computation or atom shuttling.

#### 3.1.3. VARIATIONS

As with all qubit modalities, there are many variants of quantum computers based on neutral atoms.

Dual species qubits: which combine two elements with different energy transitions, such as the combination of cesium and rubidium. This makes it possible, for example, to implement a non-destructive qubit measurement, which is essential for error correction<sup>54</sup>.

Nucleus spin qubits: with longer coherence times, in seconds. This is the technological choice made by the Atom Computing (USA), which announced it could control a thousand qubits in 2023<sup>55</sup>.

Circular Rydberg atoms: which have good stabilities of the order of 30 ms, better than those of atoms with elliptical orbitals<sup>56,57</sup>. This technology is being studied at the Collège de France and the LKB, but seems to be better suited to quantum simulations than for gate-based computing.

Fermionic computation: this is particularly well suited to chemical simulations and to solving problems involving the electronic structures of molecules. A proposal from MIT combines pairs of cold atoms in 2D optical lattices in which the quantum information is encoded in the vibrational state of the pairs of atoms<sup>58</sup>. Also worth noting is the international work involving the LPMMC in Grenoble with Anna

<sup>54 &</sup>lt;u>Dual-Element, Two-Dimensional Atom Array with Continuous-Mode Operation</u> by Kevin Singh *et al*, University of Chicago, February 2022 (11 pages).

<sup>55 &</sup>lt;u>Assembly and coherent control of a register of nuclear spin qubits</u> by Katrina Barnes, Michael Yarwood *et al, Atom Computing*, August 2021 (11 pages).

<sup>56 &</sup>lt;u>High-fidelity gates with mid-circuit erasure conversion in a metastable neutral atom qubit by Shuo</u> Ma, Guido Pupillo, Shruti Puri, Jeff D. Thompson *et al*, *Nature*, May-October 2023 (17 pages).

<sup>57</sup> Long-Lived Circular Rydberg Qubits of Alkaline-Earth Atoms in Optical Tweezers by C. Hölzl,

A. Götzelmann, E. Pultinevicius, M. Wirth, and F. Meinert, Universität Stuttgart, *PRX*, May 2024 (11 pages).

<sup>3</sup> Quantum register of fermion pairs by Thomas Hartke, Botond Oreg, Ningyuan Jia and Martin Zwierlein, MIT, 2021 (10 pages).

Minguzzi and Benoit Vermersch<sup>59</sup>. Here again, we are more in the realm of quantum simulation than FTQC.

Cat-codes: there is a way of implementing *cat-codes* with Rydberg atoms, the *cat-manifold* variety exploiting atom nucleus spins<sup>60</sup>.

Hybrid gate-based and analogue computing: a mixture of analogue and gatebased circuits<sup>61</sup>.

Atom ensembles: these are being studied in particular in China. These are sets of atoms or "super-atoms", both for gate computing and for quantum simulations<sup>62</sup>.

Hybridization of atoms and ions: in architectures combining cold atoms and trapped ions<sup>63</sup>.

#### 3.1.4. PATHWAYS TO SCALABILITY

Implementing fault-tolerant computing with neutral atoms involves a number of scientific and technological challenges:

Fermionic quantum processing with programmable neutral atom arrays by Daniel González-Cuadra, Dolev Bluvstein, Marcin Kalinowski, Raphael Kaubruegger, Nishad Maskara, Piero Naldesi, Torsten V. Zache, Adam M. Kaufman, Mikhail D. Lukin, Hannes Pichler, Benoît Vermersch, Jun Ye, and Peter Zoller, arXiv, March 2023 (13 pages).

<sup>60</sup> Fault-tolerant quantum computation using large spin cat-codes by Sivaprasad Omanakuttan, Vikas Buchemmavari, Jonathan A. Gross, Ivan H Deutsch, and Milad Marvian, University of New Mexico, Google AI, arXiv, January 2024 (22 pages).

<sup>61</sup> Quantum simulation of fermionic systems using hybrid digital-analog quantum computing approach by Nikita Guseynov and Walter Pogosov, December 2021-May 2022 (20 pages).

<sup>62 &</sup>lt;u>Rydberg superatoms: An artificial quantum system for quantum information processing and quantum optics</u> by Xiao-Qiang Shao, Shi-Lei Su, Lin Li, Rejish Nath, Jin-Hui Wu, and Weibin Li, *arXiv*, April 2024 (32 pages).

<sup>8</sup> Performing Non-Local Phase Estimation with a Rydberg-Superconducting Qubit Hybrid by Juan Carlos Boschero, TU Eindhoven, *arXiv*, February 2024 (98 pages).

Gate fidelity: we need to maintain a good quality of two-qubit quantum gates, approaching 99.9% and on a large scale. All this without losing the atoms as they move through the zones dedicated to the implementation of these quantum gates.

Atoms positioning: the challenge is to be able to control and move a large number of neutral atoms in space using optical tweezers. The preferred technique is to use SLM- and AOD-based optical tweezers on a large scale to trap up to several thousand atoms. Caltech's record, announced in March 2024, of controlling the location of 6,100 atoms with optical tweezers is worth noticing<sup>64</sup>.

Laser power: as the number of controlled atoms increases beyond a few hundreds, the lasers positioning and controlling the atoms need to be more powerful and phase stabilized. The technology being considered is fibre lasers. The related commercial offer is still limited.

Non-demolition measurement: this is necessary to implement error correction. The solutions envisaged involve the use of two species of atoms or the displacement of atoms in dedicated measurement zones in cavities.

Faster operations cycles: current quantum gates, particularly those with two qubits, are too slow, in the order of a millisecond. The same applies to the duration of a complete quantum simulation cycle, which lasts around one second. Shortening these times will require the use of more powerful, fibre lasers and optimizing conventional control electronics.

Positioning atoms by ionizing them beforehand and using electrodes: this is a technique envisaged by Atom Quantum Lab in Slovenia.

<sup>64 &</sup>lt;u>A tweezer array with 6100 highly coherent atomic qubits</u> by Hannah J. Manetsch, Gyohei Nomura, Elie Bataille, Kon H. Leung, Xudong Lv, and Manuel Endres, Caltech, *arXiv*, March 2024 (21 pages).

Interconnection between processors: although it will be possible to control several thousand atoms in a single processor to create a few dozen logic qubits, they will need to be quantum interconnected in order to scale up. The connection envisaged uses optical photons and quantum memories such as those developed by Welinq. The advantage of this type of connection is that optical photons can also be used to interact with atoms. All that remains is to ensure that the wavelengths are compatible between these different systems and to ensure that the distant quantum gates are of good quality.

#### 3.1.5. ROADMAPS

The main startups using neutral atoms are Infleqtion (formerly ColdQuanta, 2007, USA, initially working on quantum sensors), Pasqal (2019, France), QuEra (2020, USA, linked to Mikhael Lukin and Harvard), Atom Computing (2018, USA), and PlanQC (2022, Germany).

In January 2024, QuEra announced its FTQC roadmap, planning to achieve 100 logic qubits by 2027 using more than 10,000 physical qubits, and with a logical error rate of 10<sup>-6</sup> to 10<sup>-8</sup>. Support for non-Clifford (T) quantum gates is expected in 2025.

In February 2024, Infleqtion announced its 5-year roadmap. It presented its Sqorpius QPU controlling 1,600 neutral cesium atoms arranged in a  $40 \times 40$  array, with no indication of gate fidelities. It plans to support 100 logical qubits with 40,000 atoms with physical two-qubit gate fidelities of 99.95% by 2028. It will do this using two elements, rubidium and cesium<sup>65</sup> and plans to use qLDPCtype error correction codes<sup>66</sup>.

Infleqtion Unveils 5-Year Quantum Computing Roadmap, Advancing Plans To Commercialize Quantum At Scale by Infleqtion, February 2024.

<sup>66</sup> Architecture for fast implementation of qLDPC codes with optimized Rydberg gates by C. Poole, T. M. Graham, M. A. Perlin, M. Otten, and M. Saffman, Infleqtion and University of Wisconsin-Madison, arXiv, April 2024 (12 pages).

In March 2024, Pasqal announced its new roadmap, paving the way for quantum gate computing in the FTQC regime by 2028<sup>67,68</sup>. It also plans to support a hybrid analogue/digital approach<sup>69</sup>. It positions its analogue simulators using an equivalent number of quantum gates.

Recent major scientific milestones for Pasqal and the IOGS research team include the 2D control of 196 qubits<sup>70</sup>, the execution of parallel gates<sup>71</sup> and the control of 324 cold atoms in a 2D lattice, with the use of a 4 K cryostat cooling the ultra-high vacuum pump, which creates a better vacuum and extends the lifetime of the atomic qubits, as well as the vacuum chamber, to avoid the effect of electromagnetic radiation<sup>72,73</sup>. In May 2024, Pasqal and IOGS broke a record by manipulating 828 atoms in 2,000 sites in their Humber QPU using optical tweezers and two lasers using slightly different wavelengths<sup>74</sup>.

<sup>67</sup> PASQAL Announces New Roadmap Focused on Business Utility and Scaling Beyond 1,000 Qubits Towards Fault Tolerance Era, Pasqal, March 2024.

<sup>8</sup> Roadmap to Quantum Readiness with a Full Stack Approach & Transformative Use Cases by Pasqal, April 2024.

Microwave Engineering of Programmable XXZ Hamiltonians in Arrays of Rydberg Atoms by P. Scholl, Loic Henriet, Thierry Lahaye, Antoine Browaeys et al, PRX, April 2022 (10 pages).

<sup>70 &</sup>lt;u>Enhanced atom-by-atom assembly of arbitrary tweezer arrays</u> by Kai-Niklas Schymik, Antoine Browaeys, Thierry Lahaye *et al*, November 2020 (10 pages).

<sup>71 &</sup>lt;u>Pulse-level Scheduling of Quantum Circuits for Neutral-Atom Devices</u> by Richard Bing-Shiun Tsai, Loic Henriet *et al*, Pasqal, June 2022 (8 pages).

<sup>72</sup> Single Atoms with 6000-Second Trapping Lifetimes in Optical-Tweezer Arrays at Cryogenic <u>Temperatures</u> by Kai-Niklas Schymik, Sara Pancaldi, Florence Nogrette, Daniel Barredo, Julien Paris, Antoine Browaeys, and Thierry Lahaye, *IOGS, PRA*, September 2021 (8 pages).

<sup>&</sup>lt;sup>73</sup> In situ equalization of single-atom loading in large-scale optical tweezer arrays by Kai-Niklas Schymik, Adrien Signoles, Florence Nogrette, Daniel Barredo, Antoine Browaeys, and Thierry Lahaye, PRA, August 2022 (6 pages).

<sup>74 &</sup>lt;u>Rearrangement of single atoms in a 2000-site optical tweezers array at cryogenic temperatures by</u> Grégoire Pichard, Adrien Signoles, Antoine Browaeys, Thierry Lahaye, Davide Dreon *et al, arXiv*, May 2024 (6 pages).

## 3.1.6. CONCLUSION

Fault-tolerant computers based on cold atoms are promising, with the potential capacity to deliver a hundred logical qubits by the end of the decade, paving the way for a possible quantum advantage. As with all qubit modalities, the question is how far it will be possible to scale using a single processor in terms of reliable atom control. Beyond that, photonic interconnect will be essential. It remains difficult to develop, especially on a large scale, and particularly in terms of its ability to enable the execution of reliable two-qubit quantum gates between two processors connected in this way. This is essential if we are to consider creating distributed fault-tolerant computers with logical qubits distributed across several processors that can communicate with each other.

# 3.2. SUPERCONDUCTING QUBITS

Superconducting qubits are currently the most studied modality in the academic world and have the largest and best-funded community of commercial suppliers, including major players such as IBM, Google, IQM and Rigetti. By misnomer, this category also includes bosonic qubits based on entangled microwave photons trapped in cavities, with players such as Alice&Bob (*cat-qubits*) and Nord Quantique (GKP qubits).

The most common superconducting qubits are based on the use of anharmonic oscillators of superconducting currents passing through one or more Josephson junctions. The two lowest energy levels of these oscillators encode the  $|0\rangle$  and  $|1\rangle$  states of each qubit. These oscillators operate at a frequency of around 10 MHz and are controlled by microwave pulses of around 5 GHz, which generate single-qubit gates. Two-qubit gates are usually created by controlling tunable couplers linking adjacent qubits in circuits. Qubits readout use a resonator located close to the qubit. It is excited by a microwave pulse that is reflected, amplified and analyzed. The combination of its phase and amplitude is used to determine the qubit collapsed state. This measurement is non-destructive, because the qubit remains operational after its measurement, in the state generated, which is either  $|0\rangle$  or  $|1\rangle$ .

#### superconducting qubits key takeaways

highlights	challenges
<ul> <li>key technology in the academic world and with the largest and well funded commercial vendors community including large players like IBM, Google, IQM and Rigetti.</li> <li>record of 156 programmable qubits with 99.7% two qubit gates fidelities (IBM Heron r2, 2024).</li> <li>noise reduction constant progress in regular transmons with tunable couplers and with bosonic qubits which could enable a record low ratio of physical/logical qubits.</li> <li>first break-even logical qubits: Google in August 2023.</li> <li>enabling technologies are abundant with cryostats, cabling, analog electronics, amplifiers, and sensors.</li> <li>quantum error mitigation and quantum error correction known techniques to enable NISQ applications and future FTQC designs.</li> </ul>	<ul> <li>qubits variability: requires calibration and complex micro-wave frequency maps and need to contain crosstalk.</li> <li>qubit connectivity: limited to neighbor qubits in 2D structures.</li> <li>qubit coherence time: usually &lt; 300 µs with some lab records &gt;1ms.</li> <li>cryogeny: constrained technology at &lt;15 mK, but not a scientific obstacle per se, more an engineering one. Yield can be improved.</li> <li>logical qubits: are not yet under break even. They are currently worse than physical qubits in most experiments.</li> <li>cabling clutter: complexity and many passive and active electronic components to control qubits with micro-waves and other signals.</li> <li>qubit size: uneasy miniaturization limits qubit # per chips and requires QPU interconnect solutions.</li> <li>qubit fidelities: have a hard time reaching 99.9%, needed for QEC.</li> </ul>
variations	path to scalability
<ul> <li>bosonic qubits: cat-qubits, GKP, dual rail with self-correction and</li> </ul>	materials: improve elements purity, identify other promising ones.
lower QEC overhead. They are less mature but promising.	EDA: full stack electronic design automation tools.
<ul> <li>fluxonium qubits: with better fidelities but a more complicated designs and few involved vendors.</li> </ul>	<ul> <li>manufacturing: industrialization, 300 mm wafers epixaxial deposition.</li> <li>qubit mid-range coupling: to enable lower overhead qLDPC QEC.</li> <li>interconnect: using transduction and photonic gate teleportation.</li> <li>signals multiplexing or SFQ control electronics: to reduce cabling overhead and cryogenic requirements.</li> <li>QEC syndrome detection speed improvements using ASICs.</li> <li>cryostats: larger and more efficient.</li> </ul>
<ul> <li>qutrits: with larger Hilbert space, which are exotic in the commercial world.</li> </ul>	
<ul> <li>Andreev spin qubits: localized excitation of the BCS condensate that natively has only two levels, using a nanowire.</li> </ul>	
<ul> <li>hybrid quantum analog-digital architectures: to solve specific problems, not generic.</li> </ul>	

# Figure 11: Strengths, challenges, varieties and scalability options for superconducting qubits.

#### 3.2.1. ADVANCES

The first superconducting qubits appeared in the laboratory in 1999 in Japan<sup>75</sup> and then in France and the USA in the early 2000s, in particular with the creation of circuit electrodynamics (cQED) at Yale University in 2004, which led to the development of transmon qubits<sup>76</sup>.

<sup>75 &</sup>lt;u>Coherent control of macroscopic quantum states in a single-Cooper-pair box</u> by Yasunobu Nakamura, Yuri Pashkin and Jaw-Shen Tsai, *Nature*, 1999 (4 pages).

<sup>76</sup> Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for guantum computation by Alexandre Blais, Ren-Shou Huang, Andreas Wallraff, Steve Girvin and Rob. Schoelkopf, *PRA*, 2004 (14 pages) and <u>Strong coupling of a single photon to a</u> <u>superconducting qubit using circuit quantum electrodynamics</u> by Andreas Wallraff, David Schuster, Alexandre Blais, Steve Girvin, Rob Schoelkopf et al, Nature, 2004 (7 pages).

The first commercial processor of this type was brought online by IBM in 2016, with 5 qubits.

In 2019, Google made headlines with its 53-qubit Sycamore processor, with a quantum supremacy experiment corresponding to a quantum circuit sampling running in less than 3 minutes that was supposed to cost 10,000 years and then 2.5 days<sup>77</sup>.on a classical supercomputer. This type of quantum computing do not use any input data and only generated a correct result in 0.14% of the circuit runs. Four years later, this circuit sampling task was digitally simulated in 6 seconds on the US Department of Energy's Frontier supercomputer using tensor networks. Nevertheless, Google's experiment, which has since been repeated with 72 qubits of slightly improved quality, was still a technological feat. It would take 46 days to simulate it using tensor networks on the Frontier supercomputer, and it has not yet been done conventionally<sup>78</sup>.

In August 2024, Google produced its first logical qubit with higher fidelities than its physical qubits and on two processors, one with 72 qubits and the other with 105 qubits, called Willow in December 2024. Google was able to correct the error of a logical qubit in *surface* distance *code* 5 on the first in real time and on the second, with a distance code 7 and 101 qubits, a record for superconducting qubits. However, they had not yet corrected quantum logic gates on these processors.

Since 2019, developments in this type of qubit have continued. The number of operational qubits has increased reasonably, but not exponentially. Their quality has improved, although less spectacularly than that of trapped ions. Transmon qubits using tunable couplers achieved a record fidelity of 99.7%.

Quantum supremacy using a programmable superconducting processor by Frank Arute, John Martinis et al, October 2019 (12 pages) and <u>Supplementary information for "Quantum supremacy</u> <u>using a programmable superconducting processor"</u> by Frank Arute, John Martinis et al, October 2019 (58 pages)

<sup>78</sup> Phase transition in Random Circuit Sampling by A. Morvan et al, Google AI, April-December 2023 (45 pages).

for two-qubit (CZ) gates on the 133 qubits of IBM's Heron Torino processor in December 2023 and later in 2024 and 2025 on the Heron r1, r2 and r3 156-qubit chips<sup>79</sup>.

IBM and its partners then published a large number of scientific works relating to "quantum utility" with the ability to solve problems that would not be accessible to conventional computers, even by exploiting tensor networks<sup>80</sup>. However, several researchers succeeded in solving these problems using tensor networks<sup>81, 82, 83</sup>. However, many of these use cases do not yet correspond to the needs of companies, but to physical simulations in the context of fundamental research<sup>84, 85</sup>. IQM (IQM) has achieved gate fidelity of 99.5% in 2024 on 20 qubits. In China, Origin Quantum was further behind, with 97% two-qubit gate (CZ) fidelity achieved on 6 selected qubits in a 72-qubit processor<sup>86</sup>. Other players in this market, such as OQC, Rigetti and Anyon Systems, achieved even worse results.

Qubit fidelity is usually measured at several points: qubit initialization, one-qubit gates, two-qubit gates and qubit reading. In practice, we concentrate on the fidelity of two-qubit gates, which are the most difficult to obtain and which determine the quality of entanglement between qubits and, therefore, the efficiency of quantum algorithms. Other important figures of merit are the duration of these operations. Two qubit gates and reading are generally the slowest operations and various techniques have been proposed to reduce their duration.

Evidence for the utility of quantum computing before fault tolerance by Youngseok Kim, Kristan Temme, Abhinav Kandala *et al*, IBM Research, RIKEN iTHEMS, University of Berkeley and the Lawrence Berkeley National Laboratory, *Nature*, June 2023 (8 pages).

<sup>81 &</sup>lt;u>Efficient tensor network simulation of IBM's Eagle kicked Ising experiment</u> by Joseph Tindall, Matt Fishman, Miles Stoudenmire and Dries Sels, *PRX Quantum*, June 2023-January 2024 (16 pages).

<sup>82</sup> Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance by Tomislav Begušić et al, Caltech, August 2023 (17 pages).

Effective quantum volume, fidelity and computational cost of noisy quantum processing experiments by K. Kechedzhi *et al*, Google Al, NASA, June 2023 (15 pages).

Realizing the Nishimori transition across the error threshold for constant-depth quantum circuits by Edward H. Chen, Sarah Sheldon, Simon Trebst, Abhinav Kandala *et al*, September 2023 (16 pages).

<sup>85</sup> Uncovering Local Integrability in Quantum Many-Body Dynamics by Oles Shtanko, Derek S. Wang, Haimeng Zhang, Nikhil Harle, Alireza Seif, Ramis Movassagh, and Zlatko Minev, July 2023 (8 pages) on spin lattices simulation with 124 qubits.

<sup>8</sup> Demonstrating a universal logical gate set on a superconducting quantum processor by Jiaxuan Zhang, Guo-Ping Guo *et al*, Origin Quantum, CAS *et al*, *arXiv*, May 2024 (15 pages).

Bosonic qubits are usually classified as superconducting qubits. They share some of the same characteristics, although in practice the qubit uses microwave photons trapped in cavities and in superposed states that protect them against certain types of error. These are called biased error qubits. In this field, cat qubits reduce one type of error (*flip*) at the expense of moderately increasing another type of error (*phase*). They achieve stability times of several seconds<sup>87</sup>. These qubits make it possible to create logic qubits with a lower ratio of physical qubits per logical qubit, of the order of 15 to 20 to obtain 100 logical qubits and a logical error rate of the order of 10<sup>-8</sup> to 10<sup>-6</sup>, according to Alice&Bob<sup>88</sup>. In September 2024, Nord Quantique published initial results on its GKP bosonic qubits, correcting automously both flip and phase noise, with a proposed architecture for scaling up<sup>89</sup>. Amazon followed suit with a demonstration of a 5-qubit cat chip combined with four transmon auxiliary qubits using error correction to create one- and two-qubit logic gates. However, the fidelities obtained were no better than those of the best physical transmon qubits<sup>90</sup>.

Enabling technologies covering the needs of superconducting qubits are abundant, with cryostats and their dilution systems, superconducting or flexible ribbon cables, analogue and hybrid digital-analogue electronics, lowtemperature amplifiers and sensors. Each of these technologies has progressed in the last five years. Cryogenics is increasingly powerful and efficient. Control electronics are denser, starting with those that operate at ambient temperature. As a result,

<sup>87</sup> Quantum control of a cat-qubit with bit-flip times exceeding ten seconds by Ulysse Réglade, Pierre Rouchon, Alain Sarlette, Mazyar Mirrahimi, Philippe Campagne-Ibarcq, Raphaël Lescanne, Zaki Leghtas *et al.*, July 2023 (17 pages).

<sup>8</sup> LDPC-cat codes for low-overhead quantum computing in 2D by Diego Ruiz, Jérémie Guillaud, Anthony Leverrier, Mazyar Mirrahimi, and Christophe Vuillot, *arXiv*, January 2024 (23 pages).

<sup>8 &</sup>lt;u>Hardware-Efficient Fault Tolerant Quantum Computing with Bosonic Grid States in</u> <u>Superconducting Circuits</u> by Marc-Antoine Lemonde, Philippe St-Jean *et al, Nord Quantique, arXiv*, September 2024 (17 pages).

<sup>9 &</sup>lt;u>Hardware-efficient quantum error correction using concatenated bosonic qubits</u> by Harald Putterman, Fernando G. S. L. Brandão, Oskar Painter *et al*, AWS, *arXiv*, September 2024 (60 pages).

In 2022, IBM succeeded in integrating the control electronics for 433 qubits in a single rack. This record was extended to 1,121 qubits with the Condor processor in December 2023. The systems were subsequently dismantled because the associated chips had two- qubit gates errors over 2%. The amplifiers that operate at the quantum limit (TWPA) also improved in terms of efficiency and bandwidth, as has the integration of cabling, notably with flexible cables from Delft Circuits in the Netherlands.

Finally, in the NISQ regime, error mitigation techniques have been developed to achieve quantum utility, which remains limited to a few specialized use cases. They are relevant when implemented with at least fifty physical qubits. Below that, quantum algorithms can generally be emulated more efficiently in terms of both time and energy.

#### 3.2.2. CHALLENGES

Numerous challenges remain in order to continue to develop these qubits and, in particular, to enable them to be implemented in fault-tolerant architectures.

Qubit fidelity for two-qubit gates has not yet reached 99.9%. This level of fidelity is considered necessary for error correction, whether based on surface codes, their derivatives or even qLDPC codes inspired by classic LDPC-type parity detection codes. Researchers fight a large number of sources of qubit noise and decoherence, including the influence of external circuits connected to the qubit (resonators, filters...), the charge noise resulting from the movement of charges in dielectrics and oxides, critical current noise arising from charge motion in the Josephson junction, paramagnetic spin fluctuations in the superconducting loop insulator interface, tunneling poisoning of single quasiparticles inside and outside the superconducting island, dielectric losses of substrates due to tunneling, various other tunneling effects in amorphous materials, crosstalk between qubits

generated by their control signals and their routing, the effects of cosmic rays generating correlated errors and, finally, disturbances from the low frequencies generated by the compressors and pulse tubes that cool the cryostat to 4 K.

The coherence times of qubits are generally fairly short, in the order of 100 to 300  $\mu$ s, with some laboratory records exceeding one millisecond, but generally for the fluxonium variant of these qubits<sup>91</sup>. In contrast, Origin Quantum (China) has a T<sub>2</sub> (phase stability) of just 2  $\mu$ s<sup>92</sup>.

*Leakage* errors and correlated errors caused by cosmic rays and ambient radioactivity are not dealt with by the error correction codes envisaged for fault tolerance and require special treatment. These errors are difficult to avoid and correct, but solutions have been proposed, notably by Google and Michel Devoret, based on modifications to the geometry of Josephson junctions in qubits, a technique labelled gap engineering<sup>93, 94</sup>.

The first logical qubit experiments were carried out by ETH Zurich<sup>95</sup>, Google<sup>96</sup>, IBM<sup>97</sup> and Origin Quantum. They have not yet achieved "*break even*" with gate operations. These logical qubits are currently less reliable than their underlying physical qubits. The error rates of two-qubit logical gates were higher than those of physical gates until Google Willow's achievement in 2024 with a logical memory qubit.

Image: Millisecond coherence in a superconducting qubitMillisecond Somoroff, Vladimir Manucharyan et al,<br/>University of Maryland, PRL, March 2021-June 2023 (14 pages).

<sup>92</sup> Enabling Large-Scale and High-Precision Fluid Simulations on Near-Term Quantum Computers by Zhao-Yun Chen, Guo-Ping Guo et al, Origin Quantum, arXiv June 2024 (31 pages).

Direct evidence for cosmic-ray-induced correlated errors in superconducting qubit array by Xue-Gang Li *et al*, *arXiv*, February 2024 (7 pages).

Resisting high-energy impact events through gap engineering in superconducting qubit arrays by Matt McEwen, Michel Devoret, Alex Opremcak *et al*, *arXiv*, February 2024 (14 pages).

Realizing repeated quantum error correction in a distance-three surface code by Sebastian Krinner, Alexandre Blais, Andreas Wallraff et al, Nature, May 2022 (28 pages).

Suppressing quantum errors by scaling a surface code logical qubit by Rajeev Acharya et al, Google AI, Nature, July 2022-February 2023 (44 pages).

<sup>97</sup> Encoding a magic state with beyond break-even fidelity by Riddhi S. Gupta, Benjamin J. Brown et al, IBM, Nature, January 2024 (24 pages).

The qubits variability requires calibration and complex frequency plans for microwave control, to limit the effects of crosstalk. This calibration is frequent and costly. Its complexity can increase with the number of qubits used.

Qubit connectivity is currently limited to nearest neighbour qubits in 2D circuit structures, resulting in a high rate of redundancy required for error correction, the highest being in surface codes where thousands of qubits are needed to create logic qubits with the fidelity required by applications such as chemical simulation.

The size of qubits remains large, at least a mm<sup>2</sup>. This makes it difficult to miniaturize circuits and limits the number of qubits per chip, especially as the current trend with transmons is to link them together using tunable couplers, which requires more space. This means that we need to use interconnection techniques between chips, some of which are still highly uncertain as to their feasibility. IBM plans to create chips using just 156 physical qubits and then assemble them using various interconnection techniques with short to mid-range microwave links and then, later, optical photons transduction from/to microwave photons and photonic interconnect.

Qubits cabling is bulky, with two to five coaxial cables per physical qubit. This fills the entire interior of the cryostats. Cryostats are also filled with many passive and active electronic components used to control the qubits with microwaves and for their state readout. This applies in particular to the amplifiers (TWPA at 15 mK, HEMT at 4 K) and circulators (at 15 mK), which are all very bulky, and have a limited capacity for multiplexing the simultaneous reading of a maximum dozen qubits, due to the frequency multiplexing used and the bandwidth required per qubit, which is between 200 and 400 MHz. Control electronics operating at ambient temperature consume a lot of energy, currently at around 100 W per physical qubit, which is unacceptable on a large scale with millions of physical qubits. The cryo-CMOS low-temperature electronics techniques currently being considered displace the problem at the cryostat dilution level. They reduce the amount of wiring entering the cryostat, but, as these circuits heat up, they are currently positioned at 4 K where the cooling budget is greater (currently 1 W) than at 15 mK (around 20 µW) and we need to find a way of linking these 4 K components to quantum chips operating at 15 mK.

Cryogenics is often highlighted as a limiting factor. Superconducting qubits are those that require active cryogenics capable of reaching temperatures in the mK range. As seen above, the associated cryostats using dry dilutions have cooling powers of limited to a few tens of  $\mu$ W at 15 mK and a few W at 4 K. This limits the ability to use low-temperature electronics.

#### 3.2.3. VARIATIONS

As with every type of qubit, superconducting qubits are very diverse, with a large number of design proposals that require in-depth analysis.

Fluxonium qubits, which are variants of flux qubits, currently have better fidelities, but are based on more complicated designs that have not yet been tested on a large scale, and with few commercial players in the field (Atlantic Quantum in the USA). These qubits are being studied in many places (in the USA, Japan, China, and in France). Alibaba had obtained very good results<sup>98</sup>, but abandoned its quantum computing business in 2023. It was transferred to Zhejiang University, which is continuing research in the field<sup>99</sup>.

Qutrits use a larger Hilbert encoding space with three energy levels for the qubit's anharmonic oscillator. The technique is not commercially available<sup>100</sup>. Qutrits could make it possible to

Mative approach to controlled-Z gates in inductively coupled fluxonium gubits by Xizheng Ma et al, Alibaba, August 2023 (19 pages).

Achieving millisecond coherence fluxonium through overlap Josephson junctions by Fei Wang, Hui-Hai Zhao, Chunqing Deng *et al*, Zhejiang University and Z-Axis Quantum, *arXiv*, May 2024 (12 pages).

<sup>100</sup> High-fidelity qutrit entangling gates for superconducting circuits by Noah Goss, Irfan Siddiqi et al, Nature Communications, November 2022 (6 pages).

limit the effects of noise in the NISQ regime, but the theoretical foundations for implementing them in the FTQC regime with error correction are not yet in place.

Andreev spin qubits use localized excitation of a BCS condensate (Bardeen-Cooper-Schrieffer), which has only two levels, using a nanowire<sup>101</sup>. This is an exploratory laboratory technique, with no commercial players. We are therefore still a long way from FTQC.

Bosonic qubits incorporate total or partial autonomous error correction at the physical level. There are several types, such as cat *qubits* and GKP (Gottesman-Kitaev-Preskill) qubits. These qubits still require error correction, but this is easier to implement. The first physical qubits with flip physical qubit error correction were produced by Alice&Bob in late 2023 and early 2024. Nord Quantique announced a GKP logical qubit in early 2024, but its performance has not been fully documented<sup>102</sup>. Amazon, for its part, is working on several parallel avenues: cat qubits, GKP qubits<sup>103</sup> and dual-rail qubits<sup>104</sup>.

Hybrid analogue-digital quantum architectures are also proposed to solve specific problems. They are not generic and are dedicated to the NISQ regime. Few players offer them, such as Kipu Quantum in Germany, which is a quantum software player.

<sup>101 &</sup>lt;u>Circuit-QED with phase-biased Josephson weak links</u> by C. Metzger, Christian Urbina, Hugues Pothier *et al*, January 2021 (22 pages).

<sup>102 &</sup>lt;u>Autonomous quantum error correction of Gottesman-Kitaev-Preskill states</u> by Dany Lachance- Quirion, Philippe St-Jean *et al*, *Nord Quantique*, October 2023 (32 pages).

<sup>103</sup> Low overhead fault-tolerant quantum error correction with the surface-GKP code by Kyungjoo Noh, Christopher Chamberland, and Fernando G. S. L. Brandão, arXiv, March 2021-January 2022 (41 pages).

<sup>104 &</sup>lt;u>Demonstrating a Long-Coherence Dual-Rail Erasure Qubit Using Tunable Transmons</u> by H. Levine *et al*, The Hebrew University of Jerusalem, University of Chicago, Caltech and AWS, *Physical Review X*, March 2024 (21 pages).

#### 3.2.4. PATHWAYS TO SCALABILITY

Here is an inventory of the techniques envisioned to facilitate the scaling of these qubits and their implementation with fault tolerance support.

Improving manufacturing processes by setting up pre-industrial and industry quality production lines, using 300 mm *wafers* and epitaxial deposition techniques. This would make it possible to reduce the variability of qubits, reduce material deposition defects and improve manufacturing yields. These techniques are currently being studied by IMEC in Belgium, by Qolab with Applied Materials, and are planned by IQM and Alice&Bob as part of the creation of a semi-industrial production line operated by CEA-Leti in Grenoble. IMEC's initial experiments are not yet conclusive, but progress is possible<sup>105</sup>.

It is also possible to improve the purity of elements to avoid structural defects in circuits. This is an area of study by a research team in Slovenia working with Google<sup>106</sup>. Various alternative elements, such as tantalum, have been studied for the creation of circuits, particularly at resonator level, but in the end the results are no better than with niobium. Josephson junctions are always made with a sandwich of aluminum surrounding aluminum oxide.

The creation of '*full-stack*' and 'multi-physics' digital simulation tools for qubits will make it possible to better simulate the operation of circuits and optimize them. Alice&Bob is developing such a software chain as part of the i-Démo "Cat Factory" project.

<sup>105</sup> High-coherence superconducting qubits made using industry-standard, advanced semiconductor manufacturing by Jacques Van Damme, Kristiaan De Greve *et al*, IMEC, *arXiv*, March-April 2024 (17 pages).

Slovenian physicist helps unlock quantum computer utility in breakthrough finding, The Slovenia

 Times, April 2024.

Medium-distance coupling of qubits to enable the implementation of qLDPC- type error correction codes. This technique is currently being studied by IBM<sup>107</sup>, IQM<sup>108</sup>, IMEC<sup>109</sup> and MIT. Alice&Bob's cat qubits use conventional 1D qLDPC codes for phase error correction, with the plan to support up to 100 logical qubits<sup>110</sup>. Beyond that, in order to correct the residual flip error, 2D error correction codes will be required, which will entail the need to improve connectivity between qubits.

The simplification of low-temperature control electronics is being studied in a number of areas. Multiplexing qubit control signals via conversion into optical photons is being studied by Viqthor (France). Cryo-CMOS electronics are being considered by a number of players, including IQM, Google and IBM. Superconducting control electronics (SFQ, for *single-flux quantum*) are very attractive because of their low thermal load. It is proposed by SeeQC (USA) and studied in various academic laboratories in Finland (VTT), France (at Pascal Febvre's CNRS laboratory in Chambéry) and Japan (RIKEN). Finally, techniques for replacing ferrite-based circulators with superconducting circuits are also being considered, notably by Analog Quantum Circuits<sup>111</sup> Australia) and Google<sup>112</sup>.

Improving the speed of error syndrome detection for error correction will require the use of ASICs instead of FPGAs. This replacement will be justified from an

<sup>107 &</sup>lt;u>High-threshold and low-overhead fault-tolerant quantum memory</u> by Sergey Bravyi, Andrew W. Cross, Jay M. Gambetta, Dmitri Maslov, Patrick Rall, Theodore J. Yoder, *Nature*, March 2024 (11 pages).

<sup>108</sup> Long-Distance Transmon Coupler with cz-Gate Fidelity above 99.8 % by Fabian Marxer et al, PRX Quantum, February 2023 (23 pages).

<sup>109 &</sup>lt;u>A superconducting quantum information processor with high qubit connectivity</u> by Gürkan Kartal *et al, IMEC*, July 2023 (8 pages).

<sup>110</sup> LDPC-cat codes for low-overhead quantum computing in 2D by Diego Ruiz, Jérémie Guillaud, Anthony Leverrier, Mazyar Mirrahimi, and Christophe Vuillot, arXiv, January 2024 (23 pages).

<sup>111</sup> Passive superconducting circulator on a chip by Rohit Navarathna, Thomas M. Stace, Arkady Fedorov *et al*, *PRL*, August 2022-January 2023 (11 pages).

<sup>112</sup> Josephson parametric circulator with same-frequency signal ports, 200 MHz bandwidth, and high dynamic range by Randy Kwende *et al*, Google AI, May 2023.

economic point of view when the technology is stabilized and the volume of quantum computer production enables it. Riverlane (UK), a partner of Alice&Bob and IQM, is considering this move.

Interconnection between quantum processors will initially involve short-range interconnection using direct chip-to-chip links via indium-based connections (Rigetti, IBM), followed by medium-range microwave guides (IBM). Beyond this, there are plans to exploit transduction between microwaves (5 GHz) and optical photons (QPhox<sup>113</sup>, Netherlands) and to establish synchronized links using quantum memories based on sets of neutral atoms (Welinq<sup>114</sup>, France). All these techniques are still highly experimental. Optical photons links require the creation of entanglement based on Bell pairs, which are not deterministic and whose impact on computing overhead needs to be assessed. The ultimate goal is to create two-qubit gates, like a CNOT, between qubits located in two distant processors. These gates should have fidelities of around 99.9% to support error correction and be able to be established between several qubits in interconnected processors, so as not to make the algorithms too cumbersome. But depending on the error correction codes used across multiple QPUs, this fidelity constraint could be relaxed.

Cryostats will have to grow in power and capacity to support a greater number of cables and passive electronic components. The current commercial record is the Bluefors KIDE with 9 compressors and 3 dilutions, supporting a very large load. IBM plans to support thousands of physical qubits with this type of cryostat. Air Liquide is planning to create cold production systems based on liquid helium, providing greater efficiency and cooling power.

<sup>113 &</sup>lt;u>An integrated microwave-to-optics interface for scalable quantum computing</u> by Matthew J. Weaver, Simon Gröblacher, Robert Stockill *et al*, QphoX, *Nature Nanotechnology*, October 2022-October 2023 (14 pages).

<sup>114 &</sup>lt;u>Connecting heterogeneous quantum networks by hybrid entanglement swapping</u> by Giovanni Guccione, Tom Darras *et al.*, May 2020 (7 pages).

### 3.2.5. ROADMAPS

Roadmaps for commercial players that include the ability to create fault- tolerant systems have been announced by IBM, Google and Alice&Bob. These roadmaps are generally not detailed in terms of technology, due to frequent changes in the technological options chosen. Other players in the superconducting qubit market, such as Amazon, Nord Quantique, Rigetti, and Atlantic Quantum, have not yet announced a public *roadmap* for fault tolerance. IQM and OQC have announced such a roadmap in 2025.

IBM's initial roadmaps, announced between 2020 and 2022, forecasted the use of surface codes and the support of around 100 logical qubits with 100,000 physical qubits. Since their announcement of the adoption of qLDPC codes in 2023 and the use of what they call *gross codes*, the number of logical qubits achievable with this number of physical qubits has increased. Like many other market players, IBM provides the number of executable quantum gates in an algorithm as a performance indicator for these systems. This indicator is common to both NISQ and FTQC systems. IBM plans to assemble dozens of physical qubits with processors of modest size, around 156 qubits, which are connected together with at least three types of coupler: short-distance microwave, medium-distance microwave and optical photonic links (infrared) with signal transduction.

IBM has respected its *roadmap* since 2020, which was described in terms of the number of qubits. The *roadmap* was then readjusted to adopt a modular strategy, based on chips with 156 physical qubits and tunable couplers for creating two-qubit gates.

At Google, the plan has not changed much since 2020 and consists of assembling a million physical qubits by the beginning of the 2030s, to obtain around a hundred logical qubits. All this is still based on the use of surface codes. But with no further details on the implementation. This would create an inordinately large system. Google has to improve the quality of its physical qubits to respect its *roadmap*.

Considering the creation of their first logical qubit below threshold in August 2024, i.e. logical qubits with better fidelities than physical qubits, the projections do not change too much. They will need thousands of physical qubits to create a logical qubit with error rates of between 10<sup>-6</sup> and 10<sup>-14</sup>. And these physical qubits will not fit on a single chip, even accounting for their plan to create chips with up to 10K physical qubits. Their ability to scale up will therefore depend very much on the availability of interconnection techniques between chips based on microwave guides and optical photons, via transduction means, and whose error rates and efficiencies will have to be compatible with the requirements of error correction codes.

Alice&Bob plans to support around a hundred logical qubits with an error rate of 10<sup>6</sup> on a chip containing around 2,000 physical qubits by 2030, with intermediate stages. The manufacturing techniques needed to house all these physical qubits on a single chip have yet to be put in place. The number of chips that will be needed to achieve an FTQC regime of several hundred and then thousands of logical qubits has not yet been determined, nor have the interconnection techniques envisaged between these chips.

#### 3.2.6. Conclusion

There are many challenges to overcome in creating fault-tolerant architectures. As with all types of qubits, these challenges have both theoretical and experimental scientific dimensions as well as technological and engineering ones. The main one, however, is to create chips with more than a hundred physical qubits, with gate fidelities of 99.9% between two qubits, and then to link these chips together using entanglement resources while maintaining the quality of operations at the same level.

# 3.3. PHOTONS

Photons form a separate category of qubits, known as "flying qubits". They have no mass or electric charge, and travel through vacuum at the speed of light. In all other types of qubits, photons interact with atoms, superconducting circuits or electron spins to modify their quantum state or create entanglement between these qubits, whether these are photons in the microwave regime or optical photons operating in the visible or near infrared spectrum.

Photons carry quantum information with one or more of their properties: their polarization, wavelength, time of emission, number, path and/or orbital angular momentum. Photons are less prone to decoherence than other types of qubits. They also have the advantage of being used for quantum communications, particularly those that will eventually link quantum computers together to increase their power.

However, it is difficult to make them interact with each other. You also have losses at several levels: in their source, in the waveguides that carry them, in the optical systems that transform them (phase shifters, polarizers, interferometers) and finally, in their detectors. Their behavior is highly probabilistic at each of these levels.

As with other types of qubits, photons can be exploited in many ways. The first classification distinguishes photons handled with discrete or continuous variables<sup>115,116</sup>.

 Photon qubits with discrete variables where information can be encoded in different degrees of freedom (polarization, path, energy, time, etc.),

<sup>115 &</sup>lt;u>Photonic Quantum Computing by Jacquiline Romero, and Gerard Milburn, University of Queensland</u>, arXiv, April 2024 (26 pages).

<sup>116</sup> Quantum computing overview: discrete vs. continuous variable models by Sophie Choe, June 2022 (12 pages)

orbital angular momentum, etc.) which can even be combined<sup>117, 118, 119, 120</sup>. They are based on deterministic or probabilistic sources of single, indistinguishable photons. They are based on the particle dimension of photons. This is the PsiQuantum and Quandela approach.

Continuously variable photon qubits encode information in the fluctuations of the electromagnetic field, exploiting their quadrature decomposition. These qubits are called qumodes at Xanadu. They primarily exploit the wave nature of photons. These photons can also be implemented with cat qubit models and GKP states<sup>121, 122,123</sup>.

The general principle of photon-based quantum computers is based on the following steps, using discrete variables as an example:

 photon generation, which combines lasers and generators of single, indistinguishable photons<sup>124</sup>. These photons are produced at a given frequency and can then be distributed over several waveguides via delay lines. The figures of merit of these sources are the indistinguishability of the photons (they must all have the same quantum properties), the luminosity of the source (how many photons are emitted per

<sup>117 &</sup>lt;u>Programmable Photonic Quantum Circuits with Ultrafast Time-bin Encoding</u> by Frédéric Bouchard, Benjamin Sussman *et al*, University of Ottawa, *arXiv*, April 2024 (7 pages).

<sup>118 &</sup>lt;u>Time-bin entanglement in the deterministic generation of linear photonic cluster states</u> by David Bauch, Nikolas Köcher, Nils Heinisch, and Stefan Schumacher, Paderborn University, *arXiv*, March 2024 (11 pages).

<sup>119 &</sup>lt;u>Efficient qudit based scheme for photonic quantum computing by Márton Karácsony, László Oroszlány and Zoltán Zimborás, February 2023 (19 pages).</u>

<sup>120 &</sup>lt;u>Deterministic Generation of Qudit Photonic Graph States from Quantum Emitters</u> by Zahra Raissi, Edwin Barnes, and Sophia E. Economou, *PRX Quantum*, November 2022-May 2024 (20 pages).

<sup>121 &</sup>lt;u>Robust Preparation of Wigner-Negative States with Optimized SNAP-Displacement Sequences</u> by Marina Kudra, Jonas Bylander, Simone Gasparinetti *et al*, Chalmers University, *PRX Quantum*, September 2022 (12 pages).

<sup>122 &</sup>lt;u>Gottesman-Kitaev-Preskill encoding in continuous modal variables of single photons</u> by Éloi Descamps, Arne Keller, and Pérola Milman, *PRL*, October 2023-April 2024 (8 pages).

<sup>123</sup> Logical states for fault-tolerant quantum computation with propagating light by Shunya Konno, Akira Furusawa et al, Science, January 2024 (11 pages).

<sup>124 &</sup>lt;u>Near-ideal spontaneous sources in silicon quantum photonics by S. Paesani et al</u>, 2020 (6 pages).

unit time), system efficiency (the probability of at least one photon being created per pulse), single-photon purity (probability of obtaining a maximum of one photon per pulse), photon rate (rate of photon generation) and interphoton anti-correlation, measured by the second-order coherence value  $g^{(2)}(0)$ which corresponds to the probability of emitting two photons simultaneously, ideally at zero. Other operational aspects include stability and time coherence (coherence of the source over time), spectral properties, and the ability to control the wavelength of the emitted photons and the operating temperature, generally at 4 K. The purity and high probability of obtaining one photon per clock cycle make it possible to generate interferences and two-photon quantum gates. There are several types of single-photon sources<sup>125, 126</sup> : guantum-dots-based sources such as those from Quandela and Sparrow Quantum, which have very good figures of merit (purity of 99.7%, extraction efficiency of 65%, indistinguishability of 99%), and parametric photon-pair sources that exploit laser pumping in non-linear cavities and can be directly integrated into nanophotonics circuits, but have lower efficiencies. This is the choice made by PsiQuantum;

- One-qubit gates are easy to build and use simple optical devices such as phase shifters, beam splitters, polarizers and semi-reflecting mirrors. For example, a Hadamard gate uses a beam splitter, an X gate combines a beam splitter and a Hadamard gate, and finally, a Z gate relies on a 180° phase shifter.
- Two-qubit gates are more difficult to create because it is very difficult, but not impossible, to make photons interact deterministically. In the KLM<sup>127</sup> model, these gates are made using beam splitters and assembled into so-called Mach-Zehnder interferometers.

<sup>125</sup> Integrated photonic quantum technologies by Jianwei Wang et al, May 2020 (16 pages).

<sup>126 &</sup>lt;u>Solid-state single-photon sources: recent advances for novel quantum materials</u> by Martin Esmann *et al*, December 2023 (35 pages).

<sup>127 &</sup>lt;u>A scheme for efficient quantum computation with linear optics</u> by Emanuel Knill, Raymond Laflamme and Gerard Milburn, 2001 (7 pages).

The KLM model is based on a so-called post-selection mechanism which sorts operations according to their success, with an obvious loss of line at each operation; the qubits are read using single-photon detectors or *photon number resolving* detectors (PNRDs). Several technologies are available: avalanche diodes (SPADs), which detect photon occurrences but not photon numbers<sup>128,129</sup>, transition edge sensors (TESs), which detect photon numbers, and superconducting nanowires (SNSPDs), which can also detect photon numbers<sup>130, 131</sup>. These detectors are generally cooled to between 800 mK and 3 K<sup>132, 133</sup>. The main suppliers of single-photon detectors are Single Quantum, Photon Spot, Quantum Opus and IDQ, mainly using SNSPDs;

The integration of all or some of these elements into nanophotonics circuits<sup>134,135</sup>. Operations on photons are then electrically controlled at the level of individual circuits such as phase shifters and polarizers. These circuits are etched in CMOS (silicon) technology and/or in a nanophotonic circuit.

<sup>128</sup> Low-noise photon counting above 100× 106 counts per second with a high-efficiency reachthrough single-photon avalanche diode system by Michael A. Wayne *et al*, *NIST*, December 2020 (6 pages).

<sup>129</sup> Photon Number Resolving Detection with a Single-Photon Detector and Adaptive Storage Loop by Nicholas M. Sullivan *et al*, University of Ottawa, November 2023 (16 pages).

<sup>130 &</sup>lt;u>GHz detection rates and dynamic photon-number resolution with superconducting nanowire arrays</u> by Giovanni V. Resta *et al, ID Quantique*, March 2023 (26 pages).

<sup>131</sup> Optically-Sampled Superconducting-Nanostrip Photon-Number Resolving Detector for Non-Classical Quantum State Generation by Mamoru Endo, Akira Furusawa et al, arXiv, May 2024 (24 pages).

<sup>132</sup> The potential and challenges of time resolved single-photon detection based on current- carrying superconducting nanowires by Hengbin Zhang et al., October 2019 (19 pages) and Superconducting nanowire single-photon detectors for quantum information by Lixing You, June 2020 (20 pages). Dark counts are detected photons coming from the environment due to thermal or tunneling effects.

<sup>133</sup> Optimal Amplitude Multiplexing of a Series of Superconducting Nanowire Single Photon Detectors by Fabio Chiarello *et al*, March 2023 (6 pages).

<sup>&</sup>lt;sup>134</sup> See for example the InPhyNi work presented <u>High-quality photonic entanglement based on a silicon chip</u> by Dorian Oser, Sébastien Tanzilli *et al*, 2020 (9 pages).

<sup>135</sup> See The potential and global outlook of integrated photonics for quantum technologies by Emanuele Pelucchi, Dirk Englund, Jian-Wei Pan, Fabio Sciarrino, Christine Silberhorn *et al*, *Nature Review Physics*, December 2021 (no open access).

or III/V<sup>136</sup> germanium, etc.) or hybrids, with a CMOS substrate to which III/V elements are bonded. Manufacturing techniques involve compounds such as silicon nitride (SiN), lithium niobate (LiNbO<sub>3</sub>)<sup>137</sup> and III/V materials (GaAs<sup>138</sup>, InP...)<sup>139, 140</sup>. Light guides are generally made of silicon and surrounded by silicon oxide, the key point is to minimise photon losses<sup>141</sup>. Another challenge is to reduce the heat generated by phase shifters and optimise their calibration<sup>142</sup>.

The KLM model in use today does not yet provide a proven quantum advantage at scale. It is considered to be part of the NISQ regime, but not for the usual reason. The pitfall here is that the computational efficiency decreases exponentially with the number of two-qubit gates implemented because of their probabilistic aspect.

The computing paradigm for scaling and fault tolerance is MBQC, or measurement-based computing. This is an approach based on quantum teleportation, using cluster states represented by graphs, in which the vertices correspond to qubits and the edges are entanglement links. This approach does not use two-qubit gates, but only single-qubit gates and single-qubit measurements, making it particularly well suited to photonic qubits. Computing is performed by carrying out measurements whose effects propagate from one qubit to the next thanks to the entanglement.

<sup>136</sup> III/V or III-V materials: pairs combining elements from column III and column V of the classification.

<sup>137 &</sup>lt;u>High-speed thin-film lithium niobate quantum processor driven by a solid-state quantum emitter</u> by Patrick I. Sund *et al*, NBI, CeNTech, *Science Advances*, May 2023 (9 pages).

<sup>138</sup> Expanding the Quantum Photonic Toolbox in AlGaAsOI by Joshua E. Castro *et al*, May 2022 (9 pages). They implement nonlinear elements, edge couplers, waveguide crossings, couplers, and MZIs in Aluminum gallium arsenide-on-insulator (AlGaAsOI).

<sup>139 &</sup>lt;u>Roadmap on integrated quantum photonics</u> by Galan Moody, Jacquiline Romero, Eleni Diamanti *et al*, August 2021 (108 pages) is a good review paper on integrated nanophotonics.

<sup>140</sup> Advances in silicon quantum photonics by Jeremy C. Adcock et al, July 2022 (25 pages) describes the challenges of nanophotonics.

<sup>141 &</sup>lt;u>Mitigating photon loss in linear optical quantum circuits: classical postprocessing methods outperforming postselection</u> by James Mills, and Rawad Mezher, Quandela and University of Edinburgh, *arXiv*, May 2024 (31 pages).

<sup>142 &</sup>lt;u>Global calibration of large-scale photonic integrated circuits</u> by Jin-Hao Zheng, Guang-Can Guo *et al*, Hefei National Laboratory, CAS, *arXiv*, July 2024 (9 pages).

Quantum information can be preserved by virtual qubits.

In other words, it propagates from one qubit to the next over the physical qubits that are regularly measured. On the other hand, the measurement results for a set of qubits at a given time determine the measurement bases for subsequent qubits. This requires that the parameters of the physical components can be reconfigured ("*feed-forward*") very quickly.

The MBQC model burden is in the creation of large entangled states of photons, aka cluster or graph states<sup>143</sup> and in the error correction codes that also need to be implemented. These can be based on the so-called foliation process, which enables correction codes to be constructed on three-dimensional cluster states. This requires the creation of three-dimensional cluster states with many qubits. This can be achieved by merging small entangled states, or using hybrid strategies such as percolation, PsiQuantum's *fusion-based quantum computing* (FBQC) model or, in the case of Quandela, the simultaneous use of quantum dot spins and photons as information carriers. This requires the ability to manufacture quantum dots with identical properties and at scale.

<sup>143 &</sup>lt;u>Multidimensional cluster states using a single spin-photon interface strongly coupled to an intrinsic nuclear register</u> by Cathryn P. Michaels *et al*, University of Cambridge, April 2021 (11 pages).

#### photon qubits key takeaways

	highlights	challenges
	<ul> <li>stable qubits with absence of decoherence.</li> <li>ambient temperature for processing.</li> <li>emerging nano-photonic manufacturing techniques enabling scalability.</li> <li>easier to scale-out with inter-qubits communications and quantum telecommunications.</li> <li>MBQC/FBQC circumventing the fixed gates depth computing capacity and difficulty to create multiple qubit gates.</li> <li>boson sampling-based quantum advantage: starts to being programmable but a practical quantum advantage remains to be proven.</li> </ul>	need to cool photon sources and detectors: but at relatively reasonable temperatures between 2K and 10K, requiring lighweight cryogenic systems (unless also cooling the whole photonic circuit). not yet scalable in number of operations: due to probabilistic character of quantum gates and the efficiency of photon sources in most paradigms. non-deterministic cluster states and two-photon interactions. needs delay lines and optical switches. heat generated by phasers: increasing cooling budget.
		photon detectors efficiency: 89% with PsiQuantum. photon losses in nanophotonic circuits.
	variations	path to scalability
	encoding: direct variables qubits (DV), continous variables (CV) qubits, multimode photon encoding. MBQC (measurement based quantum computing) and FBQC	efficiency: improve photon sources efficiency, determinism and indistinguishability, improve photon detectors efficiency (SNSPD and PNRD).
	(fusion-based quantum computing).	cluster states: generate large cluster states, with or without heralding.
	BS/GBS: programmable Boson sampling and Gaussian boson	interactions: improve fusion efficiency in FBQC.
	sampling. hybrid approach: spin-optical quantum computing (SPOQC) with quantum dots spin qubits (Quandela), hybrid atom-photon qubits.	losses: large-scale and low-losses optical switches and wave guides, reduce photon losses in nanophotonic circuits thanks to higher precision manufacturing and new materials.
	classical photonic models: coherent Ising models, photonic waveguide arrays and interferometric systems.	energy: create low-heating phasers to minimize power consumption. nanophotonics: heterogeneous nanophotonic circuits (III-V + silicon). classical control speed: particularly with FBQC models.

Figure 12: Strengths, challenges, varieties and scalability options for photon qubits.

#### 3.3.1. Advances and key benefits

Here are the current highlights of photon-based qubits.

These are stable qubits with no decoherence, at least over short ranges.

They enable processing at ambient temperature, as is the case with Quandela. This is not the approach adopted by PsiQuantum, which cools its nanophotonics circuits to 1.8 K to 4 K, directly integrating photon sources and detectors.

Good scalability, made possible by nanophotonics manufacturing techniques that will eventually enable a large number of photon qubits to be supported on a chip.

The interconnection between photonic quantum computers will be facilitated by the fact that it will also be based on photons. However, the wavelengths used for photonic quantum computing are not necessarily the same as those that will be used to establish distant quantum links. This will probably require frequency conversions. However, these are easier to achieve than, for example, switching from microwaves to infrared photons, which is necessary for superconducting qubits. The ease with which entanglement can be mediated by photons is a point that photon qubits have in common with qubits based on neutral atoms and trapped ions.

Computation models based on MBQC/FBQC measurements will make it possible to get around the probabilistic aspect of two-qubit gates and envision the creation of photonic quantum computers that are fault-tolerant and able to scale up.

Systems based on boson sampling are beginning to be programmable, with a practical quantum advantage seemingly possible, particularly for solving various optimization problems.

#### 3.3.2. CHALLENGES

There are many of them, including:

The need to cool photon sources and detectors, but at relatively reasonable temperatures ranging from 2 K to 10 K, requiring lightweight cryogenic systems, unless the entire photonic circuit is also cooled.

Current photonic quantum computers are not yet scalable in terms of number of operations due to the probabilistic nature of quantum gates and the efficiency of photon sources in most paradigms.

The creation of deterministic, or as deterministic as possible, clusters of entangled photons for the implementation of the MBQC computing model.

The need for low-loss waveguides and optical switches, which have yet to be developed.

The heat generated by the phasers increases the cooling budget, particularly when the circuits are cooled to the same temperature as the sources and detectors.

The efficiency of photon detectors is still too low, at 89% at PsiQuantum, bearing in mind that they need to develop multiple photon detectors.

### 3.3.3. VARIATIONS

In addition to the continuous variable encoding and MBQC models already mentioned, there are other photon-based quantum or semi-quantum computing technologies.

*Boson* sampling is the brainchild of Scott Aaronson and Alex Arkhipov in a paper published in 2010<sup>144</sup>. They proposed a linear optics system that would be impossible to simulate efficiently in a classical computer. It involves solving a problem of sampling the distribution of identical and indistinguishable photons that are mixed in an interferometer and then end up in photon detectors. The classical simulation of the experiment is mathematically very expensive, as it is based on the evaluation of the permanent of square matrices<sup>145</sup>, a "#P hard" class<sup>146</sup>. Verifying the result would even exceed the capabilities of conventional computers. The first boson sampling experiments used passive optical components that could not be parameterized<sup>147</sup>.

The computational Complexity of Linear Optics by Alex Arkhipov and Scott Aaronson,

 Proceedings of the forty-third annual ACM symposium on Theory of computing 2010 (94 pages).

<sup>145 &</sup>lt;u>Lecture 3: Boson sampling</u> by Fabio Sciarrino, University of Rome, (63 slides), <u>Permanents and boson</u> <u>sampling</u> by Stefan Scheel, University of Rostock, 2018 (21 slides).

<sup>146 #</sup>P is the class of problems consisting of counting the number of solutions to NP problems.

<sup>147 &</sup>lt;u>An introduction to boson-sampling</u> by Jonathan Dowling *et al.*, 2014 (13 pages) describes well the issues involved in conducting boson sampling.

Then, progressively, in China in particular, researchers experimented parameterizable setups, with input and output data and the ability to solve various complex problems, in particular combinatorial problems involving graph theory. One variant developed in China is the Gaussian boson sampler (GBS) with Jiuzhang 1.0 in 2020 with 70 photon modes<sup>148, 149</sup>. They calculate the Hafnian equivalent, which evaluates the permanent of symmetric square matrices<sup>150</sup>. This was followed by Jiuzhang 2.0 in 2021, which was parameterizable at the level of the phase of incoming photons<sup>151</sup>, and with the resolution of graph problems in 2023<sup>152</sup>. The same year, Jiuzhang 3.0 supported the detection of 255 photons<sup>153</sup>. In 2024, another Chinese team created a programmable GBS that could solve a therapeutic molecule search problem using an RNA folding technique<sup>154</sup>. In 2024. researchers at the University of Maryland used a GBS to solve some machine learning problem<sup>155</sup>. In 2022, Xanadu also created a programmable GBS with a simpler setup, exploiting photon time-division multiplexing. The system was used to solve graph problems<sup>156</sup>. In practice, Xanadu's experiment was a first step on the road to fault tolerance in a measurement-based model (MBQC).

<sup>148 &</sup>lt;u>Quantum computational advantage using photons</u> by Han-Sen Zhong *et al*, December 2020 (23 pages) and the supplemental materials (64 pages).

<sup>149 &</sup>lt;u>Benchmarking 50-Photon Gaussian Boson Sampling on the Sunway TaihuLight</u> by Yuxuan Li *et al.*, 2020 (12 pages).

<sup>150</sup> The Second Moment of Hafnians in Gaussian Boson Sampling by Adam Ehrenberg et al, arXiv, March 2024 (34 pages).

<sup>151 &</sup>lt;u>Phase-Programmable Gaussian Boson Sampling Using Stimulated Squeezed Light by Han-Sen Zhong, Chao-Yang Lu, Jian-Wei Pan et al</u>, June 2021 (9 pages).

<sup>152 &</sup>lt;u>Solving Graph Problems Using Gaussian Boson Sampling</u> by Yu-Hao Deng *et al*, February 2023 (7 pages).

<sup>153</sup> Gaussian Boson Sampling with Pseudo-Photon-Number Resolving Detectors and Quantum Computational Advantage by Yu-Hao Deng et al, China, PRL, April-October 2023 (6 pages).

<sup>154 &</sup>lt;u>A universal programmable Gaussian Boson Sampler for drug discovery</u> by Shang Yu, Ian A. Walmsley, Guang-Can Guo *et al*, *arXiv*, October 2022-March 2024 (11 pages).

<sup>155 &</sup>lt;u>Biclustering a dataset using photonic quantum computing</u> by Ajinkya Borle, and Ameya Bhave, University of Maryland, *arXiv*, May 2024 (32 pages).

<sup>156</sup> Using Gaussian Boson Sampling to Find Dense Subgraphs by Juan Miguel Arrazola and Thomas R. Bromley, March-July 2018 (6 pages).

Coherent Ising *Machines* (CIMs)<sup>157</sup> are based on photonic waveguide networks and interferometric systems. The technique uses optical neural networks to solve combinatorial problems, which are then converted into NP-hard Ising models<sup>158</sup>. One of the largest CIMs has been set up in Japan in 2021, with problems to be solved involving 100,000 spins<sup>159</sup>. This type of solution competes with D-Wave's quantum annealing and Fujitsu's classical annealer<sup>160</sup>.

Simulations based on quantum walks that work with photons with discrete or continuous variables<sup>161,162</sup>. In 2022, a Chinese team created a quantum walk with continuous variables generating a Hilbert space of dimension 400<sup>163</sup>. Quantum walks could be produced using programmable waveguide arrays<sup>164, 165</sup>.

Various hybrid approaches exist, such as Quandela's Spin-Optical Quantum Computation (SPOQC) proposal, which uses their quantum dots spin as a quantum memory to enable the creation of entangled cluster states of photons,

<sup>157</sup> See for example https://www.bing.com/search? q=lsing % 20&qs=n&form=QBRE&sp=-1&ghc=2&lq=0&pq=ising % 20&sc=11-6&sk=&cvid=566E3737143A46EC9B64666E40F84401

<sup>158 &</sup>lt;u>Coherent Ising Machines: non-von Neumann computing using networks of optical parametric oscillators</u> by Peter McMahon, Cornell University, October 2020 (100 slides).

<sup>159 100,000-</sup>spin coherent Ising machine by Toshimori Honjo et al, September 2021 (8 pages).

<sup>100 &</sup>lt;u>Coherent Ising machines - optical neural networks operating at the quantum limit</u> by Y. Yamamoto *et al, npj Quantum*, December 2017 (16 pages).

<sup>161 &</sup>lt;u>Quantum walks of two correlated photons in a 2D synthetic lattice</u> by Chiara Esposito, Fabio Sciarrino *et al.*, April 2022 (18 pages).

<sup>162 &</sup>lt;u>Probing quantum walks through coherent control of high-dimensionally entangled photons by</u> Poolad Imany *et al.*, July 2020 (9 pages).

<sup>163 &</sup>lt;u>Large-scale full-programmable quantum walk and its applications</u> by Yizhi Wang *et al*, August 2022 (73 pages)

<sup>164</sup> Programmable high-dimensional Hamiltonian in a photonic waveguide array by Yang Yang, Robert J. Chapman, Alberto Peruzzo et al, ETH Zurich, Griffith University, Purdue University, University of Trento, Heriot-Watt University and Qubit Pharmaceuticals, *Nature Communications*, January 2024 (7 pages).

Programmable quantum circuits in a large-scale photonic waveguide array by Yang Yang, Alberto Peruzzo et al, ETH Zurich, RMIT University, University of Muenster, Heriot-Watt University, University of Trento, Qubit Pharmaceuticals. arXiv. May 2024 (15 pages).

or approaches combining neutral atoms and photons<sup>166</sup>.

### 3.3.4. PATHWAYS TO SCALABILITY

The development of fault-tolerant photon-based quantum computers requires a number of technological challenges to be resolved:

Improvements to photon sources in terms of efficiency, determinism, indiscernibility and, above all, the ability to generate large cluster states.

Improving the efficiency of photon detectors way beyond 90%.

Improving fusion efficiency in the FBQC model proposed by PsiQuantum.

Reducing photon losses in nanophotonics circuits requires more precise manufacturing and the use of new materials.

The creation of optical switches for routing photons between the various circuits and for implementing *feed-forward* in the FBQC model.

Reducing energy consumption by creating low-energy phasers.

Developing heterogeneous nanophotonics components combining III-V bricks on silicon by bonding.

Architecting classical circuit control to keep pace with photons, in particular with FBQC and MBQC models.

<sup>106</sup> Deterministic photonic quantum computation in a synthetic time dimension by Ben Bartlett, Avik Dutt and Shanhui Fan, Optica, November 2021 (9 pages).

### 3.3.5. ROADMAPS

Here we take a look at the roadmaps of the main players in the photon qubit industry, with a view to creating fault-tolerant quantum computers.

PsiQuantum (USA/UK/Australia) has set itself the challenge of creating a faulttolerant quantum computer without using a NISQ machine in KLM mode like Quandela. Its immediate goal is to create 100 logic qubits by the early 2030s, using one million physical qubits, or photons.

PsiQuantum's architecture is based on a variant of the MBQC model known as FBQC (*Fusion-based quantum computation*). MBQC relies on the creation of giant clusters of entangled photons, which are very difficult to create. The FBQC model gets round the problem by creating clusters of limited size, of the order of 4 entangled photons, and linking them by probabilistic joint measurements. The qubits are encoded "along the way" using the *dual* rail technique. The technique is tricky to evaluate, because it involves probabilistic processes at different levels: the generation of *cluster states*, the mergers between these *cluster states*, and the detection of photons.

PsiQuantum's other choice is to integrate all the photonics into silicon chips: photon sources, waveguides and calculation operations, photon routing and photon detectors, all operating at 1.8 K. The chips are manufactured on 300mm *wafers* in 22nm FD-SOI technology at GlobalFoundries in New York State. The 100 logical qubit QPU will be based on thousands of these chips linked together by photons.

In April 2024, the company described its progress in developing this chip<sup>167</sup>. It is reporting very good transaction rates for the generation of

<sup>167 &</sup>lt;u>A manufacturable platform for photonic quantum computing</u> by Koen Alexander *et al*, PsiQuantum, *arXiv*, April 2024 (8 pages).

photons and for one- and two-qubit gates<sup>168</sup>. The pitfalls lie in the efficiency of photon detection, which is only 89%, and in photon losses in the circuits. Numerous other challenges remain in photon generation, the creation of low-loss silicon nitride waveguides with more efficient phase shifters, the creation of a matrix photon routing system, directional couplers, couplings between circuits and optical fibers and the creation of multiple photon detectors exploiting spatial multiplexing of single-photon detectors of the SNSPD type<sup>169</sup>. Their roadmap envisages solving all these problems in six years, which seems rather optimistic.

Finally, an initial estimate of the power needed to power this machine is around 10 MW just for cryogenics, which would be prohibitive if the power needed to generate the thousands of logic qubits required for numerous algorithms increased linearly with their number.

Xanadu Quantum Technologies (Canada) is also developing a fault-tolerant photonic quantum computer based on continuous variable photons, but on an MBQC model and using GKP states encoding<sup>170, 171,172</sup>. This is documented in a paper published in 2021<sup>173</sup>. Its silicon nitride chips were manufactured by IMEC in Belgium, but in 2022, the company announced a partnership

<sup>168</sup> With: 99.98 %± 0.01 % state preparation and measurement fidelity, Hong-Ou-Mandel quantum interference between independent photon sources with 99.50 %± 0.25 % visibility, two-qubit fusion with 99.22 %± 0.12 % fidelity, and a chip-to-chip qubit interconnect with 99.72 %± 0.04 % fidelity.

<sup>169 &</sup>quot;The performance of the baseline technology described above is still not sufficient for useful photonic quantum computing. In particular, silicon waveguides incur too much propagation loss for fault tolerance, photon sources require complex and power-hungry tuning, and high-speed optical switching is unavoidably necessary to overcome the intrinsic nondeterminism of the spontaneous single photon sources".

<sup>170</sup> The power of one qumode for quantum computation, 2016 (10 pages), <u>Continuous-variable gate</u> decomposition for the Bose-Hubbard model, 2018 (9 pages).

<sup>171 &</sup>lt;u>Optical hybrid approaches to quantum information</u> by Peter van Loock, 2010 (35 pages).

<sup>172 &</sup>lt;u>Quantum computing with multidimensional continuous-variable cluster states in a scalable photonic platform</u> by Bo-Han Wu *et al.*, 2020 (22 pages).

<sup>173 &</sup>lt;u>Blueprint for a Scalable Photonic Fault-Tolerant Quantum Computer</u> by J. Eli Bourassa *et al*, February 2021 (38 pages).

with GlobalFoundries. It plans to reach 1,000 logical qubits, which would require 10,000 data center racks, which is somewhat worrying. In 2022, Xanadu demonstrated a quantum advantage achieved with a Gaussian boson sampler (GBS) in its Borealis QPU, achieving a record 216 photon modes based on frequency multiplexing and delay lines. His system can be parameterized in terms of the phases of the photons entering the sampler<sup>174</sup>.

Quandela (France) plans to adopt a variation of the FBQC model to create a fault-tolerant computer. Its research work is manifold, starting with the generation of highly efficient *cluster states* of entangled photons<sup>176</sup>, the polarization of the photons generated being controlled by the magnetic field to which the quantum dot generating them is exposed<sup>176</sup>. For the moment, the company is working on clusters of 4 photons, to be expanded to 24<sup>177</sup>. In November 2023, Quandela published a roadmap detailing its SPOQC (*spin-optical quantum computing*) model, which combines data qubits supported by the photon-source quantum dots, and the photons they emit as auxiliary qubits to manage two-qubit gates between the quantum dots<sup>178</sup>. These gates use the *repeat until success* (RUS) technique and photon routing to link all the quantum dots together, a challenge similar to the photon routing envisaged by PsiQuantum. The architecture is based on qLDPC-type error correction codes, which require such connectivity. Another challenge is to extend the coherence time of the quantum dots to support

<sup>174 &</sup>lt;u>The hardness of quantum spin dynamics</u> by Chae-Yeun Park *et al*, Xanadu, Sungkyunkwan University, December 2023 (30 pages).

<sup>175 &</sup>lt;u>High-rate entanglement between a semiconductor spin and indistinguishable photons by Nathan</u> Coste, Sophia Economou, Niccolo Somaschi, Alexia Auffèves, Loic Lanco, Pascale Senellart *et al, Nature Photonics*, July 2022 (17 pages).

<sup>176 &</sup>lt;u>Controlling photon polarisation with a single quantum dot spin</u> by Elham Mehdi, Pascale Senellart, Loic Lanco *et al*, December 2022 (9 pages).

<sup>177</sup> Quantifying n-photon indistinguishability with a cyclic integrated interferometer by Mathias Pont, Fabio Sciarrino, Pascale Senellart, Andrea Crespi *et al*, PRX, January-September 2022 (21 pages).

<sup>178 &</sup>lt;u>A Spin-Optical Quantum Computing Architecture</u> by Grégoire de Gliniasty, Paul Hilaire, Pierre-Emmanuel Emeriau, Stephen C. Wein, Alexia Salavrakos, Shane Mansfield, Quandela and LIP6, November 2023 (20 pages).

the generation of a sufficient number of photons which is currently far too weak to support this model. The record in this field comes from the University of Cambridge, with a coherence of 0.11 ms<sup>179, 180</sup>. Quandela has also established a technological partnership with Welinq (France) to develop an interconnection solution between quantum processors using Welinq's quantum memories based on neutral atoms.

ORCA Computing (UK) is developing a quantum photonic computing platform based on qumode continuous variable photons and quantum memory technology using delay lines and beam splitters<sup>181,182</sup>. Its chips are manufactured by Ligentec in Switzerland, as is Quandela. It is also working on the creation of *cluster states* and an FBQC model similar to that of PsiQuantum<sup>183,184</sup>.

QuiX Quantum (Netherlands) is developing low-loss photonic quantum computing circuits based on silicon nitride  $(Si_3N_4)$  using thermo-optical phase shifters and programmable optical beam splitters<sup>185</sup>. By 2022, they were supporting 20 qumodes<sup>196</sup> using 380 phase shifters and a photon source exploiting a Ti:Sapphire laser and a crystal. The company expects to be able to support 10,000 qumodes after 2030, but has not published a genuine FTQC roadmap.

<sup>179 &</sup>lt;u>Researchers find ways to improve the storage time of quantum information in a spin-rich</u> material by the Cavendish Laboratory of the Department of Physics of the University of Cambridge, January 2023.

<sup>180</sup> Ideal refocusing of an optically active spin qubit under strong hyperfine interactions by Leon Zaporski, Claire Le Gall (who is now VP of the quantum management team at Nu Quantum) et al, University of Cambridge, University of Sheffield, University of Oxford, Nature Nanotechnology, January 2023 (23 pages).

<sup>181 &</sup>lt;u>One-Way Quantum Computing in the Optical Frequency Comb</u> by Nicolas C. Menicucci, Steven T. Flammia and Olivier Pfister, April 2018 (4 pages).

<sup>182 &</sup>lt;u>High-speed noise-free optical quantum memory</u> by K. T. Kaczmarek et al, April 2018 (12 pages).

High photon-loss threshold quantum computing using GHZ-state measurements by Brendan Pankovich *et al*, Orca Computing, August 2023 (15 pages).

<sup>&</sup>lt;sup>184</sup> Flexible entangled state generation in linear optics by Brendan Pankovich *et al*, Orca Computing, October 2023 (20 pages).

<sup>185</sup> Quantum simulation of thermodynamics in an integrated quantum photonic processor by Frank H. B. Somhorst *et al*, December 2021 - March 2023 (20 pages).

<sup>186 &</sup>lt;u>20-Mode Universal Quantum Photonic Processor</u> by Caterina Taballione, June 2022 (9 pages).

# 3.3.6. CONCLUSION

The photon qubit sector is original compared with all the other sectors evaluated in this report. It presents both advantages and major challenges. Almost all of them revolve around the highly probabilistic nature of quantum operations with photons. The challenges are to create entangled states of photons that are as deterministic as possible, to lose as few photons as possible in the calculation circuits and to detect them correctly at the end of the run. It is very difficult to give an overall view of feasibility, given the large number of parameters involved.

The question arises as to whether it would be possible to create deterministic gates between photons in the optical domain, which could be done using cavity electrodynamics techniques (*cavity* QED, as in Gerhard Rempe's experiments at the MPQ in Munich, or with Rydberg super-atoms, as in Alexei Ourjoumtev's experiments at the Collège de France). These experiments work between two photonic qubits, but are very complex and are currently a long way from being scaled up. Finally, it should be noted that interactions between photons in the microwave domain work very well, and are the basis of the "cat qubits" discussed elsewhere.

# 3.4. SPINS IN SILICON

Electron spin qubits are a promising technology which, like the others, is subject to numerous variations. Its development is more recent than for qubits based on trapped ions and Josephson effect superconductors.

Qubits of this type use the spin orientation of electrons trapped in potential wells in a given direction, or spin holes corresponding to a missing electron in a structure, all under a static magnetic field. The idea of creating spin qubits came from Daniel Loss (University of Basel) and David DiVincenzo (then at IBM Research) in a paper published in 1997<sup>187</sup>. They proposed a two-qubit gate using electrical control of the tunnel barrier between neighboring quantum dots. A low trigger voltage creates a coupling between neighboring qubits. The design also proposed single-qubit gates. The concept was then extended to the use of pairs of quantum dot electron spins, one being the qubit itself, and the other capacitively coupled to the first and used to read out the first qubit. This measurement involves a conversion of spin to charge using a conductance measurement, generally with radio-frequency reflectometry using a microwave pulse, rather similar to superconducting qubits readout.

The first demonstrations of spin qubit in silicon were carried out in Australia in 2012 by Andrew Dzurak's group at UNSW. In 2016, they were produced using a semi-industrial manufacturing process by a team from CEA-Leti and CEA-IRIG in Grenoble<sup>188</sup>. This technique, known as Si-MOS, is derived from planar MOS and FDSOI.

Single-qubit gates are generated by exposing quantum dots to an oscillating magnetic field. Two-qubit gates are generated by lowering the potential barrier between two adjacent boxes with an electrical voltage. The qubit can be measured in various ways, the most common being by spin-to-charge conversion, often using a second electron paired with the one to be measured and whose spin is inverted due to Pauli exclusion principle<sup>189, 190</sup>.

<sup>187 &</sup>lt;u>Quantum computation with quantum dots</u> by Daniel Loss and David DiVincenzo, 1997 (20 pages).

<sup>188 &</sup>lt;u>A CMOS silicon spin qubit</u> by Romain Maurand, Maud Vinet, Marc Sanquer, Silvano De Fransceschi *et al.*, 2016 (12 pages).

<sup>189</sup> Rapid single-shot parity spin readout in a silicon double quantum dot with fidelity exceeding 99 % by Kenta Takeda *et al*, RIKEN and QuTech, npj Quantum Information, February 2024 (6 pages).

<sup>19</sup> Modeling and Experimental Validation of the Intrinsic SNR in Spin Qubit Gate-Based Readout and Its Impacts on Readout Electronics by Bagas Prabowo, Lieven M. K. Vandersypen *et al*, QuTech, December 2023 (16 pages).

#### quantum dots spin qubits key takeaways

<ul> <li>highlights</li> <li>qubits fidelity: reaching 99.5% for two-qubit gates in labs for a small number of qubits (RIKEN, 2022).</li> <li>simpler qubit drive than with superconducting qubits.</li> <li>good quantum gates times: 75 ns for two-qubit gates.</li> <li>qperating temperature: around 100 mK - 1K =&gt; larger cooling budget for control electronics vs superconducting qubits.</li> <li>2D architectures: usable with surface codes or color codes QEC.</li> <li>manufacturing: can leverage existing semiconductor fabs.</li> <li>cryo-CMOS developments: for scalable qubit drive.</li> <li>scalability potential; could reach millions of qubits, thanks to their size of 100x100 nm.</li> </ul>	<ul> <li>challenges</li> <li>research in the field: started later than with other qubit technologies and spread over several technologies (full Si, SiGe, atom spin donors).</li> <li>lesser funded startups: and Intel relatively modest investment.</li> <li>qubit addressing and unsetlled qubit drive techniques (ESR,).</li> <li>high fabs costs and long test cycles: 18 months average.</li> <li>only 4 to 15 entangled qubits (QuTech, UNSW, Princeton, University of Tokyo).</li> <li>long-distance coupling between qubits, which enables more efficient QEC like qLDPC.</li> <li>high qubits variability: requires calibration.</li> <li>charge noise and other sources of noise to address, and scalability</li> </ul>
<ul> <li>variations</li> <li>SiGe qubits: spins or holes, more complicated to manufacture but better gate fidelities.</li> <li>donor spins: phosphorus atom nucleus, more complicated to scale.</li> <li>carbon nanotubes : with better spin stability but harder to manufacture and drive, with only one startup (C12).</li> <li>spin on superfluid helium or neon with only one startup in the domain (EeroQ).</li> </ul>	that remains to be demonstrated. path to scalability  materials purity improvements: isotopic and element impurities.  manufacturing: improvements, faster cycling and characterization.  full-stack EDA for digital simulation.  long range qubit coupling within chips: to enable efficient qLDPC QEC.  SFQ electronics for qubit drive.  cryo-CMOS integration with qubits: for better integration.  inter-QPU connectivity: electron spin shuttling, hole-microwave photons coupling, color centers.

Figure 13: Strengths, challenges, varieties and scalability options of spin-based qubits in quantum boxes.

#### 3.4.1. ADVANCES

Here are a few recent advances or highlights in spin qubits.

The fidelity of quantum gates can reach 99.5% for two-qubit gates in the laboratory for a small number of qubits, according to a 2022 experiment carried out by RIKEN in Japan<sup>191</sup>. However, this fidelity is questionable, as it has never been obtained on a scale of several tens of qubits and measured using randomized benchmarks.

Qubit control is simpler than with superconducting qubits. Microwave pulses are still needed to generate gates.

<sup>191</sup> Fast universal quantum control above the fault-tolerance threshold in silicon by Akito Noiri, Giordano Scappucci *et al.*, 2022 (27 pages).

This is the case for single-qubit gates and for reading qubits, but two-qubit gates are simpler to implement with voltage control.

Quantum gates times is competitive compared with that of other types of qubits, with, for example, 75 ns for a two-qubit gate in spin-holes SiGe<sup>192</sup>.

Operation at a higher temperature than superconducting qubits, between 100 mK and 1.5 K<sup>193</sup>. This provides a larger cooling budget for equivalent power consumption in cryogenics, making it possible to accommodate for a larger quantity of low-temperature control electronics and therefore to control a larger number of physical qubits.

The 2D architecture of the circuits means they can be used with surface codes. This characteristic is like that of superconducting qubits. However, it has not yet been validated experimentally.

Capitalizing on the nanoelectronics industry's manufacturing experience with CMOS components, which are similar to those used for spin qubits. This experience is highlighted by all the industrial players but should be taken with a grain of salt. We do not yet have experimental qubits of sufficient quality and quantity to rival superconducting qubits, for example.

The development of low-temperature electronics such as cryo-CMOS and SFQ means that we can envisage a good scale-up at a lower cost, especially as the cryogenics budget available is higher to operate these components, particularly cryo-CMOS components, because of the higher operating temperature than for superconducting qubits.

Greater miniaturization capacity, with qubit sizes of the order of 100×100 nm, theoretically supporting the integration of millions of physical qubits on a single chip<sup>194, 195</sup>. This is a key point for fault tolerance, to delay as far as possible the need to use interconnection technologies between quantum chips.

<sup>192</sup> Fast two-qubit logic with holes in germanium by N.W. Hendrickx, Menno Veldhorst, Giordano Scappucci et al, January 2020 in Nature et on arXiv in April 2019 (6 pages).

Hotter is easier: unexpected temperature dependence of spin qubit frequencies by Brennan Undseth, Lieven M. K. Vandersypen *et al*, April 2023 (17 pages).

This miniaturization could ultimately have a positive economic impact, enabling the development of scalable quantum computers at lower cost. This is more important as operational fault-tolerant computing solutions for solving complex problems, for example in quantum chemistry, will require the same quantum circuits to be executed on several similar quantum processors in parallel.

### 3.4.2. CHALLENGES

Research into spin qubits started later than with other qubit technologies (cold atoms, trapped ions, superconductors), and spans several competing industrial avenues (Si, SiGe, atomic spin donors) with relatively few players.

There are fewer start-ups, and they are less well-funded than for other types of qubit, and Intel, the biggest player in the field, is devoting a relatively modest investment to them, at least compared with the corresponding investments by IBM and Google in superconducting qubits, or even start-ups such as lonQ and Quantinuum in trapped ions, and PsiQuantum and Xanadu in photons.

The techniques for addressing and controlling qubits are not yet well established, especially for large-scale operation<sup>196</sup>.

*Wafer* manufacturing costs are high, and test cycles are fairly long, averaging around 18 months. This is the technology with the longest design, manufacturing and test cycles. This has an indirect impact on the speed at which technology matures, despite its promise. Industrial players are using various methods to shorten these cycles: in-house manufacturing in research cleanrooms with dedicated staff (SemiQon at VTT), use of industrial manufacturing lines (Diraq, Quantum Motion and Equal1 at GlobalFoundries, Quobly at STMicroelectronics, and ARQUE at Infineon), testing of several designs on the same *wafers* and the creation of digital twins to simulate the circuits.

The number of controlled and entangled qubits is currently limited to 4 to 15 qubits (QuTech, UNSW, Princeton, University of Tokyo). Scaling up has not yet been demonstrated experimentally.

 <sup>194</sup> The path to scalable quantum computing with silicon spin qubits
 by Maud Vinet, Nature

 Nanotechnology, December 2021.
 Image: Spin qubits
 Spin qubits
 Spin qubits

<sup>195</sup> Scaling silicon-based quantum computing using CMOS technology: State-of-the-art, Challenges and Perspectives by M. F. Gonzalez-Zalba, Silvano de Franceschi, Tristan Meunier, Maud Vinet, Andrew S. Dzurak *et al*, *Nature Electronics*, November 2020-April 2023 (21 pages).

<sup>196</sup> A Crossbar Network for Silicon Quantum Dot Qubits by R Li et al, 2017 (24 pages).

The high variability of qubits, which requires calibration in the same way as superconducting qubits. In addition, qubit readout fidelity, which is currently around 99%, needs to be improved

Load noise and other sources of noise affect the performance of qubits. Work is continuing on design, materials and manufacturing to reduce them<sup>197</sup>.

### 3.4.3. VARIATIONS

Here are the main technological variations of spin qubits in silicon and, by extension, in exotic graphene and other structures.

Silicon-germanium. These are silicon, germanium and silicon-germanium heterostructures in which germanium is used to stabilize hole spins<sup>198,199</sup>, wide band gaps, improved electron mobility, better spin-orbit momentum coupling, and long coherence times. Hole spins are managed with germanium quantum boxes between two SiGe layers, while spin qubits use a silicon box between the same SiGe layers. However, these qubits are more difficult to fabricate, with gates that can be controlled at a distance from the qubits<sup>200</sup>, and with various sources of noise that have yet to be contained<sup>201, 202</sup>.

Spin-donor phosphorus atoms. This technique was proposed by Bruce Kane (UNSW) in Australia in 1998. It involves using phosphorus atoms (<sup>31</sup>P) in a silicon crystal structure<sup>203</sup>. This hybrid approach uses nuclear magnetic resonance to control the spin of the nucleus of these atoms and couple them to the spin of one of the electrons in the valence layer of the phosphorus atom, which is not linked to the neighboring silicon atoms. The qubits are controlled by electric and magnetic fields<sup>204</sup>. The main benefit lies in the long coherence time of the phosphorus nucleus spins that make up the qubits, which can last for several seconds<sup>205</sup>. The challenges lie in placing individual atoms on the silicon *wafer*, a technique now mastered by the Australian start-up SQC, which uses a scanning tunneling microscope and the implementation of one- and two-qubit quantum gates.

<sup>197 &</sup>lt;u>Decoherence of solid-state spin qubits: a computational perspective</u> by Mykyta Onizhuk, and Giulia Galli, University of Chicago, arXiv, May 2024 (25 pages).

<sup>198 &</sup>lt;u>Recent advances in hole-spin qubits</u> by Yinan Fang *et al*, October 2022 (46 pages).

<sup>199</sup> Coherent control of a high-orbital hole in a semiconductor quantum dot with near-unity fidelity by Junyong Yan et al, December 2022 (27 pages).

<sup>200 &</sup>lt;u>The germanium quantum information route</u> by Giordano Scappucci, Silvano De Franceschi *et al*, 2020 (18 pages).

<sup>201 &</sup>lt;u>Simulation of 1/f charge noise affecting a quantum dot in a Si/SiGe structure</u> by Marcin Kępa *et al*, March 2023 (7 pages).

Spatial noise correlations beyond nearest-neighbor in 28Si/SiGe spin qubits by Juan S. Rojas- Arias, Daniel Loss *et al*, February 2023 (11 pages).

A silicon-based nuclear spin quantum computer by Bruce Kane, Nature, 1998 and Silicon-based Quantum Computation by Bruce E. Kane, 2000 (14 pages).

Carbon nanotubes, the technique proposed by the French start-up C12, which provides better spin stability at the cost of slightly more complex control of qubits and two-qubit gates<sup>206, 207</sup>. The Australian start-up Archer Materials had been proposing the development of qubits based on carbon nanospheres for several years. It seems to have abandoned this approach to quantum computing and is now focusing on the creation of quantum sensors for medical applications.

Spin on superfluid helium<sup>208</sup>or neon<sup>209</sup>, a relatively exotic technique being explored by laboratories such as RIKEN in Japan and the start-up EeroQ in the USA.

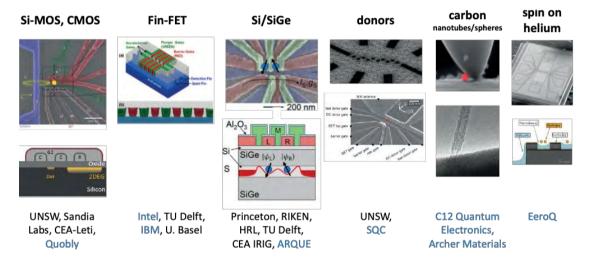


Figure 14: Different types of quantum box spin qubit. (cc) Olivier Ezratty, 2022-2023.

- 206 <u>Carbon Nanotube Devices for Quantum Technology</u> by Andrey Baydin *et al*, MDPI, February 2022 (26 pages).
- 207 Long-lived electronic spin qubits in single-walled carbon nanotubes by Jia-Shiang Chen *et al*, *Nature communications*, February 2023 (8 pages).
- Blueprint for quantum computing using electrons on helium by Erika Kawakami et al, RIKEN, QunaSys, OIST, and DLR, March 2023-September 2023 (28 pages).
- 209 <u>Single electrons on solid neon as a solid-state qubit platform</u> by Xianjing Zhou, Kater W. Murch, David I. Schuster *et al*, *Nature*, May 2022 (16 pages).

<sup>204 &</sup>lt;u>The Race To Make Better Qubits</u> by Katherine Derbyshire, Semiconductor Engineering, November 2021.

<sup>205 &</sup>lt;u>Toward a Silicon-Based Nuclear-Spin Quantum Computer</u> by Robert G. Clark, P. Chris Hammel, Andrew S. Dzurak, Alexander Hamilton, Lloyd Hollenberg, David Jamieson, and Christopher Pakes, Los Alamos Science, 2022 (18 pages).

### 3.4.4. PATHWAYS TO SCALABILITY

Here are the ways in which academic and industrial players in the spin qubit field are planning to scale up:

Improving the purity of materials, starting with the now established use of the <sup>28</sup>Si in FD-SOI to create quantum dots. Other work is focusing on the purity of the materials used<sup>210</sup>, which makes it possible to reduce charge noise<sup>211,212</sup>.

The development of industrial manufacturing techniques using the cleanrooms of GlobalFoundries, Intel, STMicroelectronics and, potentially, TSMC, which has set up an exploratory team on this subject in 2023.

The creation of '*full-stack*' electronic design and simulation solutions for circuits. This is the approach adopted by Quobly.

The coupling of long-range qubits to enable the use of qLDPC-type error correction codes, which are more efficient and require fewer physical qubits per logical qubit<sup>213</sup>.

The integration of cryo-CMOS control components close to the qubit chip, with chiplet-type solutions, while optimizing energy consumption and the capacity of cryostats adapted to operating temperatures of between 100 mK and 1 K.

<sup>210 &</sup>lt;u>Materials for Silicon Quantum Dots and their Impact on Electron Spin Qubits</u> by Andre Saraiva, Wee Han Lim, Chih Hwan Yang, Christopher C. Escott, Arne Laucht and Andrew S. Dzurak, December 2021 (22 pages).

<sup>211</sup> Low charge noise quantum dots with industrial CMOS manufacturing by Asser Elsayed *et al*, IMEC, December 2022 (22 pages).

<sup>212 &</sup>lt;u>Stabilizing an individual charge fluctuator in a Si/SiGe quantum dot</u> by Feiyang Ye, Ammar Ellaboudy, and John M. Nichol, University of Rochester, *arXiv*, July 2024 (6 pages).

<sup>213 &</sup>lt;u>Coherent Spin-Spin Coupling Mediated by Virtual Microwave Photons</u> by Patrick Harvey-Collard, Jurgen Dijkema, Guoji Zheng, Amir Sammak, Giordano Scappucci, and Lieven M. K. Vandersypen, QuTech, Physical Review X, May 2022 (16 pages).

The development of superconducting SFQ electronics to drive qubits, which is currently the preserve of a single company, SeeQC (USA). This could enable spin qubits to be scaled up with a limited cryogenics budget.

The development of interconnection techniques between quantum processors, using flying electrons<sup>214</sup>, hole-microwave photon couplings, spin couplings with coloured centres (NV centre type) and, finally, the use of optical photons. Flying selectron tests were achieved in 2024 over a distance of 10  $\mu$ m in <200 ns with an average fidelity of 99%<sup>215</sup>.

# 3.4.5. ROADMAPS

At this stage, the vast majority of industrial players in these types of qubits have published *roadmaps* which are mainly NISQ, with cautious and not very quantified developments envisaged towards FTQC:

Intel (USA) is betting on silicon-germanium qubits, which it is developing with the help of QuTech in the Netherlands. Their latest 12-qubit chip, Tunnel Falls, dates from June 2023. It was supplied to mainly American universities for integration and testing, reflecting Intel's OEM strategy in the same way as for its conventional processors<sup>216</sup>. The chip was etched on 300mm *wafers* at Intel's Hillsboro site in Oregon, using EUV ASML lithography. Intel made several presentations at the APS March meeting 2024 in Minneapolis, where it revealed that the fidelities of two-qubit gates

<sup>214 &</sup>lt;u>Coherent shuttle of electron-spin states</u> by Lieven Vandersypen *et al*, 2017 (21 pages).

<sup>215 &</sup>lt;u>High-fidelity single-spin shuttling in silicon</u> by Maxim De Smet, Lieven M.K. Vandersypen *et al*, QuTech and TU Delft, *arXiv*, June 2024 (15 pages).

<sup>216</sup> Characterization of individual charge fluctuators in Si/SiGe quantum dots by Feiyang Ye, John M. Nichol *et al*, University of Rochester and Sandia Labs, *arXiv*, January 2024 (37 pages) is about characterizing a 4-quantum dot circuit, seemingly not coming from Intel, or at least, not being Tunnel Falls.

of this chip were 92%, which is not very satisfactory<sup>217</sup>. Intel is also developing cryo-CMOS control chips, including Horse Ridge 2 in 2021, which operates at 4 K and generates all the control signals and functions for reading the state of the silicon qubits. In 2024, it will complement this chip with Pando Tree, which operates at the temperature of the qubits and enables control signals from Horse Ridge 2 to be demultiplexed, with a significant gain in control lines. Intel has not published an FTQC *roadmap*.

Quantum Motion (UK) is working in parallel on its spin qubits and cryo-CMOS control components. In 2021, together with Hitachi Cambridge, the University of Cambridge and EPFL, it presented a chip operating at 50 mK comprising both qubits and control signal routing electronics<sup>218</sup>. It completed this in 2024 with a modular qubit chip architecture and a cryo-CMOS low-noise amplifier for reading the state of the qubits<sup>219</sup>. It unveiled its architecture in August 2022<sup>220</sup> followed by its 3× 3mm<sup>2</sup> Bloomsbury chip manufactured at Global Foundries with 1024 *quantum dots*<sup>221</sup>. In June 2023, it explained how it could efficiently execute the same quantum circuit several times in parallel<sup>222</sup>. Instead of arranging N qubits in a matrix of N $\sqrt{x}$  $\sqrt{N}$  qubits, she plans to create a matrix of N× D, D being the number of times a circuit must be executed. She plans to create an FTQC computer with 100 logic qubits by 2029.

<sup>217</sup> Intel Presents 12 Quantum Research Papers at APS March Meeting 2024 by Intel, March 2024. Probably in "High-fidelity Operation of Encoded Spin Qubits on Intel Tunnel Falls" by Felix F Borjans but its content is not public.

<sup>218</sup> Integrated multiplexed microwave readout of silicon quantum dots in a cryogenic CMOS chip by A. Ruffino, Edoardo Charbon *et al*, Quantum Motion, Hitachi and EPFL, January 2021 (14 pages).

<sup>219 &</sup>lt;u>A multi-module silicon-on-insulator chip assembly containing quantum dots and cryogenic radio-frequency readout electronics by David J. Ibberson, James Kirkman, John J. L. Morton, M. Fernando Gonzalez-Zalba, and Alberto Gomez-Saiz, Quantum Motion and UCL, *arXiv*, May 2024 (3 pages).</u>

<sup>20</sup> Silicon edge-dot architecture for quantum computing with global control and integrated trimming by Michael A. Fogarty, August 2022 (13 pages).

<sup>221 &</sup>lt;u>Rapid cryogenic characterisation of 1024 integrated silicon quantum dots</u> by Edward J. Thomas *et al, Quantum Motion and UCL*, October 2023 (22 pages).

<sup>222 &</sup>lt;u>Pipeline quantum processor architecture for silicon spin qubits</u> by S. M. Patomäki, Simon Benjamin et al, *npj Quantum Information*, March 2024 (10 pages).

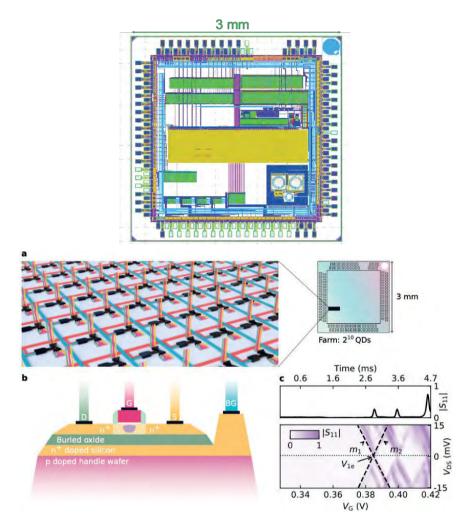


Figure 15: Quantum Motion's Bloomsbury chip manufactured by GlobalFoundries. Source: Quantum Motion, Q2B Santa Clara, December 2023.

Quobly (France) has a strategy of creating in parallel its qubit chips and lowtemperature cryo-CMOS control electronics, co-developed with CEA/LIST. The qubit chips will initially be 1D, before moving on to 2D. Three development phases are planned as of May 2025: 256 physical qubits by 2027, 20K physical qubits by 2030 enabling one billion operations and then over 1K logical qubits by 2032. Diraq (Australia) has set itself the ambitious goal of creating a quantum computer with a billion physical qubits, in stages starting with 9 and then 256 qubits. It plans to adopt the SiMOS technique similar to Quobly's, but is exploring other avenues in parallel, such as multi-electron qubits<sup>223</sup>, or even spin-donor qubits using antimony atoms (<sup>123</sup>Sb)<sup>224</sup>. It has published several advances on the fidelity of qubit initialization<sup>225</sup>, on the variability of quantum gate fidelity<sup>226</sup>, on *benchmarking*<sup>227</sup>, on qubit control<sup>228, 229, 230</sup>, on qubit reading with a parametric amplifier<sup>231, 232</sup> and on an error suppression technique<sup>233</sup>. In December 2022, it achieved fidelities of 99.96% for one-qubit gates<sup>234</sup> and 98.92% for two-qubit gates, based on an unspecified, and therefore very low, number of qubits<sup>235</sup>.

<sup>223 &</sup>lt;u>Electronic Correlations in Multielectron Silicon Quantum Dots</u> by Dylan H. Liang, MengKe Feng, Philip Y. Mai, Jesus D. Cifuentes, Andrew S. Dzurak, and Andre Saraiva, Diraq and UNSW, *arXiv*, July 2024 (6 pages).

<sup>224 &</sup>lt;u>Creation and manipulation of Schrödinger cat states of a nuclear spin qudit in silicon</u> by Xi Yu, Andrew S. Dzurak, Andrea Morello *et al*, Diraq *et al*, *arXiv*, May 2024 (40 pages).

<sup>225 &</sup>lt;u>Beating the Thermal Limit of Qubit Initialization with a Bayesian Maxwell's Demon</u> by Mark A. I. Johnson, Kohei M. Itoh, Andrew S. Dzurak, Andrea Morello *et al*, *PRX*, October 2022 (15 pages).

<sup>226</sup> Bounds to electron spin qubit variability for scalable CMOS architectures by Jesús D. Cifuentes, Andrew S. Dzurak et al, March 2023 (20 pages).

<sup>227 &</sup>lt;u>Stability of high-fidelity two-qubit operations in silicon by</u> Tuomo Tanttu, Kohei M. Itoh, Robin Blume-Kohout, Andrea Morello, Andrew S. Dzurak, March 2023 (13 pages).

<sup>228 &</sup>lt;u>On-demand electrical control of spin qubits</u> by Will Gilbert, Kohei M. Itoh, Andrea Morello, Andrew S. Dzurak *et al*, *Nature Nanotechnology*, January 2023 (21 pages).

<sup>229</sup> Implementation of an advanced dressing protocol for global gubit control in silicon by I. Hansen, Kohei M. Itoh, Andrew S. Dzurak et al, Applied Physics Reviews, September 2022 (9 pages).

Implementation of the SMART protocol for global qubit control in silicon by Ingvild Hansen, Andrew S. Dzurak et al, August-September 2021 (9 pages).

<sup>231 &</sup>lt;u>Direct detection of spin resonance with a microwave parametric amplifier</u> by Wyatt Vine, Andrea Morello *et al*, November 2022 (28 pages).

<sup>232 &</sup>lt;u>Gate-based spin readout of hole quantum dots with site-dependent q-factors by Angus Russell, Andrew S. Dzurak, Alessandro Rossi *et al*, June 2022-April 2023 (16 pages).</u>

<sup>233 &</sup>lt;u>Real-time feedback protocols for optimizing fault-tolerant two-qubit gate fidelities in a silicon spin system by Nard Dumoulin Stuyck, Andrew S. Dzurak *et al*, September 2023 (6 pages).</u>

<sup>234 &</sup>lt;u>Diraq achieves record accuracy for device manufactured by existing semiconductor</u> <u>infrastructure</u> by Diraq, June 2024.

<sup>235 &</sup>lt;u>High-fidelity spin qubit operation and algorithmic initialization above 1 K</u> by Jonathan Y. Huang, Natalia Ares, Andrew S. Dzurak, Chih Hwan Yang *et al*, *Nature*, August 2023-March 2024 (20 pages).

Equal 1 (Ireland-USA) is working on the creation of a silicon-germanium qubit chip that will eventually include millions of qubits, as well as the associated control electronics, including error correction. In practice, its Quantum Systemon-chip (QSoC) will be a component incorporating a chip for the qubits<sup>236, 237</sup> on one side and a cryo-CMOS chip for controlling the qubits on the other. For the moment, it is using TNO's research production line in the Netherlands for the qubits and GlobalFoundries' production line in Dresden, Germany, for the cryo-CMOS chip. The quantum computer that will incorporate this component will fit into a 4U server format. Its target figures of merit are two-qubit gate fidelities of 98% and one-qubit gate fidelities of 99%, with gates lasting 50 to 100 ns.

and a read time of 10 µs, which can then be reduced to 500 ns.

SemiQon (Finland) is based at VTT's premises in Espoo, near Helsinki. Its aim is to create a chip with one million spin qubits, incorporating cryo-CMOS control electronics for the qubits<sup>238</sup>. Its April 2023 prototype comprised 48 qubits with associated control electronics<sup>239</sup>.

Arque (Germany) has a *roadmap* comprising 2-4, 50, 200, 10,000 then 1 million qubits of the GaAs then SiGe types. It plans to obtain two-qubit gate fidelities of 99.9%<sup>240, 241</sup>.

A Single-Electron Injection Device for CMOS Charge Qubits Implemented in 22-nm FD-SOI by Imran Bashir, Elena Blokhina *et al*, 2020 (4 pages).

Xanoscale single-electron box with a floating lead for quantum sensing: modelling and device

 characterization
 by Nikolaos Petropoulos, Elena Blokhina *et al, arXiv*, April 2024 (7 pages).

Scalable on-chip multiplexing of silicon single and double quantum dots by Heorhii Bohuslavskyi, Alberto Ronzani, Joel Hätinen, Arto Rantala, Andrey Shchepetov, Panu Koppinen, Mika Prunnila, and Janne S. Lehtinen, arXiv, August 2022-December 2023 (30 pages).

<sup>239</sup> SemiQon's quantum processor testing and measurement facilities at the VTT MIKES premises up and running, Semiqon, April 2024. The first tested chip has 4 qubits, but no specific gate set at this point.

<sup>240 &</sup>lt;u>The SpinBus Architecture: Scaling Spin Qubits with Electron Shuttling</u> by Matthias Künne *et al*, JARA and ARQUE, June 2023 (15 pages).

<sup>241 &</sup>lt;u>Scalable Parity Architecture With a Shuttling-Based Spin Qubit Processor</u> by Florian Ginzel, Michael Fellner, Christian Ertler, Lars R. Schreiber, Hendrik Bluhm, and Wolfgang Lechner, *arXiv*, March 2024 (17 pages).

SQC (Australia) uses the spin donor atom technique, exploiting phosphorus atoms implanted in a silicon <sup>28</sup>Si substrate <sup>242</sup> and whose nucleus is coupled to the spin of the five electrons in the valence layer of the atom that is not linked to the neighboring silicon atoms by covalent bonds. She has created two- qubit gates with two phosphorus atoms 17.5 nm apart<sup>243</sup>. In 2020, it achieved 99.99% qubit gate fidelity<sup>244</sup>. In 2022, it presented a 10-qubit processor in 1D. In 2024, it executed a Grover algorithm with 4 qubits, organized with three phosphorus atoms and one electron spin with a two-qubit gate fidelity of 99%<sup>245</sup>. The company has not published an FTQC *roadmap*.

C12 (France) is working on the use of carbon nanotubes to trap electron spins in 1D. The technique has the advantage of protecting the qubits from external interference and improving their coherence time. The carbon nanotubes are made from integer-spin <sup>12</sup>C carbon, which does not interfere with the spin of the electron trapped inside. They are integrated on a silicon circuit using automated staplers<sup>246</sup>. The qubits are interconnected by microwave cavities. The company has its own facilities in Paris for manufacturing its carbon nanotubes and the silicon circuit that houses them.

<sup>242 &</sup>lt;u>Highly 28Si enriched silicon by localised focused ion beam implantation</u> by Ravi Acharya, David N. Jamieson, Richard J. Curry *et al*, University of Melbourne and University of Manchester, *Communications Materials*, May 2024 (7 pages) shows a <sup>28</sup>Si purification technique on a silicon substrate that is localized on the implanted phosphorus atom location.

<sup>243 &</sup>lt;u>Measurement of enhanced spin-orbit coupling strength for donor-bound electron spins in silicon</u> by Radha Krishnan, Michelle Y. Simmons, Bent Weber *et al*, NTU Singapore, UNSW, *arXiv*, April 2024 (8 pages).

<sup>24</sup> Exploiting a Single-Crystal Environment to Minimize the Charge Noise on Qubits in Silicon by Ludwik Kranz, Michelle Simmons et al., 2020 and <u>A two-qubit gate between phosphorus donor</u> <u>electrons in silicon</u> by Y. He, Michelle Simmons et al., 2019.

<sup>245 &</sup>lt;u>Grover's algorithm in a four-qubit silicon processor above the fault-tolerant threshold</u> by Ian Thorvaldson, Michelle Y. Simmons *et al*, *arXiv*, April 2024 (16 pages).

<sup>246 &</sup>lt;u>Nanoassembly technique of carbon nanotubes for hybrid circuit-QED</u> by Tino Cubaynes, Matthieu Desjardin, Audrey Cottet, Taki Kontos *et al*, September 2021 (6 pages).

EeroQ (USA) is developing a processor that traps electrons on a bed of superfluid helium above a silicon circuit containing control electronics<sup>247</sup>. The idea came from the University of Michigan and Bell Labs in 1999<sup>248, 249, 250</sup>. It was then developed between 2003 and 2006 to exploit electron spins with circuits resembling ion traps and exploiting shuttle electrons to create two-qubit gates<sup>251</sup>. Its first test chip in 2023 theoretically supported 2,432 qubits, all requiring just 30 control lines. The electrons fly over the superfluid helium at an altitude of around 10 nm. They are controlled by electrodes located in a CMOS chip. EeroQ plans to achieve coherence times of 10 s, and two-qubit gates with fidelities of 99.9%<sup>252</sup>. All this with the capacity to reach 10,000 qubits per chip.

### 3.4.6. CONCLUSION

Electron spin qubits hold as many promises as they do challenges. Their level of maturity is much lower than that of superconducting qubits using cold atoms and trapped ions. The support of the microelectronics industry, which is often put forward by players in this market, is a plus, but not enough. There are many difficulties to be overcome in the fields of quantum physics and materials. Resolving them, which will certainly take time, could open interesting avenues in terms of scaling up.

<sup>247 &</sup>lt;u>Electron-on-helium qubit</u> page on Wikipedia.

<sup>248</sup> Quantum Computing with Electrons Floating on Liquid Helium by P. M. Platzman and M. I. Dykman, Science, June 1999 (3 pages).

<sup>249 &</sup>lt;u>Quantum computing using floating electrons on cryogenic substrates: Potential And Challenges</u> by Ash Jennings *et al*, RIKEN, October 2023 (25 pages).

Integrating superfluids with superconducting qubit systems by Johannes Pollanen, Kater Murch et al,

 Michigan State University, Washington University in Saint Louis, 2019 (11 pages).

<sup>251</sup> Spin-based quantum computing using electrons on liquid helium by Steve Lyon, PRA, 2003- 2006 (12 pages).

<sup>252 &</sup>lt;u>Coulomb interaction-driven entanglement of electrons on helium</u> by Niyaz R. Beysengulov, Johannes Pollanen *et al*, October 2023 (19 pages).

# 3.5. TRAPPED IONS

Trapped-ion qubits are among the oldest around, with the first experiments achieved in the mid-1990s. The scientific community in this field is very large and covers every continent, including the USA, Europe, Taiwan and China. The companies in the sector are all start-ups, including lonQ and Quantinuum in the USA, Universal Quantum and Oxford Ionics in the UK, eleQtron and NeQxt in Germany, AQT in Austria and Crystal Quantum Computing in France among others.

The ions in these quantum computers are trapped above electronic circuits containing electrodes that enable them to be positioned precisely using electric fields, in addition to the lasers that are used to cool them initially. They are organized in one- or two-dimensional structures, depending on the case. The  $|0\rangle$  and  $|1\rangle$  states of these ions correspond to energy levels separated by frequencies that determine how they are controlled. The most common are ions with a hyperfine structure controlled by microwaves of a few GHz and laser-based Raman transitions<sup>253</sup>. Next come ions controlled by optical frequencies of a few hundred THz and, more rarely, ions excited to Rydberg states by lasers emitting in the ultraviolet range.

The ions are chosen according to their electronic configuration and their energy transitions. They are generally alkali metals from group 2 or IIA of Mendeleev's table, including beryllium, magnesium, calcium, strontium, barium and certain rare earths such as ytterbium<sup>254, 255</sup>. Some systems use two elements, such as calcium and strontium, or barium and ytterbium, one for cooling and the other for computing qubits.

<sup>253 &</sup>lt;u>Robust and resource-efficient microwave near-field entangling <sup>2</sup>Be+ gate by G.</u> Zarantonello, November 2019 (6 pages).

<sup>254</sup> Introduction to Trapped Ion Quantum Computing by Gabriel Mintzer from MIT, February 2020.

<sup>255 &</sup>lt;u>Ion-Based Quantum Computing Hardware: Performance and End-User Perspective</u> by Thomas Strohm, Sebastian Luber *et al*, Bosch, *arXiv*, May 2024 (44 pages).

lons provide long coherence times of up to several tens of seconds, quantum gates with the best-known fidelities - in excess of 99.9% for two-qubit gates - and better connectivity between qubits, as two-qubit gates can be created between any qubit in a set of ions. Their main drawbacks are the slowness of their gates and various difficulties in scaling up.

Spatial stabilization of the ions is provided by ion traps. There are two main types: Penning traps, which use a magnetic field and an electric quadrupole, and Paul traps, which use an oscillating electric field. Positioning and control of the ions commonly combine direct currents and microwave fields. Lasers are used to precool the ions using the Doppler effect and *sideband cooling*, to limit the phononic effects of ion vibration and to read out the state of the qubits using fluorescence. Quantum gates are created using microwave or laser fields, depending on the case<sup>256, 257</sup>.

The most common ion traps are now QCCDs (*Quantum Charge-Coupled Devices*). They allow ions and their control fields to be distributed in different zones in 1D and 2D<sup>258</sup>. The zones are linked together with ions that circulate from one zone to another (*shuttling ions*)<sup>259</sup>. Players in this market generally rely on commercial cleanrooms such as those of GlobalFoundries and Infineon. They use the same argument as the silicon qubit players: the use of professional cleanrooms in the microelectronics industry.

Individually Addressed Quantum Gate Interactions Using Dynamical Decoupling by M.C. Smith,

 A.D. Leu, M.F. Gely, and D.M. Lucas, PRX Quantum, August 2024 (9 pages).

<sup>257</sup> Quantum information processing with trapped ions by Christian Roos, 2012 (53 slides).

<sup>258 &</sup>lt;u>Architecture for a large-scale ion-trap quantum computer</u> by D. Kielpinski, C. Monroe, D. J. Wineland, *Nature*, June 2002 (3 pages).

Multi-zone trapped-ion qubit control in an integrated photonics QCCD device by Carmelo Mordini, Jonathan P. Home *et al*, ETH Zurich, *arXiv*, January 2024 (13 pages).

#### trapped ions qubits key takeaways

highlights	challenges
<ul> <li>first two-qubit gate fidelities reaching 99.9% (Quantinuum, Oxford lonics).</li> </ul>	<ul> <li>unproven scalability options beyond 60 qubits (ions shuttling, 2D architectures, photon interconnect, micro-Penning traps).</li> </ul>
<ul> <li>first logical qubits above break-even.</li> <li>high ratio between coherence time and gate time: supports deep algorithms in number of gate cycles.</li> <li>low qubits variability given the ions are all the same.</li> <li>entanglement possible between all qubits on 1D architecture which speeds up computing, avoiding SWAP gates.</li> <li>requires some cryogeny at 4K to 10K.</li> <li>QPU interconnect can directly use photon entangled resources.</li> </ul>	<ul> <li>slow computing: due to long quantum gate times and ions shuttling which may be problematic for deep algorithms in a FTQC regime despite better qubit many-to-many connectivity at small scale.</li> <li>two-qubit gate times increase with ion distance in some laser-driven 1D and 2D settings.</li> <li>many-to-many connectivity demonstrated only at small scale.</li> <li>control signals variability: microwave, lasers, etc.</li> <li>ions heating phenomenon: it is not yet explained yet and really contained.</li> </ul>
<ul> <li>variations</li> <li>microwave/DC drive instead of laser drive.</li> <li>connectivity: many-to-many within zones and ions shuttling between zones.</li> <li>dual species like ytterbium + barium for computing and cooling.</li> <li>Rydberg ion qubits for avoiding phonon and heating effect.</li> </ul>	path to scalability <ul> <li>2D QCCD and ions shuttling.</li> <li>QCCD tiling (Universal Quantum).</li> <li>multi modules ion traps with intermodules ions shuttling.</li> <li>multi-layer ion traps to enable long-range microwave based entanglement.</li> </ul>
hybrid neutral atoms-ions platforms.	photonic interconnect to entangle qubits from different QPUs.

Figure 16: Strengths, challenges, varieties and scalability options of trapped-ion qubits.

### 3.5.1. ADVANCES

Qubits based on trapped ions have been around for almost three decades. The pioneers were NIST in Boulder, the University of Innsbruck and then the University of Maryland.

The fidelities of quantum gates based on trapped ions are the best of all qubit modalities. They reach 99.99% for one-qubit gates and 99.9% for two-qubit gates, particularly at Quantinuum and for 20 qubits. In June 2024, Quantinuum achieved a record of 99.84% fidelity for two-qubit gates on 56 qubits<sup>260</sup>. This makes it possible to run NISQ algorithms that are relatively deep in their number of quantum gate cycles.

Trapped ions were the first with which logical qubits were created above the *break-even* point, i.e. with fidelities of of logical qubits better than that of the underlying physical qubits.

The computational power of random quantum circuits in arbitrary geometries by Matthew DeCross, Marco Pistoia, et al, arXiv, June 2024 (36 pages).

This is particularly true of Quantinuum, which is the most advanced in this respect<sup>261</sup>. In 2024, it achieved teleportation of logic qubits with a fidelity of  $97\%^{262}$ . It created three logical qubits with two-qubit gate fidelity of 99.91% on its 20-qubit H1-1 QPU and 99.8% on the 32-qubit H2-1<sup>263</sup>.

lons have a high ratio between their coherence time and the duration of the quantum gates. This is around 10<sup>6</sup> compared with 10<sup>3</sup> for superconducting qubits and 200 for cold atoms. This enables ions to support algorithms with a larger depth in terms of the number of gate cycles, without error correction, and in the NISQ regime.

Like atoms, ions have low variability. It depends essentially on the control circuit.

Trapped ions allow the creation of arbitrary quantum gates between pairs of qubits, avoiding the numerous SWAP gates that are otherwise necessary in qubit architectures where qubits are only connected to their nearest neighbor, such as superconducting qubits. However, this technique only works on a small scale of a few dozen qubits, with lonQ (39) and Quantinuum (56). The further apart the qubits are, the slower and poorer the quality of the quantum gates connecting them are. Beyond that, connectivity is only made possible by "shuttle ions" that are moved from one block of qubits to another and enable these blocks to be linked together by two-qubit gates. Peer-to-peer connectivity is therefore not always generalized to all the ions in a processor.

Benchmarking logical three-qubit quantum Fourier transform encoded in the Steane code on a trapped-ion quantum computer by Karl Mayer, Russell Stutz et al, Quantinuum, arXiv, April 2024 (13 pages).

<sup>262 &</sup>lt;u>High-fidelity and Fault-tolerant Teleportation of a Logical Qubit using Transversal Gates and Lattice Surgery on a Trapped-ion Quantum Computer by C. Ryan-Anderson, D. Hayes et al, Quantinuum, arXiv, April 2024 (13 pages).</u>

<sup>263 &</sup>lt;u>Benchmarking logical three-gubit quantum Fourier transform encoded in the Steane code on a trapped-ion quantum computer by Karl Mayer, Russell Stutz *et al*, Quantinuum, *arXiv*, April 2024 (13 pages).</u>

lons require cryogenics at 4 K, which is simpler to implement than for superconducting qubits at 15 mK. As with neutral atoms, cooling the vacuum chamber containing the ions improves the quality of the ultra-high vacuum and avoids the effects of thermal photons.

Finally, ions have the advantage of being relatively easy to entangle with photons to ensure interconnection between processors, since part of their control relies on photons in the visible or infrared range, based on lasers.

### 3.5.2. CHALLENGES

Here are the challenges to be overcome in order to make progress with these qubits in fault-tolerant architectures.

The scalability options have not yet been tested beyond around fifty qubits. They are all based on 2D QCCD circuits integrating blocks of ions with ions floating from one block to the next. The *blueprints* for this have been around for over two decades, but have yet to be implemented in practice.

The slow quantum gates are detrimental to computing speed, particularly in faulttolerant conditions where the circuits are deep with a large number of quantum gate cycles. This can range from one to a few hundred microseconds. This slowness is increased by the use of floating ions that are physically moved from one place to another in the circuits to establish entanglements between blocks of qubits. Trapped ions quantum gates can be around 1,000 times slower than with superconducting qubits. This is an enormous handicap in terms of computing speed, which is not necessarily compensated for by good connectivity between qubits. It could however be potentially compensated by a lower price and energetic footprint for the QPUs.

*Many-to-many* connectivity operates in a limited way in blocks of qubits comprising between 4 and 12 physical qubits. In certain types of microwave-based control circuits such as those used by Oxford Ionics, connectivity is currently limited to the nearest neighbor and in 1D.

Two-qubit gate times increase with the distance between ions in 1D and 2D circuits and laser ion control, as in lonQ.

Residual errors arise mainly from variations in the voltage or current of the ion control circuits, variations in the frequencies of the control lasers, the effects of optical crosstalk in the control circuits<sup>264, 265</sup>, inaccuracies in the optical tweezers<sup>266</sup>, and finally, the effect of ion heating, which remains partially unexplained. Ions are also sensitive to magnetic field variations, creating phase noise and collisions between ions and residual gases.

### 3.5.3. VARIATIONS

Here are some of the main technological variants concerning quantum computers based on trapped ions.

The qubits can be driven by lasers or by microwaves, radio frequencies and direct currents. Laser control is the historical solution chosen by lonQ, for example. It works well on a small scale, enables precise control of the ions, and is suitable for high-energy transitions and the generation of long-distance entanglements, but it is more difficult to operate on a large scale. Most of the other players, such as Quantinuum, Universal Quantum and Oxford lonics, prefer to drive their quantum gates with microwaves<sup>267</sup> or DC voltages, which are directed to the ions via QCCD circuits. This solution eliminates the need for some of the optical components required for laser control,

<sup>264 &</sup>lt;u>Low-Crosstalk, Silicon-Fabricated Optical Waveguides for Laser Delivery to Matter Qubits</u> by Clayton L. Craft, David Hucul *et al*, AFRL, HRL Laboratories, *arXiv*, June 2024 (9 pages).

<sup>265</sup> Physical coherent cancellation of optical addressing crosstalk in a trapped-ion experiment by Jeremy Flannery, Roland Matt, Luca Huber, Kaizhao Wang, Christopher Axline, Robin Oswald, and Jonathan P. Home, ETH Zurich, arXiv, June 2024 (11 pages).

<sup>266 &</sup>lt;u>Alignment and Optimisation of Optical Tweezers on Trapped Ions</u> by M. Mazzanti, R. X. Schüssler *et al*, QuSoft, *arXiv*, June 2024 (8 pages).

<sup>267 &</sup>lt;u>Robust and fast microwave-driven quantum logic for trapped-ion qubits</u> by M. A. Weber, D. M. Lucas et al, University of Oxford, arXiv, February 2024 (6 pages).

These are limited to cooling the ions and reading their state. Ion addressing is easier on a large scale and consumes less power.

Connectivity between qubits for executing two-qubit gates varies from one technology to another. Laser and microwave gate control generally allows arbitrary two-qubit gates to be constructed between all the pairs of qubits in a processor. In practice, this connectivity is limited to clusters of ions ranging from a few to around fifty. Beyond that, connectivity between clusters is provided with shuttling ions moving from one cluster to another to establish entanglement between adjacent ion clusters. These clusters can be organized in one dimension (lonQ) or in two dimensions on orthogonal rails (Quantinuum). The arbitrary connectivity between ions is therefore local to these clusters and not global to the whole of a computer managing several ion clusters.

The use of two elements such as ytterbium and barium, the first for qubit operations and the second for ions cooling. This is the technique used by Quantinuum. IonQ, for its part, is in the process of migrating from ytterbium to barium, because barium can be driven by photons in the 1,550 nm telecom band, which facilitates photonic interconnection between quantum processors.

The use of ions in Rydberg states to avoid the use of phonons and the effects of ion heating. This is the technique adopted by Crystal Quantum Computing (France). The price to pay is the absence of *many-to-many* connectivity, as Rydberg states only allow ions to be linked to their immediate neighbors.

Hybrid platforms using neutral atoms and ions are currently being researched in Germany<sup>268</sup>.

<sup>268 &</sup>lt;u>Emulating Solid-State Physics with a Hybrid System of Ultracold Ions and Atoms</u> by U. Bissbort, R. Gerritsma *et al*, PRL, August 2013 (8 pages)

### 3.5.4. PATHWAYS TO SCALABILITY

Here are the technological avenues chosen by the various players in the market for scaling up with trapped ions.

2D QCCD circuits and shuttle ions are the first way envisaged to scale up the number of physical qubits<sup>269, 270</sup>.

Multimodule computers with ion transfer between modules. This is the approach adopted by Universal Quantum. This transfer would take place with very high fidelity<sup>271</sup>. This requires optimized quantum code compilation techniques<sup>272</sup>.

Multilayer ion traps to enable long-range entanglement using microwaves. One solution being considered, which reduces the effects of heating, is to use superconducting circuits to deliver microwaves to drive the qubits, as RIKEN is studying in Japan<sup>273</sup>.

Blueprint for a microwave trapped ion quantum computer by Bjoern Lekitsch et al, Science Advances, February 2017 (11 pages). See also <u>A Shuttle-Efficient Qubit Mapper for Trapped Ion</u> <u>Quantum Computers</u> by Suryansh Upadhyay et al., April 2022 (7 pages).

High-Fidelity Transport of Trapped-Ion Qubits in a Multi-Layer Array by Deviprasath Palani, Tobias Schaetz et al, University of Freiburg, May 2023 (8 pages).

<sup>271 &</sup>lt;u>TITAN: À Distributed Large-Scale Trapped-Ion NISQ Computer</u> by Cheng Chu *et al*, Lawrence Berkeley National Laboratory, *arXiv*, February 2024 (6 pages) transports ions from one QCCD circuit to another in 0.4 ms and with a fidelity of 99.999 993%. It uses ytterbium <sup>171</sup>Yb+ ions for calculation and barium <sup>138</sup>Ba+ ions for transfers.

<sup>212 &</sup>lt;u>Shuttling for Scalable Trapped-Ion Quantum Computers</u> by Daniel Schoenberger, Stefan Hillmich, Matthias Brandl, and Robert Wille, Infineon, *arXiv*, February 2024 (7 pages).

<sup>213 &</sup>lt;u>Superconducting surface trap chips for microwave-driven trapped ions</u> by Yuta Tsuchimoto *et al*, University of Tokyo, RIKEN, Inamori Research Institute for Science, *arXiv*, July 2024 (27 pages).

Photonic interconnection to interleave the qubits of different QPUs. This is the lonQ strategy<sup>274, 275</sup>. Links of up to 400 meters have been tested at the University of Innsbruck in Austria, with calcium ions separated by 230 meters by a 520 m long fibre, which can be extended to 50 km<sup>276, 277</sup>. There is still a great deal of research to be done to achieve this, given the still low success rate in creating these entanglement-based photonic links, with Bell pairs obtained at a very low error rate of 2.18 × 10<sup>-4</sup> 278, 279</sup>.

# 3.5.5. ROADMAPS

Trapped ion vendors have almost all communicated their roadmap towards faulttolerant quantum computers. In general, they plan for a continuous transition from the NISQ to the FTQC regime, due to the fact that ions have good fidelities and make it possible to envision the use of quantum error mitigation techniques with more operations than with other qubit modalities.

lonQ published its *roadmap* in 2020. It has also adopted an in-house *benchmark* of "algorithmic qubits" (AQ = *Algorithmic Qubits*), using the log<sub>2</sub>of IBM's quantum volume and based on a set of intermediate-level application *benchmarks* from the QED-C consortium<sup>280</sup>. It has successively reached a level of 22 algorithmic qubits in 2022, then 29 and 35 in 2023<sup>281</sup>. It then plans to reach 64 AQ in 2025 by

<sup>274 &</sup>lt;u>Large Scale Modular Quantum Computer Architecture with Atomic Memory and Photonic</u> <u>Interconnects</u> by Christopher Monroe *et al*, 2014 (16 pages).

<sup>275</sup> Integrated photonic structures for photon-mediated entanglement of trapped ions by F. W. Knollmann *et al*, MIT and Sandia National Laboratories, *arXiv*, January 2024 (17 pages).

<sup>276 &</sup>lt;u>Entanglement of Trapped-Ion Qubits Separated by 230 Meters</u> by V. Krutyanskiy, Maria Galli, Nicolas Sangouard, Tracy Northup *et al*, PRL, February 2023 (22 pages).

<sup>277 &</sup>lt;u>Atom-photon coupling with trapped ions</u> by Tracy Northup, 2022 (40 slides).

High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network by L. J.

 Stephenson *et al*, University of Oxford, PRL, 2020 (6 pages).

Ion Trap with In-Vacuum High Numerical Aperture Imaging for a Dual-Species Modular Quantum

 Computer by Allison L. Carter, Christopher Monroe et al, UMD, October 2023 (8 pages).

<sup>280</sup> Quantum Economic Development Consortium - https://quantumconsortium.org/

<sup>281</sup> How We Achieved Our 2024 Performance Target of #AQ 35 by IonQ, January 2024.

exploiting logical qubits to 16 physical qubits, then 256 and 1024 between 2026 and 2028. This will only be possible by interconnecting 64-ion quantum processors with photons<sup>282,283,284,285</sup>. This led the company in 2023 to acquire the Canadian start-up Entangled Networks, which specializes in this field. It then acquired IDQ, Lightsynq and Capella Space in 2025 to pursue its development of a broad quantum telecommunications offering.

Quantinuum's *roadmap* began in 2020 with 1D circuits (H1), then moved on in 2023 to H2 '*racetrack*' circuits (similar to racecourses), supporting 56 qubits. Beyond that, the company plans to move on to a circuit with a crossover (Helios) with 96 qubits, then to 2D circuits in orthogonal rails and based on shuttle ions passing from block to block of qubits to link them by two-qubit gates<sup>286</sup>. However, it has not communicated precisely how many physical and logical qubits it plans to achieve with its H3, H4 and H5 QPUs<sup>287</sup>. The company updated its *roadmap* with a few more information in September 2024. It plans to assemble several thousand physical qubits and hundreds of

of logical qubits by 2030, with logical error rates of 10<sup>-16</sup>.

AQT (Austria) is also planning to implement an FTQC architecture. It produced its first CNOT logical gate in May 2022 using 16 physical qubits and a 7-qubit correction code.

<sup>282 &</sup>lt;u>Scaling the ion trap quantum processor</u> by Christopher Monroe and J. Kim, Science, 2013 (7 pages).

Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects by Christopher Monroe, Robert Raussendorf *et al*, PRA, 2013 (16 pages).

<sup>284</sup> IonQ Achieves Critical First Step Towards Developing Future Quantum Networks by IonQ, February 2024.

<sup>285</sup> Enabling Networked Quantum Computing with Ion-Photon Entanglement by IonQ, February 2024.

<sup>286</sup> Transport of multispecies ion crystals through a junction in an RF Paul trap by William Cody Burton et al, June 2022 (6 pages) where they describe how they can transport ytterbium and barium in 2D structures.

<sup>287 &</sup>lt;u>Scalable Multispecies Ion Transport in a Grid Based Surface-Electrode Trap</u> by Robert D. Delaney *et al*, Quantinuum, *arXiv*, March 2024 (11 pages).

It also produced a fault-tolerant T-gate using 16 qubits<sup>288</sup>. In 2023, it published a *blueprint* for implementing an FTQC<sup>(289</sup> architecture.

In the UK, Universal Quantum is working on interconnecting several QCCD- type trapped ion control modules with ytterbium ions operated at the hyperfine level of the 2S<sup>290</sup> orbital. Non-destructive measurement of the qubits is based on auxiliary ions of the barium type. In 2022, it has announced its intention to reach one million qubits thanks to its modular approach<sup>291</sup> (Figure 17). Its modules conduct heat well, enabling them to dissipate it efficiently. In July 2024, it announced the creation of an ASIC chip supporting its UQCOnnect interconnection platform <sup>292</sup>. It supports the transfer of ions between modules with a fidelity of 99.999,993%. It operates at 70 K.

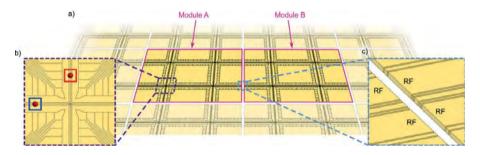


Figure 17: Universal Quantum's approach to interconnecting ion trapping modules, with ions able to move from one module to another. Each module has a size of 20 mm in their first generation. Source: <u>A high-fidelity quantum matter-link</u> <u>between ion-trap microchip modules</u> by M. Akhtar, W. K. Hensinger *et al*, Nature Communications, February 2023.

<sup>28</sup> Demonstration of fault-tolerant universal quantum gate operations by Lukas Postler, Rainer Blatt, Thomas Monz *et al*, *Nature*, November 2021 and May 2022 (14 pages).

Strategies for practical advantage of fault-tolerant circuit design in noisy trapped-ion quantum computers by Sascha Heußen *et al*, January 2023 (36 pages).

<sup>290</sup> Microfabricated Ion Traps by Winfried Hensinger et al, 2011 (28 pages).

<sup>291</sup> How Universal Quantum is rising to the million-qubit challenge, Universal Quantum, February 2022.

<sup>292 &</sup>lt;u>Universal Quantum develops key enabler of million-qubit quantum computer</u> by Universal Quantum, July 2024.

Oxford lonics is based on the WISE (*Wiring using Integrated Switching Electronics*) principle, a module supporting up to 1,000 ions in a 2D matrix arrangement (Figure 18). The chips are optimized for routing qubit control signals. In July 2024, the company published its first results exploiting qubit control with simple electrodes and voltages<sup>293</sup>. It obtained two-qubit gates with fidelities of 99.97%, but on a small scale, over 10 ions and 7 zones of 2 ions. Its chips are manufactured by Infineon in Germany. In a few years' time, it aims to achieve 256 qubits with this technology<sup>294</sup>.

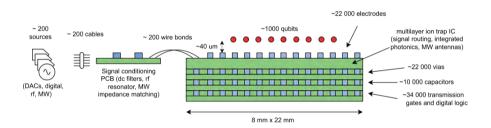


Figure 18: Multi-layer wiring diagram of Oxford Ionics' ion control chip capable of supporting up to 1,000 qubits. The chip comprises capacitors and routing circuits, fed by 200 wires, and covers an area of 8 mm× 22 mm. Source: <u>How to wire a 1000-qubit trapped</u> ion quantum computer by M. Malinowski *et al*, Oxford Ionics, *PRX Quantum*, October 2023.

In France, Crystal Quantum Computing's ambition is to create a quantum computer based on strontium ions with excited ions in Rydberg states. Their control requires UV lasers operating at 243 nm, using infrared lasers followed by frequency doubling and THz-directed microwave antennae<sup>295</sup>. It has not yet demonstrated these qubits or published an FTQC *roadmap*.

<sup>283 &</sup>lt;u>Scalable, high-fidelity all-electronic control of trapped-ion qubits</u> by C. M. Löschnauer, C. J. Ballance, C. Matthiesen, M. Malinowski, T. P. Harty *et al*, *arXiv*, July 2024 (12 pages).

<sup>294</sup> Oxford Ionics breaks global quantum performance records by Oxford Ionics, July 2024.

<sup>295 &</sup>lt;u>A microwave dressed Rydberg ion by Fabian Pokorny, Stokholm University, 2020 (146 pages).</u>

# 3.5.6. Conclusion

Trapped ions offer great potential with the best fidelities currently available. They make it possible to create fault-tolerant computers with reasonable ratios of physical qubits to logical qubits. Scaling them up will involve a number of major challenges: the ability to control the placement and movement of 2D ions on interconnected chips, then interconnecting these chips with resources of entangled photons or *shuttling ions*, and finally, somehow getting around the slowness of their quantum gates.

# Chapter 4

# SCALING UP QUANTUM COMPUTERS

Over the past three decades, experimental demonstrations of quantum computing have gone through three phases. Initially, research groups working on different technologies demonstrated that it was possible to create and control one, then a few qubits, with increasing fidelity. Then, when it was possible to control more than 50 qubits, though in an imperfect manner, demonstrations of quantum computational advantages were made: these are sampling tasks (boson sampling and circuit sampling) that are difficult to perform for a conventional supercomputer, but easy for a quantum processor, even a small and non- universal one<sup>296,297,298,299,300</sup>. Finally, from 2023 onwards, a few laboratories managed to control several hundred qubits well enough to produce a logical qubit (see Chapter 2). The next step will be to implement an architecture that scale above 1,000 qubits, and to demonstrate its necessary building blocks. Scaling up to control millions of qubits, will require combining several smaller modules interconnected by quantum links, as we explain in this chapter.

<sup>2%</sup> F.. Arute et al. "Quantum supremacy using a programmable superconducting processor", Nature 574, 505-510 (2019)

<sup>297</sup> H.-S. Zhong et al. "Quantum computational advantage using photons", Science 370,1460-1463 (2020)

<sup>298</sup> Y. Wu et al "Strong Quantum Computational Advantage Using a Superconducting Quantum Processor", Physical Review Letters 127, 180501 (2021)

<sup>299</sup> L.S. Madsen *et al.* "Quantum computational advantage with a programmable photonic processor", *Nature* 606, 75-81 (2022)

<sup>30</sup> The exact comparison of the performances for these sampling tasks, when noise is taken into account, between a classical computer and a quantum computer, is an on-going topic of discussion, as both classical and quantum computing are progressing, but it is outside the scope of this report.

### 4.1. SCALABILITY THROUGH MODULARITY

### 4.1.1. LIMIT TO MONOLITHIC INTEGRATION

A fault-tolerant guantum computer capable of implementing useful algorithms, for instance for chemical simulation, will require millions of physical qubits. However, current quantum processors cannot grow to scale up to his number as they stand. Because some of their physical properties depend on their size, and deteriorate when more gubits are added. For example, crosstalk is a phenomenon where a control signal applied to a gubit disrupts another gubit with a parasitic signal and can cause correlated errors on several gubits. When crosstalk increases with the size of the system, which can be the case for superconducting qubits or neutral atoms, the error rate in turn increases, and it becomes impossible to stay below the error correction threshold. For superconducting qubits, a second example of a physical effect becoming more important as the system size increases is of cosmic rays, which can cause a catastrophic error destroying coherence on the scale of the entire chip: the probability of which increases with the size of the chip. Photonic losses are a third example: photonic quantum processors suffer from photon losses, and when the manipulation of an increasing number of photonic gubits requires the use of an interferometer of increasing depth, the per-photon loss rate increases with the number of gubits. Moreover, there exist manufacturing constraints for circuits, particularly superconducting circuits, where gubits can take up a significant amount of space - around a quarter of a square millimetre per qubit, and a little more for cat gubits - particularly because of the size of the resonator used to read their state. The size of the chips is limited by the lithography systems and the size of the wafers typically 8 inches. These constraints could be partly overcome in the future by using manufacturing techniques on 12-inch (300 mm) wafers and by reducing the size of the resonators.

In addition, a number of components in current quantum processors is proportional to the qubit counts, and it would be physically impossible to increase this number to over a million in a unique integrated system. These include the cabling for the control electronics and the connection between the qubits and cryogenics (see Chapter 3 for details of the physical properties limiting the scaling up of qubit technology).

There are techniques to improve the hardware so that more qubits can be integrated despite these limiting factors, such as integrated electronics of the conventional cryo-CMOS or superconducting (SFQ or SQUID) type and multiplexing of control signals. Cryo-CMOS circuits appear to be viable for silicon qubits that operate at temperatures between 100 mK and 1 K, where the cryogenics budget is highest. It is more difficult to use them close to superconducting electronics, which operate with very low energy loss compatible with the load that the cryostat can support at this temperature. However, these circuits send different signals based on pulses, and generate parasitic effects on the qubits (*back action*). This technique is being promoted by the American start-up SeeQC.

Concerning the wiring problem, a superconducting qubit requires 2 to 5 control lines, depending on the case. That's just as many lines to be created between the processor chip and its control environment, including in the case of a neighboring superconducting chip. Unless we use multiplexing solutions integrated into these chips, this will be problematic – when to scaling up the number of qubits per chip. It gets even trickier when the control signals are created at room temperature. They require the use of flexible, coaxial cables that are connected to the chip. Today, we do not know how to produce flexible superconducting cables with low heat dissipation, so they are generally only installed in cryostats above 4 K, where the cooling budget is greater than at 15 mK. Below this level, coaxial cables should always be used.

In addition, to hardware improvements, there exist software methods known as circuit knitting<sup>301</sup> which make it possible to take advantage of several separate quantum circuits to simulate a single larger quantum circuit, and thus go beyond the maximum size of current processors at the cost of an overhead handled by classical computing. Although interesting to temporarily increase the capabilities of current processors and to run NISQ algorithms on larger instances, these techniques are not *scalable*, because the simulation overhead is, in the general case, exponential.

#### 4.1.2. LINK BETWEEN PHYSICAL QUBIT CONNECTIVITY AND ERROR CORRECTION

For a computing architecture to be fully scalable, and thus compatible with error correction and fault tolerance techniques, it must be composed of modules whose characteristics are constant, ie independent of the number of modules that are assembled. In this way, and only in this way, the basic principle of error correction, which consists in adding physical qubits to reduce the logical error as much as necessary when the physical error rates are below the threshold, applies.

These characteristics depend on both the qubit technology and on the system architecture which determines how qubits function collectively. A key element to implement logical qubits is the connectivity between the physical qubits. In this respect, the surface code (see Chapter 2) is a valuable quantum error correcting code, as it requires only 2D connectivity between nearest neighbors: each qubit is placed on a two-dimensional grid and must only interact with its 4 nearest neighbors, independently of the size of the code. However, this type of local code cannot match the performance of the best codes in terms of overhead cost in physical qubits per logical qubits and in terms of distance, for which non-local interactions are required. In this case, for the architecture to be scalable,

<sup>301</sup> C. Piveteau, D. Sutter. "Circuit Knitting With Classical Communication", IEEE Transactions on Information Theory, 70 (4) 2734-2745 (2024)],

the number of connections between qubits must be limited, to prevent connectivity requirements from exploding with the number of qubits. This corresponds to the family of quantum low-density parity check (LDPC) codes.

#### 4.1.3. COUPLING QPU WITH OPERATION ON REMOTE QUBITS

To overcome the limits of monolithic integration, one must not only assemble modules whose physical properties scale as required, but also ensure a quantum interconnection between these modules that is compatible with the quantum error correcting codes and fault-tolerant protocols one wishes to implement. Indeed, only exchanging classical information between quantum processors does not enable full-scale and universal fault-tolerant quantum computing on all the logical qubits in these processors. In other words, two quantum processors, each containing N qubits and connected by a channel of classical communications, do not constitute a 2N-qubit quantum processor.

One way to exchange quantum information between two processors is based via quantum teleportation, which use entanglement to transfer a state or to teleport a gate between distant qubits. In that case, the interconnection between the quantum processors is usually photonic.

To implement a CNOT gate between two remote qubits C and T of two distant processors, one can distribute qubits A and B of a Bell pair

$$|\psi_{AB}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$$

to each of the two processors, then implement local CNOT gates on the pairs (A,C) and (B,T), then measure the distributed qubits A and B, which implements, after correction by single-qubit gates, a CNOT gate on the remote qubits<sup>302</sup>. The corresponding circuit is shown in (Figure 19).

<sup>302</sup> We use the CNOT gate as an example, because it, together with a few single-qubit gates, forms a universal gate set.

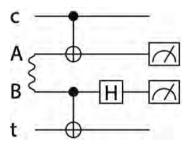


Figure 19: Teleportation of a CNOT gate. A CNOT gate is implemented between two distant qubits c and t using a Bell pair on qubits A and B. (represented by a wavy line), with qubit A (resp. B) distributed close to qubit c (resp. t). CNOT gates are then applied locally to the A and c qubits, on the one hand, and B and t, on the other. The A and B qubits are then measured. For simplicity, local corrections are omitted.

A variant consists in creating, for each of the two qubits on which the operation is to be performed, an additional qubit entangled with the initial qubit, then bringing these two additional qubits together in order to perform a Bell state measurement on them. Again, with single- qubit corrections, this is equivalent to applying a CNOT gate between the distant qubits. This is possible, for example, when the qubits are encoded in a degree of freedom of a single-photon emitter, which can emit a photon entangled with this degree of freedom.

A second possibility to implement two-qubit gates between two different modules, when the data qubits can be moved, is to move the qubits of each of the two modules, such as they are placed in a zone where they can interact directly. This approach is compatible with processors based on neutral atoms<sup>303</sup>, as well as with trapped ions, using the "ion shuttling" technique. It is promoted in particular by the British start-up Universal Quantum, which in 2024 demonstrated the high efficiency of this technique for moving an ion from one control chip to another.

<sup>303</sup> Q. Xu *et al.* "Constant-overhead fault-tolerant quantum computation with reconfigurable atom arrays", *Nature Physics* (2024)

A third approach, though more suited to shorter-range interactions, is to use a quantum bus based on a microwave resonator. The two distant qubits on which the operation is to be performed are coupled to the bus, and the quantum gate between the two qubits is performed globally thanks to this interaction with the bus. This method is particularly well suited to superconducting qubits.

#### 4.1.4. INTER-QPU LOGIC OPERATIONS

In order to implement the error correcting circuit and the fault-tolerant protocols, the proposed architecture must enable this type of remote qubit interactions between a certain number of physical qubits. The precise orchestration of logical operations between modules into a set of intra- and intermodule physical qubit interactions is determined by the architecture, the error-correcting code and the fault-tolerance protocol. For example, if one aims to encoding the logical qubits with a surface code and to perform the logical operations using lattice surgery techniques<sup>304</sup>, which make it possible to efficiently implement a universal set of gates on 2D topological codes while preserving local connectivity, one will have to be able to implement 2-qubit operations on the qubits located at the border of each of the lattice corresponding to a logical qubit, called a patch. This is illustrated by the example of a CNOT gate between two logical qubits in Figure 20<sup>305</sup>.

<sup>304</sup> D. Horsman *et al.* "Surface code quantum computing by lattice surgery", *New Journal of Physics* 14 123011 (2012)

<sup>305</sup> Other techniques, such as *braiding* or *twist deformation*, can be used to create a logical CNOT gate while preserving connectivity limited to the nearest neighbors.

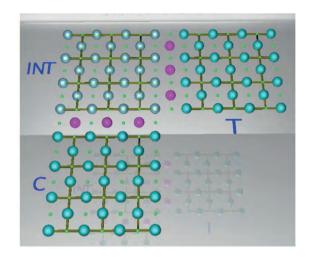


Figure 20: Creation of a logical CNOT gate using mesh surgery (in D. Horsman *et al.* 2012 New J. Phys. 14 123 011, CC BY-NC-SA 3.0). The physical qubits correspond to the larger circles, and each of the three grids ("patch") is a logical qubit encoded in surface code. To create a CNOT gate between qubits C and T using the INT intermediate qubit, one must perform two-qubit operations between the qubits located at the border of each patch and a line of intermediate physical qubits, in pink.

When qubits are not limited to connectivity between nearest neighbors, other architectures can enable a two-qubit logical gate to be implemented transversalely: the logical gate is performed by implementing the separate two-qubit gates between all the physical qubits of the two logical qubits, see Figure 21. To do this, one must be able to move the qubits so that they can interact two by two. This approach is compatible with mobile qubit platforms, such as neutral atoms or photons<sup>306</sup>.

<sup>30</sup> Not however that, it is impossible to implement all the gates of a universal set gate set in a transversal manner (Eastin-Knill theorem). The most common way of achieving universality is through distillation and magic state injection.

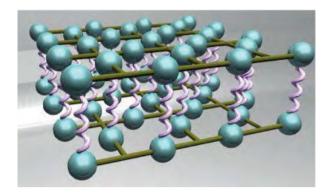


Figure 21: Implementation of a transversal logical CNOT gate (in D. Horsman *et al.* 2012 New J. Phys. 14 123 011, CC BY-NC-SA 3.0). The two patches correspond to the two logical qubits encoded using the surface code. The logical gate is performed by implementing a physical CNOT gate between each physical qubit (pink lines).

#### 4.2. PHOTONIC INTERCONNECTION

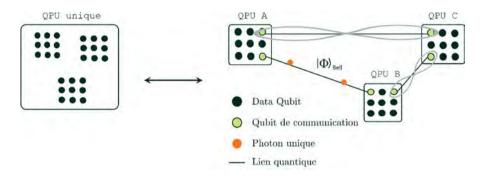
Efforts in the field of quantum computing in recent years have naturally focused on the development of increasingly large quantum processors (QPU, for *Quantum Processing Unit*) to enable the execution of algorithms demonstrating an exponential advantage in computing time compared with a conventional processor. Such an advantage, however, requires tens of thousands or hundreds of thousands of qubits, depending on the algorithm<sup>307,308</sup>. The technological challenges associated with developing a machine of this size are numerous: technical limitations, for example in terms of cryogenics requirements or circuit density depending on the qubit modality, but also the ability to reduce noise, which is itself linked to the error-correcting code used and the additional resources required.

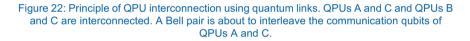
In this context, the solution of interconnecting QPUs in a modular computing architecture is currently attracting a great deal of interest, as it opens up a realistic route to scaling up quantum computing.

<sup>307</sup> E. Gouzien *et al.* "Factoring 2048-bit RSA Integers in 177 Days with 13 436 Qubits and a Multimode Memory", *Phys. Rev. Lett.* 127 140503 (2021).

<sup>308</sup> L. Clinton *et al*, "Towards near-term quantum simulation of materials", *Nature Communications*, 15: 211 (2023)

within a few years. This interconnection consists of interleaving qubits belonging to separate QPUs so that a much larger number of qubits than those in a single QPU is available for calculation. A distributed algorithm is then run on all these interconnected QPUs, which is equivalent to a quantum computer of larger size (Figure 22).





We are interested here in the model of computation based on quantum circuits. In this model, an algorithm is given in the form of an abstract quantum circuit, involving a certain number of qubits acted upon by unitary operators with one or more qubits, represented by quantum logic gates.

In order to create such gates between qubits belonging to distinct QPUs, two main protocols are used, called TeleData and TeleGate (see paragraphs below)<sup>309</sup>. These protocols are both based on quantum entanglement. Their use leads to the definition of two categories of qubits contained in each QPU to be interconnected: (1) the data qubits, which are used to carry out operations within the QPU itself, and (2) the communication qubits, which are used to establish the

<sup>30</sup> Van Meter, "Arithmetic on a Distributed-Memory Quantum Multicomputer", ACM Journal on Emerging Technologies in Computing Systems, 3 4 1-23 (2008)

entanglement between QPUs. Establishing entanglement between two communication qubits requires the intermediary of a Bell pair, each member of which will be entangled with one of the communication qubits (Figure 22). The hardware system that makes the interconnection is called a quantum link. Establishing an interconnection between two QPUs with a quantum link therefore necessitates, in an ideal case, a local qubit in each QPU and consumes a Bell pair.

#### 4.2.1. PARTITIONING QUANTUM ALGORITHMS

Partitioning a quantum algorithm is a problem inherent in the distributed quantum computing model. In this sense, the ability to partition a circuit is a major challenge for scaling up quantum computing, whatever the size of a QPU (*Quantum Processing Unit*) that it will be possible to achieve in the future, because the continuous increase in computing resource requirements will lead to the need for interconnection between QPUs.

#### 4.2.1.1. CLASSICAL SCORE

The first approaches were to classically split the initial quantum circuit into a set of mutually independent sub-circuits. The general idea, given the availability of a QPU of limited size and a quantum circuit requiring more qubits than those available on the QPU, is to divide the circuit into sub-circuits whose size corresponds to the number of qubits available on the QPU, replacing the operations rendered non-local during slicing with local operations. The subcircuits are activated one by one on the QPU, and the results are recombined using conventional post-processing techniques to reproduce the result of the initial circuit<sup>310</sup>.

This partitioning technique, and in particular the optimization of cuts, is currently the subject of active research<sup>311</sup>. However, attempting to reconstruct the quantum result from the classical intermediate results leads to a significant

<sup>310</sup> Piveteau *et al*, "Circuit knitting with classical communication", *IEEE Transactions on Information Theory*, 70 4 (2023)

<sup>311</sup> Ibid.

additional cost in terms of traditional computing resources. This extra cost increases exponentially with the number of cuts<sup>312</sup>, preventing realistic implementation beyond a few dozen cuts. Preserving entanglement when cutting a circuit and allowing non-local operations promises a more usable and optimal partition. This is what quantum interconnection of QPUs is all about.

#### 4.2.1.2. QUANTUM PARTITIONING PRESERVING ENTANGLEMENT

Quantum interconnection of QPUs provides the possibility of establishing entanglement between different QPUs and thus enables logical operations to be carried out between separate QPUs, exploiting the two protocols mentioned earlier that are the building blocks of distributed quantum computing: TeleData and TeleGate.

#### a) Teledata

We want to transfer the quantum state of a qubit in one QPU to a qubit in another QPU (Figure 31). This protocol consists of a sequence of four phases:

- generation of the entanglement of two communication qubits belonging to two distinct QPUs;
- 2. logical operations on the qubits of the initial QPU (CNOT gate and one-qubit control-rotation and Hadamard gates);
- intermediate measurements of the state of the qubits involved in the initial QPU and conventional communication of the result of these measurements to the second QPU;
- 4. local conditional operations in the second QPU (one-qubit control-rotation gate).

Harada *et al.* "Doubly optimal parallel wire cutting without ancilla qubits" *arXiv*: 2303.07340 (2023).

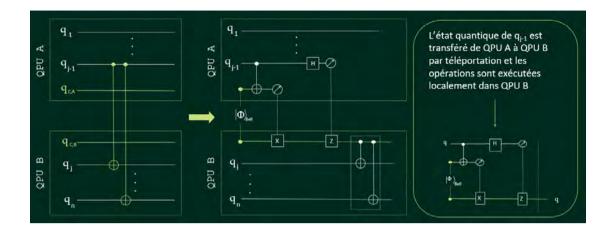


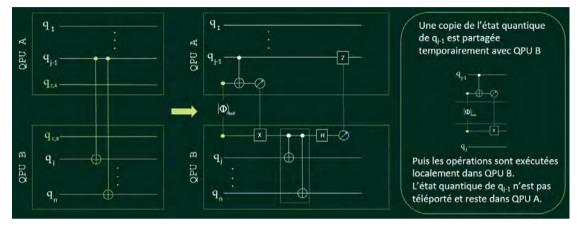
Figure 23: Schematic diagram of the TeleData protocol used to implement a non-local twoqubit gate of the control-rotation type between QPU A and QPU B. A Bell pair  $\Phi$  is consumed by communication qubits to establish entanglement between QPU A and QPU B. The quantum state of the qubit  $q_{(i)(-1)}$  is then teleported to QPU B. The logic gates are then executed locally in QPU B between the qubit and  $q_i$ . At the end, the information is no longer contained in QPU A.

With this protocol, like a quantum teleportation protocol, the quantum state of a qubit belonging to QPU A is transferred to a qubit belonging to QPU B, and further operations are carried out locally in the latter. The first QPU no longer possesses information about the initial qubit.

### b) TeleGate

We want to create a two-qubit gate between QPUs A and B (Figure 24). The TeleGate protocol can be broken down into three main phases and has the particularity of exploiting an intermediate quantum state which shares a copy of the initial state with the target QPU, without teleporting the quantum state. Taking the example of a two-qubit control-rotation gate (CX for example), the three stages are:

 entanglement of the quantum state of the control qubit with a communication qubit of the second QPU. At this stage, a copy of the control qubit is shared locally on the second QPU, transforming the communication qubit of the QPU-B into an intermediate control qubit; 2. local operations between the intermediate control qubit and the target qubits;



3. restoration of the initial quantum state which can then be used.

Figure 24: Schematic diagram of the TeleGate protocol used to implement a non-local twoqubit gate of the control-rotation type between QPU A and QPU B. A Bell pair Φ is consumed by communication qubits to establish entanglement between QPU A and QPU B. The quantum state of the qubit q<sub>i-1</sub>is then momentarily shared with QPU B and the logic gates are executed locally in QPU B between the communication qubit and q<sub>i</sub>. At the end, the intermediate quantum state is reset and the information is no longer contained in QPU B.

With this protocol, the quantum state of the control qubit of QPU A has only been temporarily shared with QPU B, during the time it takes to create the distributed gates concerned, without being teleported. In the end, the quantum information remains in QPU A and is not present in QPU B.

We also note that in the context of interconnecting a large number of QPUs, an alternative route to using the TeleData and TeleGate protocols is to develop protocols involving multi-partite entangled quantum states (for example, maximally entangled GHZ states, for *Greenberger-Horne-Zeilinger*). This could optimize the cost of the entanglement distribution as the number of

interconnection links and the number of QPUs to be interconnected<sup>313,314</sup> increases. However, this case requires in-depth investigation, taking into account all the necessary resources.

#### 42.13. Optimal partitioning of a quantum circuit

The partition optimization criterion we use is to minimize the number of interconnection quantum links so as not to introduce too great a hardware overhead.

For this optimization, it is necessary to have a distribution compiler adapted to the interconnection of QPUs<sup>315</sup>. This software brick takes as input the quantum circuit to be distributed as well as the characteristics of the network of QPUs (number of QPUs, number of qubits per QPU, number of quantum links, number of communication qubits per QPU, connectivity, etc.) within the QPU, network topology, etc.). It outputs the partition of the initial quantum circuit into a set of sub-circuits linked together by a number of TeleData and TeleGate operations, minimizing the number of Bell pairs required.

As shown in Figure 25, the first step in the distribution compiler is to map the initial quantum circuit data to a graph structure in which the qubits become the vertices and the two-qubit gates become the edges.

In a second step, the graph representing the quantum circuit is partitioned using graph partitioning techniques. Each partition corresponds to a sub- circuit, a fragment of the initial algorithm, to be executed on one of the interconnected QPUs.

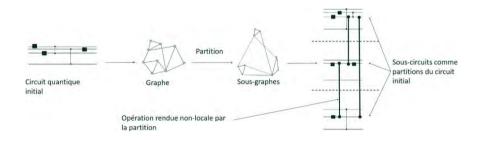
<sup>313</sup> Meignant *et al.* "Distributing graph states over arbitrary quantum networks", *Phys. Rev.* À 100 052333 (2018).

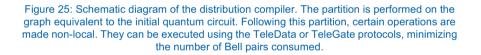
<sup>314</sup> Barral *et al.* "Review of Distributed Quantum Computing. From single QPU to High Performance Quantum Computing", *arXiv*: 2404.01265 (2024)

<sup>315</sup> Tomesh *et al.* "Divide and Conquer for Combinatorial Optimization and Distributed Quantum Computation", *IEEE International Conference on Quantum Computing and Engineering (QCE)* 1 1-12 (2023).

The third step is to identify the multi-qubit operations which, after partitioning, involve qubits belonging to different sub-circuits. These are distributed operations that need to be carried out using the protocols described above.

The distribution compiler iterates these steps to minimize the number of Bell pairs used. It then provides the optimal graph partition. In a final step, the sub-graphs are then transformed back into sub-circuits involving distributed operations. Each of these sub-circuits will be sent to a QPU and the interconnection quantum links will perform the necessary operations according to the optimal partition.





#### 4.2.1.4. LINK TO ERROR CORRECTION

An important aspect of distributed quantum computing is error correction, which is an essential step if quantum computing is to be fully adopted. It is in fact necessary to anticipate the fault-tolerant version of the algorithms as input to the distribution compiler. This will lead to an additional optimization step compared with the one presented previously, taking into account the additional cost of intermediate measurements that must be carried out during the circuit, depending on the error-correcting codes chosen. The possibility of interleaving qubits at an arbitrary distance, even between distinct QPUs, also opens the way to new error-correcting code architectures that promise to drastically reduce the ratio between the number of physical qubits and the number of logical qubits<sup>316,317</sup>. This field of exploring error correction facilitated by quantum interconnection is still in a preliminary phase, but has interesting potential.

#### 4.2.2. HARDWARE IMPLEMENTATION

From a physical point of view, the most suitable way of creating quantum interconnection links is to use pairs of entangled photons. This takes advantage of the great maturity of photonic technologies.

Distributed quantum computing on a network of interconnected QPUs will require quantum links that allow for varied and flexible architectures. These links will have to create entanglement on demand, quickly enough not to limit the execution of quantum algorithms, and preserving the fidelity of the computation's logical operations.

### 4.2.2.1. QUANTUM INTERCONNECTION LINKS

A quantum link is a system that creates photonic entanglement to establish the interconnection between the communication qubits of the QPUs. Depending on the technology of the QPUs, it is interesting to have a quantum link that absorbs a Bell pair emitted by the quantum registers to be interconnected or a quantum link that produces a Bell pair on demand and makes it available. In both cases, the speed at which entanglement is established and the resulting fidelity of the TeleData and TeleGate protocols are the two main parameters of the link.

<sup>316</sup> Cohen *et al.* "Low-overhead fault-tolerant quantum computing using long-range connectivity", *Science Advances* 8 20 (2022)

<sup>317</sup> Pecorari *et al.* "High-rate quantum LDPC codes for long-range-connected neutral atom registers", *arXiv*: 2404.13010 (2024)

#### a) Absorptive quantum bond

For certain types of QPU, such as those with superconducting qubits, the emission of a single photon by the communication qubit is the simplest approach. In this case, when the photon is emitted, it is entangled with the communication qubit that emitted it. As shown in Figure 26, the quantum link collects the photons emitted by each of the QPUs to be interconnected, and causes them to interfere by performing a Bell measurement. The fact that these photons interfere at the same time on the Bell module creates entanglement between these photons, which results by transitivity in the entanglement of the communication qubits belonging to the QPUs.

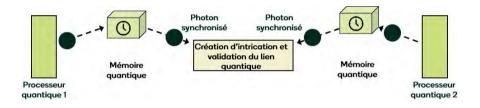


Figure 26: Schematic diagram of an absorptive quantum link. The quantum link collects the single photons emitted by the communication qubits of the QPUs to be interconnected and causes them to interfere to establish entanglement between the communication qubits. At the end of the process, the Bell pair is consumed.

The emission of the photon entangled by the QPUs is probabilistic, and it is necessary to use a quantum memory to synchronize the arrival of the photons collected on the measurement module.

#### b) Emissive quantum bond

In certain situations, it is more advantageous to have a quantum link that produces Bell pairs on demand. In this case, the quantum link produces a photonic entanglement resource that is made available to the QPUs at an arbitrary time that can be adjusted by the computation scheduler. A pair of single entangled photons is created within the link. Each of these photons is then stored in a quantum memory, which synchronously re-emits a photon to the QPUs to be interconnected. This pair of entangled photons forms the Bell pair that will be absorbed by the communication qubits of the QPUs, which will establish the interconnection between QPUs (Figure 27).



Figure 27: Schematic diagram of a quantum emissive link. The quantum link prepares a pair of entangled photons (Bell pair) which is made available to the QPUs. to be interconnected. Each of these photons is then absorbed by a communication qubit. Thanks to quantum memories, this type of interconnection link can provide a photonic entanglement resource on demand.

In this situation, the role of on-demand quantum memories is essential. They store the entanglement resource that has been created in advance, until the moment when the Bell pair is ready to be consumed by the QPUs to implement a TeleData or TeleGate protocol. It is thus possible to adjust the time delays between the arrival of the Bell pairs and the execution of the logic gates in order to minimize the duration of the protocols and thus maximize the effective computing time.

For example, in the case of the interconnection of neutral atom QPUs, we want to be able to provide a photonic entanglement resource at least at each QPU cycle just before the start of the calculation phase, i.e. at a cycle time of the order of 100 ms<sup>318</sup>.

#### 4.222. EFFICIENT ON-DEMAND QUANTUM MEMORY

An optical quantum memory on demand is a system that accepts a photonic quantum state as input, stores it for a period of time and later re-issues it on command from an external trigger signal.

<sup>318</sup> L. Henriet et al, "Quantum computing with neutral atoms", Quantum 4 327 (2020)

In a quantum link, the memory acts as a buffer and allows the emission of single photons to be adjusted over time, which is an essential feature for optimal interconnection.

A quantum memory is characterized by four main physical parameters:

- memory efficiency, which corresponds to the probability of recovering an output photon. This parameter is essential for efficient interconnection. The state of the art in efficiency is currently around 90%<sup>319</sup>. This level of performance enables quantum links to establish entanglement with fidelities greater than 99.5% and cycle times of the order of 50 ms;
- the storage time, which corresponds to the length of time the photon can be stored before the efficiency of the memory is reduced by a factor of 1/e. For local interconnection, a storage time of a few tens of microseconds is sufficient;
- memory fidelity, which compares the quantum state of the output photon with the quantum state of the input photon. Today, quantum memories have fidelities of over 99.5%;
- multiplexing, which indicates how many photons a quantum memory can store at a time.

There are several physical platforms for producing a quantum memory (cold atoms, vapour atoms, doped crystals, etc.). A cold atom quantum memory has favourable characteristics for interconnection

of QPUs and benefits from the maturity of cold atom technology to turn it into an industrial-standard product that can be used outside laboratories.

<sup>319</sup> Cao *et al.* "Efficient reversible entanglement transfer between light and quantum memories", *Optica* 7 10 (2020)

At present, cold atom quantum memories hold the world record for efficiency at 90%, making it possible to interconnect QPUs using quantum links.

#### 4223. INTERFACE BETWEEN THE QUANTUM LINK AND THE QPUS

The qubit technology of the QPUs to be interconnected (neutral atoms, trapped ions, photons, etc.) will determine the physical characteristics of the Bell pair photons used to create entanglement. In particular, their wavelength and their temporal spread are decisive. In order to obtain high- performance TeleData and TeleGate protocols, it is necessary to adjust the interface between the quantum memory and these photons.

It is possible to adjust the optical wavelength using a frequency conversion stage. For this mature technology based on the use of non-linear crystals, it will be necessary to maximize conversion efficiencies. In the case of superconducting qubits, it is necessary to add a transduction stage, which converts a qubit in the microwave domain into a qubit in the optical domain. Several players are currently working on developing this type of transducer, which will be available as a subsystem<sup>320</sup>.

As far as the time spread of photons is concerned, it is possible to adjust the physical parameters of the quantum memory so that it accepts shorter or shorter photons. In particular, cold atom memories offer the versatility needed to adapt to different qubit modalities.

#### 4.2.2.4. Non-ideal case

In practice, the Bell pairs produced to establish entanglement are not perfect. In addition, the losses introduced during coupling between the single photons and the communication qubits must be taken into account.

Weaver *et al.* "An integrated microwave-to-optics interface for scalable quantum computing", Nat. Nanotechnol. 19 166-172 (2024).

It will therefore be necessary to add an entanglement<sup>321</sup>and redundancy purification step at the quantum link level. This indicates that it will be useful to allocate several communication qubits per QPU for a TeleData or TeleGate protocol and to have multiplexed quantum links.

# 4.2.3. CONCLUSION

Interconnecting QPUs using photonic quantum links based on efficient quantum memories is a promising way of scaling up quantum computing. By networking QPUs of intermediate size, it will make it possible to envisage an effective quantum computer with a sufficient number of qubits, quality of qubits and coherence time to execute non-trivial quantum algorithms.

The current development of multiplexed quantum memories will enable more complex quantum network topologies, the implementation of quantum link redundancy schemes to improve the fidelity of distributed quantum gates, and pave the way for the efficient distribution of multipartite entanglement.

To ensure that quantum computing can be scaled up through the interconnection of QPUs and that fault-tolerant distributed quantum computing can be fully adopted, it is necessary to anticipate the integration of quantum link technologies into intensive computing environments, both in terms of software and hardware. These interconnection technologies constitute a quantum resource in their own right, in the same way as QPUs, and will enable clusters of interconnected quantum processors to be put into service.

### 4.3. SCALABILITY AND CONNECTIVITY OF SUPRACONDUCTOR QUBITS

The fidelity of one- and two-qubit gates is steadily increasing in superconducting qubit platforms. In the architecture most widely used worldwide (transmon-type qubits), the fidelity of the CNOT gate is is now between 99.3% and 99.7% (IBM's Heron r1 and r2 processors will be launched in 2023 and 2024).

<sup>321</sup> Bennett *et al.* "Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels", *Phys. Rev. Lett.* 76 722 (1996)

In addition, the factors limiting measurement fidelity are in the process of being understood. We are therefore now below the theoretical thresholds for error correction codes (*surface code*, codes, etc.).

LDPC...).

In this section, we give an overview of the potential problems when scaling up supra-architectures (for processors with 1,000 or even 1 million qubits) and the efforts being made to solve these problems. We also look at ways of increasing the connectivity (number of qubits with which each qubit can perform a logic gate) of these architectures.

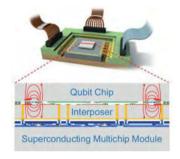
# 4.3.1. 3D INTEGRATION AND ELIMINATION OF PARASITIC COUPLING DURING SCALING

Most processors with a small number of qubits are built using a planar architecture. When the number of qubits is increased, two problems arise:

- if the qubits are arranged in a grid, it is more difficult, but not impossible, to control all the qubits independently without the control lines crossing each other or the interconnections between qubits (when these lines cross, a signal tends to leak from one line to the other). By increasing the size of the chip to incorporate more qubits, we lower the frequency of electromagnetic modes defined by the metal sample holder containing the chip (whose field is mainly contained in the substrate). These modes are responsible for *cross talks* between qubits that are not nominally connected. One solution to these two problems is to use vias. A via is a tunnel piercing the chip, the edges of which are metallised (TSV: *Through Silicon Vias*). They are used to pass control lines orthogonal to the chip;
- define conductive boundaries to avoid the appearance of resonant modes at low frequencies.

The TSVs must be made of superconducting materials, which we do not yet know how to do. An example is given in Figure 28. A similar technology is developed in<sup>322, 323</sup>. This technology reduces *cross talks* between qubits that are not nominally connected.

of~ -40 dB (for the qubits that are physically closest, this isolation is supposed to increase exponentially with distance). This makes it possible to The principle also allows distant qubits to be connected to the chip, and not just their nearest neighbors.



#### Figure 28 Source : Yost, Donna-Ruth W., *et al.* "Solid-state qubits integrated with superconducting through-silicon vias." npj Quantum Information 6.1 (2020): 59.

It should be noted that:

- with a large number of qubits, we are exposed to the phenomenon of *spectral crowding*: it becomes increasingly likely that two qubits will have a very close frequency. If the residual coupling is of the order of frequency mismatch, the two qubits hybridize and we obtain an unwanted interaction (gate) between qubits that are not nominally connected. To avoid these "frequency collisions" on a chip with a large number of qubits, cross talks will have to be strongly suppressed at long distance, as the frequency of the qubits is difficult to predict

323 Spring, Peter A., *et al.* "High coherence and low cross-talk in a tileable 3D integrated superconducting circuit architecture." *Science Advances* 8.16 (2022): eabl6698.

<sup>322</sup> Acharya, Narendra, *et al.* "Integration of through-sapphire substrate machining with superconducting quantum processors." *arXiv* preprint *arXiv*: 2406.09930 (2024).

at better than~ 1% at the time of manufacture, due to the lack of reproducibility of the critical current of Josephson junctions (despite recent advances in adjusting this critical current using local annealing);

- To switch interactions between nominally connected qubits on or off, the main players in the field use adjustable couplers: a flux-adjustable circuit is placed between the two qubits to be connected. This circuit meditates an interaction that is added to the natively present interaction. When we want to switch off the interaction, we choose a flux for which these two contributions cancel each other out exactly<sup>324,325</sup>. This circuit takes up a space equivalent to that of a qubit and about two are needed per qubit (121 couplers in Google's 72-qubit Sycamore chip), which in fact multiplies the chip area by almost three for an equal number of useful qubits.

#### 4.3.2. FOOTPRINT AND ENVIRONMENTAL RESOURCES

#### 4.3.2.1. CIRCUIT FOOTPRINT

If the problems of *cross-talking* and 3D addressing are resolved, there is, *a priori*, no obstacle to increasing the size of chips and the number of qubits. In terms of order of magnitude, a transmon has a typical footprint of less than 0.5 mm× 0.5 mm, which enables 300,000 qubits to be assembled on a *wafer* 12 inches in diameter. The footprint per transmon could be drastically reduced by using *parallel plate capacitor* technologies.

<sup>324</sup> Arute, Frank, *et al.* "Quantum supremacy using a programmable superconducting processor." *Nature* 574.7779 (2019): 505-510.

<sup>&</sup>lt;sup>325</sup> Zhang, Helin, et al. "Tunable Inductive Coupler for High-Fidelity Gates Between Fluxonium Qubits." PRX Quantum 5.2 (2024): 020326.

### 4.3.2.2. MICROWAVE COMPONENT FOOTPRINT

In a *full-stack* quantum system, a major contribution comes from the "environment" components used to generate and detect the control signals and route them to each qubit. In particular, the (attenuated) input signals must be physically separated from the (amplified) output signals by non-reciprocal components (circulators). Most commercially available components (circulators, amplifiers,

attenuators, etc.) are bulky (~ cm<sup>3</sup>) and expensive. Major efforts are being developed by the community to miniaturize and integrate these components on chip<sup>326,327</sup>.

### 4.3.2.3. CLASSIC ELECTRONICS RESOURCES

Another *hardware* component whose footprint may pose a problem is the electronics used to generate the control pulses, detect the measurement pulses and perform more or less complex mathematical operations for error correction. The cost of these systems is high at the moment (around 10,000 euros to control a qubit), but is set to fall as a result of competition from

programmable open source software (Qubic, Qick, etc.). The recent ability to generate arbitrary microwave pulses numerically (as opposed to a

analogue processing of a monochromatic signal mixed with low-frequency pulses) also opens up the possibility of extensive multiplexing to share a single control line between several qubits. On-chip reconfigurable microwave routing systems could further increase multiplexing capabilities. Finally, it may eventually be possible to carry out conventional information processing operations at low temperature in a system integrated with the quantum architecture (low-energy consumption RSFQ technology)<sup>328</sup>.

<sup>326</sup> Macklin, Chris, *et al.* "A near - quantum-limited Josephson traveling-wave parametric amplifier." *Science* 350.6258 (2015): 307-310.

<sup>327</sup> Chapman, Benjamin J., et al. "Widely tunable on-chip microwave circulator for superconducting quantum circuits." *Physical Review* X 7.4 (2017): 041043.

Liu, Chuan-Hong, *et al.* "Single flux quantum-based digital control of superconducting qubits in a multichip module." *PRX Quantum* 4.3 (2023): 030310.

#### 4.3.2.4. CRYOGENICS RESOURCES

Recent announcements have shown that it is possible to significantly increase the size and power of helium dilution cryostats used to cool superconducting qubits to 10 mK. In particular, two cryostats have been connected by a superconducting link<sup>329</sup> and a cryostat capable of accommodating several thousand qubits using current microwave components (Bluefors KIDE). If the on-chip integration of components continues, cryogenics should not be a limiting factor.

At the opposite end of the spectrum, we are seeing the emergence of smaller cryostats that are less expensive, require less helium-3 and consume less energy, driven by new players (Maybell, Qinu) as well as market leaders (Bluefors). Other industrial approaches, such as those proposed by Air Liquide, offer the possibility of creating very high-power cryostats to support a large number of physical qubits.

#### 4.3.3. CONNECTIVITY

The connectivity of an architecture denotes the number of qubits with which each qubit can implement a logic gate. It can be defined more broadly as the number of elementary logic gates for interleaving two arbitrary qubits.

The *surface code*, which is the most widely studied error-correcting code, has the advantage of having a low connectivity structure (4) with interactions between nearest neighbors only according to a planar geometry. Other codes offer better performance (in terms of correction threshold or number of encoded logic qubits per physical qubit) with greater connectivity and/or a more elaborate topology (structure with more than two dimensions). LDPC codes are a promising example of high-density logic qubit codes.

<sup>329</sup> https://bluefors.com/products/kide-cryogenic-platform/

Another element to take into account is the read chain, which is currently based on the use of circulators, low-noise amplifiers operating at 15 mK (notably the TWPAs from *Silent Waves*), HEMT amplifiers operating at 4 K and amplifiers operating at ambient temperature. Together, these devices can frequency multiplex read signals of around 8 to 10 qubits when low-noise amplification is based on a TWPA. But each of these elements takes up a significant amount of space, with parallelepipeds measuring around 2 cm on each side and less than 1 cm thick. Techniques aimed at miniaturizing these components are being developed. In particular, circulators can be replaced by superconducting circuits that take up much less space.

#### 4.3.3.1. SINGLE-CHIP CONNECTIVITY

Compared with other platforms where the physical systems encoding quantum information can be moved around, *all-to-all* connectivity seems unrealistic for superconducting circuits at the current stage. However, it seems possible to establish a few "long-distance" connections on the chip (for example according to the diagram described in figure 29 to enrich an architecture with connections between close neighbors. This makes it possible to drastically reduce the number of elementary gates needed to perform an operation on 2 arbitrary qubits, or to efficiently implement a precise code<sup>330</sup>.

Ultimately, an *in situ* reconfigurable on-chip microwave routing architecture would make it possible to increase connectivity in an arbitrary way.

#### 4.3.3.2. INTER-CHIP OR LONG-DISTANCE CONNECTIVITY

The connection of several chips enabling logic gates between qubits a metre apart has been demonstrated, with gate fidelities approaching those on a single chip<sup>331</sup>. Although these

<sup>330</sup> Bravyi, Sergey, *et al.* "High-threshold and low-overhead fault-tolerant quantum memory." *Nature* 627.8005 (2024): 778-782.

<sup>331</sup> Niu, Jingjing, et al. "Low-loss interconnects for modular superconducting quantum processors." Nature Electronics 6.3 (2023): 235-241.

Although the demonstrations are small-scale for the time being, and have a large hardware footprint, this type of *modular* architecture could ultimately offer decisive advantages: replacement of defective or out-of-specification modules, correction of *chip-scale failures*, elimination of the need to install new modules, etc.

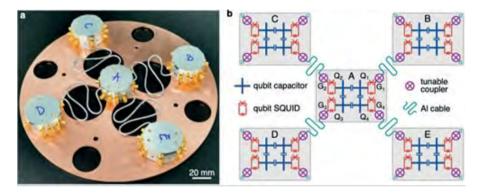


Figure 29 Source: Niu, Jingjing, *et al.* "Low-loss interconnects for modular superconducting quantum processors." Nature Electronics 6.3 (2023): 235-241.

Given the high microwave losses in microwave transmission lines (attenuation per metre much higher than in the optical domain, even for superconducting lines), long-distance entanglement will certainly require transduction to the optical domain. The quantum efficiency and speed of transducer systems are increasing steadily, but these systems are not yet mature.

# 4.3.4. OTHER CHALLENGES FOR OPERATING ARCHITECTURES WITH A LARGE NUMBER OF QUBITS

#### 4.3.4.1. QUASI-PARTICLE POISONING AND CHIP-SCALE FAILURE

The widespread failure of all the qubits on a chip (or at least a significant degradation in the coherence time of the majority of qubits) is a phenomenon observed on all current large processors on a typical timescale of 1 s to 1 min. These errors cannot be corrected. They are attributed to high-energy impacts from cosmic rays or from the natural radioactivity of materials. These impacts generate cascades of phonons, which in turn generate a large number of quasiparticles that can cause qubits to relax when they tunnel through Josephson junctions. Shielding against high-energy impacts seems useful, but insufficient (even in laboratories buried under a mountain). The trapping of phonons<sup>332</sup> or quasiparticles themselves<sup>333</sup> seems to be a promising approach, but for the moment it is very inadequate given the execution time of quantum algorithms. A modular architecture could solve this problem (possibly with intermediate-sized modules [100-1,000 qubits] to limit the total footprint).

#### 4.3.4.2. Two-level systems

The term *Two-Level System* (TLS) refers to a group of parasitic systems coupling to superconducting qubits. Their origins are not well known and are probably multiple (electrical fluctuators trapped at the interfaces, etc.). substrates/air, substrate/metal or metal/air, low-frequency magnetic fluctuators, etc.). When a TLS is resonant or quasi-resonant with a qubit, it can degrade its T1 relaxation time and modify its

resonance. It has also been observed that this "TLS bath" reconfigures itself in an unpredictable way on time scales ranging from one second to several hours.

<sup>332</sup> Henriques, Fabio, *et al.* "Phonon traps reduce the quasiparticle density in superconducting circuits." *Applied physics letters* 115.21 (2019).

<sup>333</sup> McEwen, Matt, *et al.* "Resisting high-energy impact events through gap engineering in superconducting qubit arrays." *arXiv* preprint *arXiv*: 2402.15644 (2024).

This is probably due to the interaction of TLS resonating in the GHz range with low-frequency TLS<sup>334</sup>.

TLSs are a major challenge for architectures with a large number of qubits, because their reconfiguration requires frequent recalibration of the chip to compile an algorithm avoiding degraded qubits, readjusting couplers to turn off unwanted interactions (see 1 b) and generally limiting the fidelity of logic gates.

The community is focusing its efforts on :

- materials engineering to limit the number of TLS ;
- rapid recalibration of system parameters, possibly using machine learning;
- the design of new qubits that are less sensitive to electrical fluctuations.

# 4.3.5. CONCLUSION

There are no major obstacles identified to the advent of large-scale superconducting circuits if the technologies continue to progress rapidly. The most serious current challenges are probably those of quasiparticle poisoning and TLS baths. It should be noted that these are also the challenges that the community has become aware of most recently, and the architectures for dealing with them have not yet been optimized to the same extent as for other sources of error.

We have focused on transmon qubits, which a r e the most widely used at the moment, but the other types of qubit that are maturing (fluxoniums, bosonic qubits, etc.) are facing the same challenges.

<sup>334</sup> Faoro, Lara, and Lev B. loffe. "Interacting tunneling model for two-level systems in amorphous materials and its predictions for their dephasing and noise in superconducting microresonators." *Physical Review* B 91.1 (2015): 014201.

We'll finish with an overview of the players. At global level, the and (to a lesser extent) Europe (including Israel) have developed considerably over the last few years, with major industrial players (IBM, Google, Amazon, Northrop Grumman), major mid-sized start-ups (Rigetti computing, IQM, Alice&Bob, etc.), companies offering foundry services (Quantware, Imec), dedicated electronics (Keysight, Zurich Instruments, QBlox, Quantum Machines), microwave components (LNF, Silent Waves for low-noise amplifiers)... The leading academic groups in the field are located in the United States (MIT, Yale, Stanford, Berkeley, etc.) as well as in Europe. public institutes (National Labs) providing academic groups with Lincoln Labs' TWPA amplifiers). Japan has a leading academic group (Riken) as well as high-level electronic component industries. The Chinese ecosystem, which is partly uncoupled, is not well known. The Chinese Academy of Science and Tsinghua University are among the leading academic players, and at least one company, Origin Quantum, seems to be positioned in the same niche as the major American groups. In France, the main academic groups are located in Saclay, Paris and Grenoble, but have not invested heavily in scaling up superconducting qubits. Alice&Bob is the leading start-up in this field and is due to start scaling up its cat qubits in the near future. Silent Waves, a start-up based in Grenoble, offers TWPA amplifiers for quantum measurements.

# Chapter 5

# Additional information

# 5.1. INTRODUCTION

The economic and political reasoning behind the development of a radically innovative technology such as FTQC can be structured around a number of questions. We will review a number of them. Over and above the technical analysis of the specific applications envisaged for FTQC, the aim is to outline the basis on which the considerable investment currently being made could be justified.

This chapter also looks at the link with High-Performance Computing, the need for *benchmarks* and human capital requirements.

# 5.2. TECHNICAL AND ECONOMIC ANALYSIS OF THE ECOSYSTEM - HISTORICAL COMPARISONS

It is not really relevant to compare the FTQC project with the development stages of mature technologies such as nuclear energy or generations of integrated microcircuits. At this stage, we can refer to other projects with techno-scientific components representing totally new challenges at their time or currently.

### 5.2.1. NUCLEAR FUSION

The impact of eventual success is very important, both strategically and universally. Two technical approaches are possible, and there is considerable uncertainty about the feasibility of both. The cost of the ITER magnetic confinement fusion project has been revised to  $\in$ 25 billion (Wikipedia). The combined cost of the Megajoule and NIF inertial confinement projects (initially aimed mainly at defence strategy) is of the order of ten billion dollars for a similar TRL.

### 5.2.2. SPACE RESEARCH (MANNED FLIGHTS)

In the space sector, the most relevant comparison concerns manned flights, and more specifically the Apollo, Shuttle and International Space Station programmes, whose respective cumulative costs are (source Wikipedia): Apollo (19 billion dollars up to 1970, i.e. 150 billion discounted), Space Shuttle (190 billion discounted), International Space Station (115 billion discounted). At the time, there was talk of justifications based on concrete applications, such as crystal growth or the production of medicines in microgravity. Regardless of the fact that the benefits have not been demonstrated. It is clear that the investment was out of all proportion to t h e potential contribution of such objectives. The real justification is geostrategic. This is of course true from a defence point of view, but in general space exploration is a new frontier and the prospects beyond the horizon are indeterminate and potentially immense. The world's major players have an obligation not to let themselves be overtaken by their competitors, irrespective of any economic considerations.

### 5.2.3. The Genome Project

The scientific, societal and medical impacts were predictable and rapidly materialized. Uncertainty was limited, even if feasibility required the development of new technologies. With a cost of 3 billion dollars, this is clearly one of the major technical-scientific projects that has been the most profitable.

### 5.2.4. CERN (LHC ACCELERATOR)

It is worth adding to these examples a purely scientific project such as the LHC accelerator, which cost around 8.9 billion euros (Wikipedia).

Indeed, it can be argued that in the absence of FTQC research is leading to fundamental advances, both theoretical and experimental, in the profound nature of quantum physics. From this perspective, the question of the relative value of these advances, compared for example with the discovery of the Higgs boson, cannot avoid being asked, even if it is probably undecidable.

### 5.2.5. COMPARISON WITH THE FTQC

The potential economic impact that can be identified today is much lower than for fusion and the genome. From this point of view, the FTQC is more comparable to the example given for space research. Areas of application can be identified, but massive investment in an uncertain area seems disproportionate to the potential impact. However, we have no way of knowing what will happen in the long term, in a world where quantum machines exist. The players, whether states or mega-companies, are seeking to stay ahead of their competitors through geostrategic reasoning that goes beyond short-term economic reasoning.

The level of uncertainty is high and there are more possible channels than in the previous examples, leading to a dispersal of resources, which should cease at the development stage. The level of venture capital spending will be in the region of 1.2 billion worldwide in 2023, down from the record levels of 2.35 billion in 2021 and 2022 (it would appear that some investors are turning their attention to generative AI). Governments have already spent more than twenty billion. Based on an optimistic scenario of around fifteen years before the first production run with competitive performance, a total global cost of around fifty billion seems a minimum estimate.

The space programme was financed by governments. The FTQC is partly financed by private contributions, on the one hand by venture capitalists investing in start-ups, and on the other by industrial groups with massive cash resources. There is certainly reason to fear that the first source will run out as the deadline approaches. As for the second, we have to assume that these industrial supergroups have strategic objectives that extend over decades.

In addition to the strategic motivation, there is also the desire to recruitment of brilliant scientists (see below). If the scale and uncertainty of the task lead to the strategic objective being called into question, the recruitment of brains will in any case have strengthened these groups, whose competitiveness depends on their human capital.

#### 5.2.6. The brain drain problem

Despite all the uncertainties about the possible outlets for FTQC, the subject is extraordinarily attractive to talented young people. One of the socio-economic reasons for continuing to invest in this area is the fear that this talent will evaporate if the subject regresses in France, even though it is a valuable resource, whatever its future activity. Silicon Valley projects have always been a magnet for the world's best researchers and engineers, long before ICTs or Al materialised their promises.

### 5.3. A FEW QUESTIONS

### 5.3.1. IS IT NECESSARY TO GO THROUGH DEVELOPMENT TO ESTABLISH FEASIBILITY?

Here again, space is the closest example. In other cases, technological demonstrations at research level demonstrate feasibility, even if there are still major uncertainties about development costs and times. Interestingly, IBM's semiconductor logic chain in the 1970s demonstrated the gap between the feasibility of a competitive computer and laboratory demonstrations of very good performance.

It seems that for most FTQC production methods, the performance objective will be difficult to achieve without a context of standardised manufacture outside the laboratory: feasibility will undoubtedly require industrial development. Paradoxically, the carbon nanotube approach, which is currently the least advanced in terms of results, has the ambition of carrying out rapid iterations in the laboratory to create a complete system whose individual elements would be sufficiently tested to achieve feasibility. The consequence of this need to develop and the transition to industrial production is that it would require massive investment, with no guarantee of feasibility.

# 5.3.2. WILL A FIRST GENERATION AT THE LIMIT OF THE QUANTUM ADVANTAGE BE USEFUL?

This is doubtful. The contribution is likely to be mainly scientific. Here again, the analogy with space is obvious: Sputnik was useless. It was the first step towards significant military and civil developments much later. Remember that in the former USSR, the launchers in the space programme were a by-product of the much larger strategic defence system.

# 5.3.3. VERTICAL OR HORIZONTAL?

Building an FTQC will require not only a basic technology capable of performance, but also major advances in peripheral developments (communications, input/output, cryogenics, etc.).

This raises the question of verticality. If we can envisage that, while the leading manufacturer with substantial financial resources takes financial risks, this is hardly the case for manufacturers supplying these highly specific peripherals.

For example, IBM's research was initially carried out using dilution refrigerators from the Finnish company Bluefors. It appears that the increase in the number of qubits will lead to new designs for high-power, high-performance cryogenic systems.

Can Bluefors carry out this research using its own funds? Will IBM develop refrigerators in-house? Will IBM acquire Bluefors, depriving the rest of the players of access to these technologies?

The example of the space industry points towards the creation of a sector. The need for a large number of peripheral industrial players makes it difficult to set up such a sector at national level, Europe being the most appropriate level. The political decision will be all the more difficult because of the wide variety of possible technical options, which will remain so for a long time, and the associated uncertainties.

How should R&D investment be framed in terms of budget? In the absence of economic justification at the outset, how should the scientific contribution of an FTQC be viewed in relation to, for example, CERN's international LHC project? What multiplication factor should be added for the hope of economic spin-offs? The R&D component of space developments runs into billions, most of which is not self-financed.

While innovative start-ups play a major role today, do they have the capacity to integrate into the development process, either vertically or downstream? We can only hope that they will find a way to exploit the vast intellectual capital they have amassed in activities with shorter-term economic potential, before investors become discouraged.

#### 5.3.3.1. EVALUATION

Current activities aimed at creating *benchmarks* (see §5.5.) on the progress of FTQCs will have the advantage of countering the frequent announcements of breakthroughs, which mask the fact that the objective is still distant and uncertain. Even if these *benchmarks* concern problems that are unlikely to be the final outlets for FTQCs, they help to maintain a realistic political vision of the field.

#### 5.3.3.2. GENERAL QUESTIONS

From a general point of view, the FTQC over the next twenty years or so is a project with mainly scientific outlets. The first question is: do these outlets justify the investment, with high-energy physics and space research as the benchmarks?

In a few years' time, it is likely that research will be funded mainly by governments, even if sponsoring groups will remain involved longer than venture capitalists. For a country like France, it is doubtful that this will be the future of re-industrialisation.

However, the very high level of French research in this field means that it can be seen as a promising area for the future, in both industrial laboratories and startups, begs the following question:

- Should this research be funded to counter the brain-drain?

If we succeed in countering the *brain-drain* and if a real economic opportunity is discovered:

- Will those who have carried out the development have a decisive advantage, as was the case for silicon microelectronics and the internet and so on, but not for III-V microelectronics?
- 5.4. DEVELOPMENT OF COMPETING TECHNOLOGIES ON THE HORIZON OF THE FIRST FTQC?

Competition from quantum computers will come mainly from traditional computers, which will continue to progress along three axes: integrated circuit hardware technologies, their architectures and software, all three components having progressed in parallel since the beginning<sup>335</sup>. A new situation has recently arisen with the lightning progress, albeit of a different nature, of artificial intelligence.

Integrated circuits continue to progress, but more slowly than according to Moore's Law, with the number of transistors in *central processing units* (CPUs) doubling every two years or so. This growth will not take place as it did at the beginning, thanks to increasingly fine planar lithography. We have gone from lithographic line widths of a few microns to generations today.

<sup>335</sup> John L. Hennessy and David A. Patterson, A New Golden Age for Computer Architecture, *Communications of the ACM* 62, 48 (2019)

of the order of 3 or 2 nm<sup>336</sup>. The initial gain in computing performance increased as the cube of the size reduction factor, according to Dennard's scaling laws. The physical limit imposed by the minimum control voltage of the field-effect transistor has led to a very marked slowdown in the reduction in the expected supply voltage, from 5 V up to 1985 to 0.6 V today, whereas the scaling laws would have led to a 60-fold reduction, depending on the evolution of the lithographic line width and the scaling laws. Another consequence is that a limit has been reached for the electrical power consumed, at 20 W or 300 W for laptop or desktop microprocessors respectively. This power limitation plays a dimensional role in the choice of components (balance between quantities of logic circuits and memories on a circuit) and architectures. Despite these limitations, which appeared around 2002, marking the end of Dennard's laws and leading to the stagnation of clock speeds at around 3 GHz, the number of components is expected to continue to increase thanks to 3D component architectures, the use of new concepts for logic and memory components, the multiplication of highly efficient specialized circuits for a specific task on a single chip, and the stacking of circuit wafers to avoid the energy dissipation required by long-distance interconnections in a plane. The industry consensus, set out every two years in the reports of the International Roadmap on Devices and Systems (IRDS)<sup>337</sup>, predicts that chip performance will continue to increase at the same rate for at least a dozen years, without coming to a halt at the end of that period (it is a constant feature of the IRDS not to predict solutions beyond a dozen years), rising from more than 100 billion transistors to more than 1,000 billion.

In fact, the planar density of circuits is increasing only very slowly: instead of the lithographic line decreasing by a factor of 2 every four years (4-fold increase in density, initial Moore's law), we have only progressed by a factor of 3 in twenty years (!), because of manufacturing difficulties. Yet there is still more integration thanks to new transistor structures and circuit architectures in the vertical dimension, with dimensions that are no longer linked to lithography. As a result, we have introduced "lithographic" generations (7 nm, 5 nm, 3 nm, 2 nm, 1 nm) which have no connection with true line width (currently around 12 nm), but which describe new technological generations with functionalities equivalent to those of planar technologies of 7 nm, 5 nm, 3 nm, 2 nm, 1 nm.

<sup>337</sup> Available at https://irds.ieee.org/editions

The recent arrival of artificial intelligence (AI) poses a far greater threat to FTQC applications than advances in traditional hardware and software alone: the irruption of AI is due to the combination of two essential components, big data to feed the learning of AI systems, and deep neural networks enabling the prediction accuracy of the first neural networks to be greatly improved. In this context, rather than solving a problem exactly from the description of a model, a neural network is built by training on data samples. This constitutes a statistical model of the domain in which the problem is posed, and is approximately solved by inference from the network. The use of graphics processing units (GPUs), equipped with vector or matrix operations, significantly accelerates the training and inference of networks and the optimized calculation of synaptic weights. This trend has been reinforced by the emergence of new specialized processors such as the NPU (Neural Processing Unit), from Intel or Apple for example, or by the emergence of neuromorphic processors<sup>338</sup> which seek to imitate the human brain. The result is extremely powerful tools for tackling a very large number of problems in a statistical way. As the power of specialized processors continues to outstrip that of general-purpose processors, FTQC computers will have to prove their worth in a difficult market. In addition, learning methods are based on the use of data archived in mass memories, which are slow-access by design. The amount of input/output on these memories means that quantum processors are not predisposed to providing significant acceleration.

As long as we content ourselves with statistical results and approximations, and take the necessary precautions to obtain reliable results, we obtain an extremely powerful tool for tackling a very large number of problems. As the power of GPUs/NPUs continues to outstrip that of general-purpose processors, FTQC computers will have to prove their worth in an increasingly competitive market.

<sup>338</sup> Cf. Report by the Académie des technologies : Les technologies matérielles supports du numérique futur - Trois questions 2022 (<u>https://www.academie-technologies.fr/publications/ lestechnologies-materielles-supports-du-numerique-futur-trois-questions/</u>)

difficult for them. In addition, learning methods are based on the use of data archived on mass memories which, by construction, have slow access. The amount of input/output on these memories means that quantum processors are not predisposed to providing significant acceleration.

To determine the future applications where quantum computers are considered to be indispensable, it is therefore essential to take into account the likely developments in classical processor and specialized processor technologies, as well as learning methodologies, and to monitor their progress continuously, as we will be chasing moving targets. This is essential if we are not to disperse our efforts when it comes to implementing the first FTQCs in relevant applications in ten to fifteen years' time.

### 5.5. APPLICATION-ORIENTED BENCHMARKS

With the support of LNE's MetriQs-France national programme for measurements, standards and evaluation of quantum technologies, which is part of France's national quantum strategy, the BACQ project is dedicated to application-oriented benchmarks for quantum computing. The consortium, which includes Thales, Eviden, CEA, CNRS, Teratec and LNE, aims to establish performance evaluation benchmarks that are meaningful for industrial users. The aim of the BACQ project (Benchmarks applicatifs des calculateurs quantiques - Application Benchmarks for Quantum Computing) is to provide a suite of application-oriented benchmarks for an objective, multi-criteria evaluation of quantum computing performance, which is key to industrial applications.

Measuring progress towards the quantum advantage and the realization of its promises, objectively and reliably, is of great interest to potential end-users and crucial to the future development of the field, which is now subject to intense media hype and competition. The challenges, particularly when it comes to making comparable measurements, stem from the diversity of hardware platforms, their specific characteristics in terms of physics and applications, their maturity (which may still be low), and the potential rapid evolution of technologies. There are a number of initiatives to compare the performance of quantum computers. Examples include IBM's Quantum VOLUME and CLOPS, Super-Tech's SupermarQ and Berkeley Lab's Quantum LINPACK. The metrics used in the previous approaches are highly technical and require familiarity with the technology. It is therefore not possible to derive operational performance indicators for the different families of algorithms running on the various existing quantum computers. Dedicated to the entire value chain, from the development of quantum hardware to industrial use cases, the BACQ project complements *benchmarking* initiatives that focus solely on low-level physical hardware criteria. The envisaged benchmark suite will be based on solving several classes of problems covering important application areas of quantum computing that are significant for industrial users: simulation of quantum physics models, optimisation, solving linear systems and prime number factorization. Machine learning could be included in the field of optimization.

These types of applications are generic and could concern different branches of industry and services (chemicals, aeronautics, electronics and energy, for example). Criteria will be defined for solving each problem, some of which will be hardware-independent and others hardware-dependent (low-level): calculation time, latency, problem size, rate, etc.

approximation, probability of resolution, precision, fidelity... Above all, the project also takes energy criteria into account when assessing the consumption performance of the machines.

The proposed methodology involves the aggregation of low-level technical metrics and a multi-criteria analysis using the MYRIAD-Q tool in order to provide operational performance indicators for the various quantum computing solutions and to pinpoint the qualities of service of interest to end users. The aggregation of criteria and multi-criteria analysis enable fully explainable and transparent ratings, comparisons between different quantum machines and with classical computers, and identification of the practical advantages of each quantum machine in relation to specific applications. The BACQ project also covers gate-based machines, in particular *Fault Tolerant Quantum Computing* (FTQC).

The practical approach is to have a suite of benchmarks, adaptive to a certain extent, adapted to the capabilities of the machines available and capable of demonstrating their respective advantages, including, in the longer term, an exponential acceleration of specific algorithms on FTQC machines.

As part of the project, a first action has already been launched concerning Q-Score, which Eviden created based on the MAXCUT optimization problem, to test and validate the *benchmark* on different types of quantum machines.

Sharing the reference suite as widely as possible is an important objective for establishing common reference measurement methods and ensuring the absence of bias with the inclusion of all technologies. Consultations with technology suppliers as well as end-users are essential to develop an instrument that meets needs. Once specified and developed, the reference suite will be available for use by the community.

To achieve a universal tool, it is necessary to establish cooperation with similar initiatives around the world, such as the TQCI (*Teratec Quantum Computing Initiative*) seminars dedicated to the "Overview of forthcoming application-oriented *benchmarks* for quantum computing in France and abroad", which are organized each year by Teratec, Thales and the LNE. A first seminar was organized in 2023 at Thales TRT in Palaiseau and a second seminar in 2024 in Reims.

Standardization will be another means of achieving consensus and widespread adoption, given the European CEN-CENELEC committees JTC 22 WG3 on quantum computing and simulation, the international ISO/IEC committees (JTC1 WG14 on quantum information technologies and JTC-Q on quantum technologies) and the IEEE working groups (P7131 "QC Benchmarking", P3329 "Quantum Energy Initiative" and P3120 "Quantum Energy Initiative"). "Architectures for QC).

In France, standards for quantum technologies are covered by AFNOR/CN QT.

Looking to the future, it is essential to measure progress towards realizing the promise of quantum computing. Application-oriented *benchmarks*, enabling the real performance of quantum computing to be assessed from the user's point of view, seem useful in this perspective. The challenge lies in the diversity of hardware platforms, their specific physical characteristics and applications, their maturity and the potential rapid evolution of the technology. We need to develop an objective, long-term and widely-shared measurement tool that will serve as a common reference.

The evaluation of the practical performance of quantum computing will be envisaged through *benchmarks* close to real applications, significant for industrial (as well as academic) end-users. The main objective is to measure progress towards a practical quantum advantage. In this respect, comparisons should be made between different quantum computing solutions as well as comparisons with current classical computers. This benchmarking initiative will eventually make it possible to highlight the advantages of each quantum computing solution for specific applications.

The benchmark suite will be maintained by LNE, an independent and trusted third party. Through interaction with the end-user community, this tool will be exploited to analyze the results obtained from machines tested with the various benchmarks using the explicable aggregation tool. LNE will draw up a performance catalogue, maintain it over time and update the test definitions. In addition, French, European and possibly international partners will be encouraged to take ownership of this initiative by developing communication tools on the approach adopted and visual representations of the aggregation of results obtained by the machines from the various benchmarks. International dialogue and collaboration on the subject of benchmarking quantum computers will be encouraged so that the approach, supported by MetriQs-France, becomes and remains an international reference. The benchmark encourages the development of international standardization concerning the methods used to evaluate the specifications of quantum machines.

Through the BACQ project, collaborations have been initiated as well as a dialogue with other international benchmarking initiatives: Fraunhofer IKS & FOKUS (BenchQC Project) in Germany, TNO & TU Delft (QPack Project) in the Netherlands, Qilimanjaro (CUCO Project) in Spain, QuIC (Use cases WG) in Europe, HamLib Project (Intel, LBNL *et al.*) in the United States, Fonds unitaire (Metriq project) in the United States and QED-C (QC Benchmarking WG) in the United States.

The results of benchmarks also feed into the discussions of standardization initiatives: the AFNOR National Committee on QT in France, CEN-CENELEC (JTC22 QT/WG3 Quantum Computing & Simulation) in Europe, ISO & IEC (JTC1/WG14, JTC3) at international level and IEEE (P7131 "QC Benchmarking", P3329 "Quantum Energy Initiative" and P3120 "Quantum Energy Initiative"). "Architectures for QC") in the United States.

The European Commission has set up the EQCBC (*European Quantum Computing Benchmark Coordination Committee*), which is working on a *white paper* to be published shortly on the subject of "*Systematic benchmarking of quantum computers: status and recommendations*". Thales, LNE and Alice&Bob are the French co-authors of this publication. The aim of the EQCBC is to coordinate national actions in the European Union, and in particular France, Germany and the Netherlands, to encourage the emergence of a European *benchmark* for quantum computers.

#### 5.5.1. Recommendations concerning benchmarking

- Use the national BACQ project (Application *Benchmarks* for Quantum Computers) of LNE's MetriQs programme to monitor the performance of FTQC quantum computers.
- Coordinate *benchmarking* activities for FTQC calculators at European level, by promoting cooperation between France and Germany (Fraunhofer IKS) and the Netherlands (TNO).

### 5.5.2. References

For further information, please consult the following references:

BACQ: Application-oriented Benchmarks for Quantum Computing, *arXiv*: 2403.12205v1 [quant-ph]; <u>https://doi.org/10.48550/arXiv.2403.12205</u>

Bench-QC: *Application-driven benchmarking of quantum computers*, Fraunhofer Institute IKS, Germany; <u>https://www.iks.fraunhofer.de/en/projects/ bench-qc-application-driven-benchmarking-of-quantum-computers.html</u>

QPack: *Benchmark for quantum computing*. Delft University of Technology, Delft, The Netherlands; <u>https://github.com/koenmesman/QPack</u>? tab=readme-ov-file

CUCO project 'Computacion Cuantica En Industrias Estrategicas', Spain ; https://t/www.cuco.tech/en/projec

Metriq "Community-driven quantum benchmarks", Unitary Fund; <u>https://</u>metriq.info/

QED-C benchmark; <u>https://github.com/SRI-International/QC-App-Oriented-Benchmarks</u>

DARPA QBI (*Quantum Benchmarking Initiative*); <u>https://www.darpa.mil/ news-events/2024-07-16</u>; <u>https://www.darpa.mil/news-events/2024-08-15</u>

HamLib : A library of Hamiltonians for benchmarking quantum algorithms and hardware, in 2023 IEEE International Conference on Quantum Computing and Engineering (QCE), 2023. *arXiv*: 2306.13126v3 [quant-ph]; <u>https://doi.org/10.48550/arXiv.2306.13126</u>

E. Vergnaud, F. Schopfer, F. Barbaresco, first TQCI Seminar "Overview of upcoming application-oriented benchmarks for quantum computing in France and abroad", Thales Research & Technology, Palaiseau, France, 11 May 2023; <a href="https://teratec.eu/Seminaires/TQCI/2023/Seminaire\_TQCI-230511.html">https://teratec.eu/Seminaires/TQCI/2023/Seminaire\_TQCI-230511.html</a>

E. Vergnaud, F. Schopfer, F. Barbaresco, 2nd TQCI Seminar "Overview of upcoming application-oriented benchmarks for quantum computing in France and abroad", 4th & 5th June 2024, Reims, France ; <u>https://teratec.eu/</u>activites\_quantiques/TQCI\_240604\_programme.html

Jeanette Mirian Lorenz, Thomas Monz, Jens Eisert, Daniel Reitzner, Félicien Schopfer, Frédéric Barbaresco, Krzysztof Kurowski, Ward Van Der Schoot, Thomas Strohm, Cécile M. Perrault, Martin Knufinke, Systematic benchmarking of quantum computers: status and recommendations, EQCBC (European Quantum Computing Benchmark Coordination Committee) preprint, to be published

Anastasiia Butko, Katherine Klymko, Daan Camps, Nicolas SawayaAuthors Info & Claims, HamPerf: A Hamiltonian-Oriented Approach to Quantum Benchmarking, CF '24 Companion: Proceedings of the 21st ACM International Conference on Computing Frontiers: Workshops and Special Sessions, Pages 133-138; <u>https://doi.org/10.1145/3637543.3653431</u>

5.6. LINKS WITH HIGH-PERFORMANCE COMPUTING

We will use the term HPC (*High Performance Computing*) for "traditional computing", with the general meaning of "computing resources that go beyond standard equipment such as desktops or laptops", up to extreme scales in specialized computing centers (currently "exascale" for the largest systems)<sup>339</sup>.

A quantum computing system needs classical control computing, which is in a way low-level, close to the control electronics of the system components. We exclude it from the field of HPC.

HPC can play a role in quantum computing at three levels:

1. As a resource and method for simulating/emulating quantum systems, including quantum circuits;

339 Top 500: https://top500.org/

- 2. As a means of performing auxiliary processing for quantum computing, typically in connection with error correction or compilation;
- 3. As a co-processor for certain problems that can combine a classical algorithm and a quantum part ("hybrid" calculation).

When we talk about the FTQC/HPC combination, we usually think of the third level, but we think it would be useful to comment briefly on the other situations, as the destinies of classical and quantum computing seem to be permanently linked.

### 5.6.1. EMULATION/SIMULATION

In general, physical systems modelled in a quantum manner can benefit from HPC simulation, particularly the qubits that make up quantum circuits and gates. At present, the most advanced emulations deal with the order of 55 qubits exactly (in "state vector" mode), and the order of a few hundred with approximations (based on tensor networks) - the largest calculations carried out are of almost exascale class. In this way, we can gain some insight into how quantum computers or algorithms should work.

Specialized hardware and software make this type of approach more affordable, and enable quantum computers to be emulated down to a few dozen qubits without requiring the largest HPC computers. The memory of a computing cluster dedicated to such an approach is the first factor to be taken into account (*a priori*, each additional qubit requires the memory to be doubled). Such clusters benefit from HPC methods for SMP approaches (computing clusters with large shared memory) and can also include GPUs to accelerate processing. This is the case, for example, with EVIDEN's Qaptiva systems, 'HPC appliances' that can also emulate noise and errors, in very compact physical formats of the order of a computing cabinet. (Qaptiva emulates up to 41 qubits in state vector mode and over 75 qubits in noisy mode)<sup>340</sup>.

<sup>340</sup> Qaptiva™ - Atos (https://atos.net/en/solutions/high-performance-computing-hpc/quantumcomputing-qaptiva)

Another example is Qubit Pharmaceutical's Hyperion emulator, which is based on Nvidia DGX servers and supports up to 40 qubits in state vector mode<sup>341</sup>. This type of approach makes it possible to virtually execute complete, ideal FTQC quantum programs, which may even be noisy, and to prepare and develop them in the absence of, or in advance of, the availability of a FTQC quantum computer. A suitably constituted environment makes it possible to switch from emulation mode to real mode using the same source and intermediate quantum codes.

By definition, these computational approaches will be overtaken by the achievement of a real quantum advantage, which cannot be simulated/emulated conventionally in practice. But they will always be of interest, for studying and developing quantum sequences and algorithms, in particular, for observing the evolution of data internal to the calculation during the process, which is not possible on a quantum computer.

### 5.6.2. AUXILIARY TREATMENTS

As mentioned in section 3.2.3, data processing for error correction (detection of error syndromes) can be very intensive and will require conventional hardware optimizations, such as the addition of GPUs, FPGAs and even ASICs dedicated to quantum processing. Riverlane is an example of a company focusing on this aspect<sup>342</sup>.

The organization of realistic quantum computation sequences may also require circuits or portions of circuits to be decomposed and compiled/transpiled on the fly, for iterative approaches by nature or because a complex circuit does not fit into a single execution pass (physical limits on the width and depth of circuits that can be realized). This may also require the dedication of conventional computing resources.

<sup>341</sup> Qubit Pharmaceuticals and Sorbonne University achieve a major scientific breakthrough by simulating quantum calculations at more than 40 qubits on conventional computers | Sorbonne Université| Sorbonne université (sorbonne-universite.fr) - <u>https://www.sorbonne-universite.fr/en/presse/qubit-pharmaceuticals-and-sorbonne-university-achieve-major-scientificbreakthrough</u>

<sup>342 &</sup>lt;u>https://www.riverlane.com/quantum-error-correction-stack</u>

It must also consider that compilation, which includes transpilation and optimization, represents a variable operating cost, unlike classical computing where it is generally a fixed cost. Since quantum computation is the equivalent of *in-memory processing*, with data loading using quantum gates directly integrated into the body of the algorithm (*ansatz* of a variational algorithm, Hamiltonian used in the exponentiation of a phase estimate), the code is almost always specific to the dataset used, with the exception of oracle-based algorithms using qRAM memory access to classical data.

These resources and approaches are similar to HPC, but in a variant "Embedded", where the HPC is subordinate to the quantum system, to ensure sufficient precision. Beyond low-level control constraints, and before we can talk about truly hybrid computing, quantum computing needs this kind of classical logic integrated as closely as possible.

This explains, for example, the partnerships between Nvidia, Quantum Machines and Alice&Bob. On large FTQC computers, it is not out of the question to add clusters with a large number of DGX or equivalent servers.

### 5.6.3. Hybrid HPC/FTQC computing

What we are talking about here is hybridization at the level of algorithmic processing, with a sort of parity between classical and quantum computing, each invoked where it is best suited to contribute to the execution of a task: two radically different computing models coexist. The reason for this is pragmatic: the aim is to make the best possible use of resources.

It is also linked to the very nature of certain quantum algorithms, such as phase estimation, which requires fairly costly classical preparation (based, for example, on the DFT method). It is also linked to the very nature of certain quantum algorithms, such as phase estimation, which requires a fairly costly classical preparation (based, for example, on the classical DFT method).

The granularity and coupling in terms of space and time of the breakdowns and distribution of roles are variable. To date, they have been suggested in many cases, but in practice there has been little or no experience of them.

The simplest forms are, for example, the case of Shor or Simon algorithms, with classical pre- and post-processing of the data, with a "simple" quantum processing core. This is not a highly intertwined hybrid, but it does illustrate the inseparable nature of HPC and FTQC in practice.

Hybridization is more marked in the case of variational quantum algorithms (VQE, QAOA, QML). Here, the vision is that of a classical processor (or cluster) calling a quantum (co-) processor to solve a sub-problem and optimize a cost function iteratively, with the general logic, as well as the significant processing involved in the resolution, being taken into account in a classical way.

In fact, Shor's algorithm is very hybrid, involving a lot of classical arithmetic in the sense of a global factor search loop and using quantum in each loop solely for function period estimation (based on modular exponentiation and an inverse QFT, with the quantum computation cost concentrated in modular exponentiation).

The VQE approach can be used in quantum chemistry, quantum simulation and optimization. It is an algorithm that alternates between classical and quantum to find the state (usually fundamental) of a given physical system. From a guess, the quantum processor calculates the expectation value of the system with respect to an observable, often the Hamiltonian, and a classical optimizer is used to improve the estimate by modifying the parameters of the quantum circuit (known as the *ansatz*).

The similar QAOA approach expresses optimization problems in terms of finding the highest energy configuration of a spin Hamiltonian, and also alternates classical parameter optimization with quantum cost function estimation.

Where this type of approach is concerned, we find many of the elements of HPC:

- management of computing resources in a resource-sharing environment (allocation of CPU/GPU [\*PU] and QPU, their availability, task scheduling, etc.);
- communication between \*PUs CPU, GPU... conventional and QPU;
- ways of expressing classical/quantum alternations general logic uses a classical programming language, which can be augmented by quantum directives [pragma] or calls to specialized libraries;
- underlying *runtime* support.

Within the FTQC framework, chemical simulations based on the phase estimation algorithm [QPE] require significant classical pre-calculation. The quantum algorithm needs to be fed with a classically calculated estimate that is as close as possible to the desired result. The more accurate this calculation, the more expensive it is classically. And the less accurate the calculation, the more the quantum calculation will diverge probabilistically from the desired result.

A great deal of work and experimentation has been carried out to date, but no clear standardized solutions or approaches have yet emerged.

It should also be noted that at this early stage of development, classical and quantum processors are packaged and configured in separate systems [racks,

chassis...]. Hybrid applications can certainly

accommodate 'remote' connections with workflows running

partly on different sites, as long as the data exchanges are limited and not timecritical. However, locating \*PUs and QPUs on the same HPC site has undeniable advantages:

- Scale-up of communications facilitated if necessary by local network;
- easier joint maintenance of systems;
- common IT security in the computing centre's HPC bubble.

As a quick conclusion to this section, the cohabitation of classical processing and quantum processing is here to stay. In its form of HPC/FTQC hybridization at an algorithmic and application level, there is still a lot of development and experimentation to come:

- development of *middleware* and adaptation of HPC+QC software stacks.
- experimentation and feedback on the various qubit implementations, as they mature and become available, in a "production" HPC environment. We need a better understanding of the potentially variable suitability of qubit technologies for the targeted applications, as well as the details of the behavior of qubits and circuits, particularly in terms of time, with an impact on the way in which they are coupled with conventional computing, depending on the duration and latency of circuit preparation and execution, compared with the duration of HPC computing loops and tasks.

### 5.7. HUMAN CAPITAL

To support and accelerate the development of FTQC in France, human capital is essential. Training courses need to be adapted and developed to meet the growing need for skills in quantum technologies, whether in terms of engineers, researchers, teacher-researchers, technicians or managers. This involves developing initial training courses or including new modules in existing courses, as well as developing continuing training courses, in partnership with vocational training and industrial players: quantum technologies are in fact creating a renewed need for lifelong training in order to learn or re-learn the fundamentals of quantum physics or the constant developments in quantum technologies. As the QuantEdu France project financed by France 2030 has already done, we need to create a strong training dynamic between those involved in training, research and industry.

What's more, there is a shortage of skilled workers in quantum computing, whether in software engineering or research in physics and mathematics, is rare.

In some areas, such as error correction, the need far exceeds the human resources available. This tension is encouraging a war for talent, both between different players in the same region [between the academic sector and companies, or between companies themselves] and between different regions of the world, first and foremost between Europe and North America. In this highly competitive environment, a well-developed and dynamic ecosystem of start- ups, major industrial groups and universities is essential to attracting and retaining skilled people, especially as it is difficult for some players to offer salaries that are competitive with those in other parts of the world. In this respect, regional clusters that bring together the players present in a given area, such as the quantum houses, can be an asset. A more detailed analysis will be carried out in 2025 and developed in the new version of the report.

### 5.8. FINANCING STRATEGIES

Public and private investment in quantum computing in France is relatively high and is driven by a variety of factors. It is crucial to understand this and to compare the French situation with the rest of Europe and internationally in order to assess the performance of public action in this area. It should be noted, however, that in spring 2024 France set up a specific funding programme for the FTQC, called PROQCIMA.

#### PROQCIMA: an original tool for financing disruptive innovation in France

Financed by the "France 2030" plan of the General Secretariat for Investment and placed under the supervision of the Digital Defence Agency [AND] [part of the Directorate General for Armaments, Ministry of Defence], PROQCIMA will provide two fault-tolerant quantum computer prototypes with 128 logic qubits by 2032. With a budget of €500m, PROQCIMA is setting up long-term framework agreements (more than eight years) with five start-ups from French research, each developing a specific technology: Alice&Bob, C12, Pasqal, Quandela and Quobly. PROQCIMA takes the form of a competition to select the best-performing technologies in three stages:

- In 2028, only the three best-performing technologies will continue to benefit from PROQCIMA funding to develop the best qubits capable of going to scale;
- In 2032, only two prototypes of 128 qubit logic [fault-tolerant] computers will be selected;
- In 2035, the ambition is to go to industrial scale with 2048 logic qubits.

In a context where it is impossible, given the current state of knowledge, to select the technologies that will ultimately deliver the best performance, which is likely to require major long-term research and development programmes, PROQCIMA makes it possible to put competitors back on the same starting line, to avoid dispersing public funding for FTQC in France, to assume simplified governance via the AND and to confirm the need to demonstrate the transition to industrial scale, while accepting the risk of failure.

### Conclusion

The working group, which was set up in mid-2023, has initiated a collective process of reflection involving a large part of the French ecosystem around the Académie des technologies. The mass of information and data collected has enabled us to identify key scientific and technical challenges, such as scaling up to the necessary quality and quantity of physical qubits, developing error-correcting codes and interconnecting these elements.

In particular, we still need to analyze the state of development of the algorithms and applications, as well as the compilation environments. The same applies to the technical architecture of the systems to be built. The economic challenges, such as costs and value creation, enabling technologies, the respective roles of the various public and private players, comparisons with other European and international dynamics, and skills requirements remain to be explored in greater depth.

In order to make fault-tolerant quantum computers a reality in the long term, a great deal of scientific and technical progress is still needed, not least in order to establish realistic industrial strategies.

It is therefore essential to continue the collective reflection, and to complete and regularly update the information collected to ensure a complete and up-to- date understanding of the challenges associated with fault-tolerant quantum computing. A second report will be drawn up in 2026.

# Appendices

## A. Glossary

Algorithm: a method of solving a problem using a finite sequence of operations or instructions. The word comes from the name of the ninth-century Persian mathematician Al-Khwârizmî.

Hybrid quantum algorithm: an algorithm that combines conventional processing on traditional computers with processing on quantum computers, where necessary.

Variational algorithm: generic hybrid quantum algorithm using a classical optimizer which is used to train a parameterized quantum circuit. Its variants are VQE [*variational quantum eigensolver*] for chemical simulation, QML or quantum machine learning and QAOA [*Quantum Approximate Optimization Algorithm quantum*] for optimizations.

AOD: *Acousto-optic* light deflector which modifies the orientation of light beams to create optical tweezers used to move and trap atoms at will in a vacuum, and generally in two-dimensional geometric structures.

ASIC: Application-Specific Integrated Circuit, a type of electronic circuit designed and optimized to perform a particular task or application. It is specially manufactured to meet the needs of a single application or a set of specific applications. This makes it extremely efficient in terms of speed, power consumption and cost per unit when mass-produced. ASICs could replace FPGAs for the electronic control of qubits, particularly superconductors.

Rydberg atoms: excited state of an atom with one or more electrons and whose principal quantum number n (index of the electron layer in of the atom, which is an integer between 1 and the number of electron layers in the atom) is high, in excess of ten.

The transitions between levels are in the microwave domain, and they are therefore manipulated in this frequency range.

Cold atoms: atoms cooled to very low temperatures, generally using techniques based on lasers and the Doppler effect. They are used in certain types of cold atom quantum computers. The atoms used are neutral (not ionized) and quite often rubidium, an alkali metal.

Quantum advantage: occurs when a quantum computer executes a process faster than its optimum equivalent adapted to a supercomputer. This advantage can be applied to aspects other than computing time. For example, a quantum energy advantage relates to energy consumption instead of computing time.

BACQ: Application *benchmarks* for quantum computers.

Blueprint: a technology provider's roadmap for the coming years.

Boson (simulation): solving a quantum physics problem using quantum computation or an analogue quantum simulator.

BPP: *Bounded-error Probabilistic Polynomial*, class of decision problems that can be solved by a classical probabilistic algorithm operating in polynomial time with a bounded error rate.

BQP (problem class): complexity class of problems that can be handled by quantum algorithms. Stands for "*bounded-error quantum polynomial time*". This is the class of problems that can be solved in polynomial time relative to the size of the problem, with a probability of obtaining an error of no more than a third of the results. This class lies between the P class (problems that can be solved in polynomial time on a classical machine) and the NP class (problems for which a solution can be verified in polynomial time on a classical machine).

Break-even: threshold above which, with a given quantum error correction code, we obtain logical qubits with an error rate lower than the error rate of its physical qubits.

Chiplet: specialized integrated circuit, usually integrated inside a single package. Unlike a traditional monolithic chip, which groups together all the functions, a chiplet system divides the functions between several independent semiconductor modules, each optimized for a specific task (calculation, memory, communication, etc.).

Circulator: three-input isolator which prevents the signal amplified by a TWPA after reading a qubit from disturbing the qubit being measured.

CLOPS: *Cycle Layer Operations Per Seconds*. Number of quantum gate cycles a quantum computer can execute per second. In IBM's quantum computers of 2024, this will be around 4,000 cycles per second.

Cluster state: the starting point for an MBQC (*Measurement Based Quantum Computing*) calculation, with a set of interacting qubits on which measurements are then taken.

Cluster states: intricate states of n-qubits.

CMOS: a common semiconductor manufacturing technique used to produce processors and memory, and reused to create qubits that manipulate electron spins.

CNOT: two-qubit gate which inverts the sign of the second target qubit according to that of the control qubit.

Bosonic code: hardware system that implements built-in error correction through the use of bosonic modes, which are quantum harmonic oscillators with continuous energy levels. This includes *cat-codes* and GKP qubits.

Surface code: type of quantum error correction code that supports relatively high errors and is adapted to two-dimensional qubit structures where the qubits are connected to their nearest neighbor.

Error-correcting codes: describes both logical methods and physical architectures for circumventing errors generated by noise in universal quantum computing and a variant of post-quantum cryptography.

### CPU: Central Processing Unit of conventional computers

Crosstalk: decoherence resulting from the influence of qubits on neighbouring or distant qubits, either through electronic control signals or physical interactions between qubits.

Cryo-CMOS: silicon CMOS electronic circuit designed to operate at very low temperatures. It is generally used to provide part of the electronic control of qubits located on another chip.

Cryogenics: cooling technique. Very low temperature cryogenics is used in a large proportion of quantum computers, all those based on electrons or cold atoms. The temperatures required to stabilize qubits and reduce their error rate are very close to absolute zero: between 5 and 20 mK. The systems most commonly used are dilution refrigerators using helium-3 and helium-4. Cryogenics is also used for low-noise, high-performance photon-based qubit reading systems and for photon generation.

Decoherence: marks the end of the coherence of a quantum object or qubit, i.e. the memory of its initial quantum state. It is caused in particular by interactions between qubits and their environment. The expression coherence time (time during which qubits are in a state of superposition and entanglement with other qubits) or decoherence time (time at the end of which this superposition and entanglement come to an end) is often used interchangeably, which amounts to the same thing.

Crosstalk: see crosstalk.

Dilution: part of a cryostat used to cool superconducting qubit chips down to a temperature of around 15 mK.

Quantum emulator: software and/or hardware system using a conventional computer to run software intended for a quantum computer. This makes it possible to test quantum programmes without having a quantum computer. Execution speed is not as good as on a quantum computer, especially once you exceed a few dozen qubits. And beyond around fifty qubits, the capacity of conventional machines is insufficient to carry out this kind of emulation. Emulation should not be confused with quantum simulation. The latter simulates quantum physics phenomena using an analogue quantum computer.

FBQC: *Fusion-Based Quantum Computation*, a variant of the measurementbased calculation method used in photonics and proposed by PsiQuantum. It is based on cluster states grouping together entangled photons, on which the calculation is based on successive photon measurements, and on mergers between clusters, which consist in performing Bell tests between clusters, equivalent to two-qubit gates generating entanglement.

Fidelity: characterizes in percentage terms the probability of obtaining a correct result from an operation on qubits. This refers to an action such as initializing a qubit, a quantum gate with one or two qubits, and measuring the state of a qubit. We are generally interested in the fidelities of two-qubit gates, which are lower, compared with the fidelities of one-qubit gates, which are always better. The fidelity of the reading determines the quality of the calculation results and the effectiveness of error correction, which is also based on measurements.

Figure of merit: indicator characterizing a system. The fidelity of a two-qubit gate characterizes a qubit chip.

Fluxonium: type of superconducting qubit.

FPGA: *Field-Programmable Gate Array*, an integrated circuit that can be configured after manufacture, unlike ASICs (*Application-Specific Integrated Circuit*), which are designed to perform a fixed task. An FPGA is made up of a network of interconnected programmable logic blocks, enabling it to execute almost any digital algorithm depending on the configuration applied. However, these circuits are less powerful than ASICs in terms of computing speed and energy efficiency.

FTQC: *Fault-Tolerant Quantum Computer*. Fault-tolerant quantum computer. *A priori*, with a very large number of qubits and error correction codes.

GHZ: superposed Greenberger-Horne-Zeilinger state with three or more qubits, which makes it possible to demonstrate the non-existence of hidden variables in the quantum entanglement of at least three particles and with a finite number of measurements. The concept dates from 1989 and its experimental validation from 1999.

GKP: type of self-correcting bosonic qubit. Corrects the qubit's flip error and phase error at the same time.

Grover (algorithm): quantum algorithm for finding an element in an unindexed array.

Hadamard (gate): gate used to generate a superimposed state between 0 and 1 in a qubit.

Heuristics: method used to solve problems in a practical, quick and approximate way, when exact or optimal solutions are difficult or impossible to find within a reasonable time.

Hilbert (space): real or complex vector space with a Euclidean or Hermitian scalar product, used to measure distances and angles and to define orthogonality. It is a generalization of the concept of three-dimensional Euclidean space. In quantum mechanics, the state of a quantum is represented by a vector in a Hilbert space with as many dimensions as the number of basic states (or observables) of the quantum. These geometric spaces are used in particular to measure generalized lengths and angles, make projections onto dimensions and define orthogonality between abstract 'vectors'.

HHL: Harrow-Hassidim-Lloyd, quantum algorithm for solving linear equations.

### HPC: High-Performance Computing.

Entanglement: link between two quantum objects that are connected in such a way that a change in one causes a change in the other. This process is used to link qubits together using two- or three-qubit quantum gates in quantum computers. It is also used in quantum cryptography and telecommunications systems based on entangled photons, which are used in QKDs. Their measurement is correlated, but random. Entanglement cannot be used to define information at point A and entangle it with point B. However, entanglement can be used to teleport the state of a qubit from point A to point

B. Since this requires the use of two conventional bit information channels in addition to the optical channel of entanglement, the information about the qubit cannot be transmitted faster than light.

lon: a non-neutral atom with a positive or negative electric charge. It is negative if its number of electrons exceeds that of the protons (anions) and positive in the opposite case (cations).

Trapped ion: these are ions used in certain types of quantum computer. They are generally magnetically or electrically trapped and their state is controlled using lasers.

Josephson effect: superconducting effect used in the qubits of superconducting quantum computers such as those made by IBM and Google.

LDPC: *Low-Density Parity Check*, a conventional error correction code used to transmit data reliably over noisy communication channels. These codes are called "low parity density" because their parity check matrix has a low number of non-zero elements, which makes them computationally efficient to process. Their quantum equivalent is qLDPC.

MAXCUT: combinatorial optimization problem consisting of dividing the vertices of a graph into two distinct subsets so as to maximize the number (or weight) of edges linking these two subsets.

MBQC or MQCM: *Measurement Based Quantum Computing*, a quantum computation method invented in 2001 by Robert Raussendorf and Hans Briegel that uses a large number of qubits embedded in two-dimensional grids and in which qubit state readings are taken to modify the grid structure. These measurements are also used to guide the algorithm.

MOT: magneto-optical trap, used to trap cold atoms in a vacuum chamber using lasers.

NISQ: *Noisy Intermediate-Scale Quantum*, the name given to current quantum computers and those to come in the near future, which are intermediate in size in terms of the number of qubits (a few dozen to hundreds) and subject to quantum noise which limits their capabilities. This name was coined by the American researcher John Preskill.

NP (problem class): class of problems whose solution can be verified in a time that is polynomial with respect to the size of the problem. This includes so- called exponential or intractable problems, whose solution time is exponential with respect to their size. A quantum computer can be used to solve some NP problems.

NP-complete (problem class): decision problem for which a solution can be verified in polynomial time and for which all problems in the NP class can be reduced to this one via a polynomial reduction. This means that the problem is at least as difficult as all the other problems in the NP class. The travelling salesman and backpack filling problems are NP-complete problems.

NP-hard (problem class): a problem to which any problem in the NP class can be reduced by a polynomial reduction. If it is also in the NP class, it is said to be an NP-complete problem. If  $P \neq NP$ , then NP-hard problems cannot be solved in polynomial time.

NPU: Neural Processing Unit.

Unitary operation: operation on a vector which preserves its length. In the case of qubits whose vector always has a length of 1, unitary quantum gates apply a transformation to it which preserves this length. In the representation of qubits in the Bloch sphere, the operation rotates the vector representing the state of the qubit in this sphere.

P (problem class): problem that can be solved in polynomial time with respect to its size, on a deterministic Turing machine.

Permanent: the permanent of a square matrix is an operation that associates a number with this matrix by combining its elements in a very structured way. We consider all the possible paths that run through each row and each column once, and for each of these paths, we multiply the corresponding elements of the matrix. The permanent is then obtained by adding up all these multiplications. This operation is particularly important in quantum physics, as it is used to describe certain interference probabilities between identical particles, such as photons, in optical systems. It is extremely difficult to calculate for large matrices.

Phasers: electrically controlled optical devices that modify the phase of an optical signal.

Universal quantum gates: sets of quantum gates from which all other quantum gates can be reproduced.

Quantum gates: operations that manipulate the state of qubits by acting on one or more qubits. Multi-qubit gates (Toffoli, Friedkin, etc.) exploit the principle of quantum entanglement. Operations

of quantum gates are generated by physical actions on the qubits

depending on their nature. For superconducting qubits, this involves sending microwaves between 5 and 10 GHz via electrical conductors. For trapped ions, laser-controlled operations are used. For CMOS qubits, electrical voltages are used. For qubits based on particles with mass (electrons, ions, cold atoms), quantum gates are used.

act on the qubits, but they do not move. In the case of photon-based qubits, the photons circulate and pass through quantum gates that modify their state (phase, frequency, etc.).

T-Gate: a gate which performs a  $45^{\circ}$  phase transformation, making it possible to construct a set of universal gates. This gate is corrected using a special error correction code, which is very expensive compared with the error correction codes for X, Z and CNOT type gates.

Toffoli gate: a three-qubit gate that can be used to create a universal set of quantum gates.

Post-selection: technique for selecting the result of an operation by eliminating results considered to be poor.

PQC: *Post Quantum Cryptography*, cryptography resistant to algorithms designed for quantum computers. It is based on the use of public keys that cannot be decomposed using conventional computers or quantum computers. This is linked to the fact th a t it is a problem of "NP-difficult".

PSPACE: complexity class in theoretical computer science that groups together decision problems that can be solved using a quantity of memory (space) that is polynomial with respect to the size of the input. Unlike other classes such as P or NP, the PSPACE class focuses on the memory constraint used, independently of the time needed to perform the calculations.

QAOA: *Quantum Approximate Optimization Algorithm*, variational algorithm used to solve combinatorial optimization problems.

QCCD: *Quantum Charge-Coupled Device*, an architecture for quantum computers based on trapped ions that exploits the ability of certain ions to interact by electrical coupling to carry out quantum operations, while allowing great modularity and scalability.

QFT: *Quantum Fourier Transform*. Quantum variation of the Fourier transform. The classical Fourier transform is used to break down a signal (as in audio) into frequencies (or frequency spectrum). The QFT does this on a sequence of integers and determines its highest observable frequency.

qLDPC: *quantum Low-Density Parity Codes*. Very fashionable error correction codes that require a low number of physical qubits per logical qubit at the cost of better connectivity between remote qubits.

### QML: Quantum Machine Learning.

QPE: *Quantum Phase Estimate*. FTQC algorithm used to obtain the phase of a Hamiltonian and indirectly the energy of the rest state of an N-body quantum system.

### QPU: Quantum Processing Unit.

Bosonic qubit: A superconducting bosonic qubit is generally based on a low-loss linear microwave resonator acting as a bosonic quantum memory. Alice&Bob cat qubits are a special case of a bosonic qubit.

Logical qubit: assembly of physical qubits implementing a hardware or software error correction device. From the software developer's point of view, it exhibits the behaviour of a physical qubit whose fidelity is better than that of the latter. The fidelity of logic qubits depends in particular on the number of physical qubits they contain, the quality of the error correction codes and the stability of the fidelity as the number of qubits increases.

Qubit or physical qubit: elementary unit of information in quantum computers. It stores a quantum state associating two distinct states of a particle or a quantum system based on several particles (electron spin, superconducting state of a group of electrons, energy level of an atom or a trapped ion, polarization or other property of a photon). Its mathematical representation is a vector comprising two complex numbers. QUBO: *Quadratic Unconstrained Binary Optimization*. A technique for formulating an optimization problem which can then be used by automatic conversion by analogue computers such as D-Wave's *quantum annealers* or Pasqal's quantum analogue computers.

Qudit: is a generic form of qubit that has n possible quantum states instead of two. The approach is rarely used, at least in quantum computers outside research laboratories.

Qutrit: this is a form of qubit which, instead of having two possible quantum states, has three. It is a special case of qubits.

Dilution refrigerator: name given to the cryostats in most quantum computers, which are used to cool the quantum computing chip to less than 20 mK. Dilution is linked to the fact that these systems use a mix of two helium isotopes: 3 and 4, which are diluted in each other in the refrigeration loop, the two isotopes having slightly different properties. A helium 4 cryostat only goes down to 4 K, a helium 3 cryostat goes down to 300 mK, while a cryostat using both generates a temperature that goes down to 10 mK. It should be noted that the most common variant is the "dry" dilution refrigerator, as opposed to the "wet" version. This version uses less helium and leaves more space in the 'candlestick' for electronic and quantum equipment.

Register: set of bits or qubits.

Rydberg (atoms): excited state of an atom possessing one or more electrons and whose principal quantum number n (index of the electron layer in the atom, which is an integer between 1 and the number of electron layers in the atom) is very high. These atoms are generally large, proportional to n2, and have very strong interatomic interactions. These interactions allow the entanglement of atomic subsets or even single atoms. These atoms have been used by Serge Haroche's team to non-destructively detect the presence of a photon in a cavity, and so study quantum decoherence. But hydrogen can also

be a Rydberg atom if it is excited to high energy levels, causing its electron to move to a higher numbered quantum layer.

SAT: class of logic problem or Boolean satisfiability problem, of order 0 logic. It is a decision problem, which, given a propositional logic formula, determines whether there is an assignment of propositional variables that makes the formula true.

SFQ: *Single Flux Quantum*, electronic circuit technique operating at low temperature and superconducting. Considered for low-temperature control of qubits such as superconducting and silicon qubits. Reduces the amount of wiring used to control qubits.

SGPI: General Secretariat for Investment

Shor (algorithm): quantum integer factorization algorithm invented by Peter Shor in 1994. In theory, it could be used to break RSA public keys by decomposing them into prime numbers.

Silicon 28: silicon isotope used to create silicon *wafers* suitable for creating silicon qubits. Silicon 28 has zero spin, which has no effect on the spin of the trapped electrons used to manage the qubits. It is purified in Russia and can then be deposited as a thin layer in the gas phase on conventional silicon.

SLM: spatial light modulator that controls the phase of photons emitted by lasers. Used to control cold atoms.

Spin: state of a particle describing its rotation on itself or a magnetic moment. Applies to electrons, neutrons and atoms. The spin of composite particles is the sum of the spin of its components. A proton and a neutron have a spin of  $\frac{1}{2}$ . An electron has a spin of +1/2 or -1/2.

Nucleus spin: spin of the nucleus of an atom.

SMP: *Symmetric Multiprocessing*, an architecture in which several processors share a common processor space.

memory and access this memory symmetrically, i.e. with equal access rights.

SQUID: Superconducting Quantum Interference Device, an extremely sensitive device capable of detecting very weak magnetic fields. It is based on the principles of superconductivity and the Josephson effect. It is used in particular for the precision capture of magnetism in the circuits of certain quantum sensors.

Superposition: fundamental principle of quantum mechanics according to which a quantum system can be prepared in the form of a combination of several basic states, such as the up or down spin of an electron in a given direction or the excitation levels of an atom. This is linked in particular to wave-particle duality and the fact that two wave functions of a quantum object can be combined linearly to create a new state of that object.

Superconductivity: the ability of matter to conduct electricity without resistance. It generally occurs at low temperatures. Qubits of certain types, particularly electron-based qubits, are cooled to very low temperatures to enable this effect, either at the qubit level for superconducting qubits, or for devices and cables for reading the state of the qubits. Note the false friend: in English, we don't say "super", but "superconductivity".

Quantum supremacy: describes a situation where a quantum computer can perform a calculation that is inaccessible to the best supercomputers of the moment in a humanly reasonable time. The difference in computing time between quantum computing and classical computing must be several orders of magnitude or more than a human lifetime. Supremacy may or may not concern a useful calculation. For example, the quantum supremacy claimed by Google in October 2019 concerns an algorithm for generating and verifying random numbers that is of no practical interest. Another debate concerns the appropriateness of the term "supremacy", which echoes the controversial theme of "white supremacy". The term was coined by John Preskill in 2011.

SWAP: quantum gate that inverts the state of two qubits.

TDA: *Topological Data Analysis*, an automatic *clustering* method based on machine learning.

Enabling technology: "technology that enables another to function", technology that is essential for a quantum computer (e.g. cryogenics, cabling and control electronics).

TeraQops: Tera Quantum Operation per second.

Threshold: threshold above which, with a surface code, the logical qubit will have better fidelity than the physical qubits. In practice, physical qubits need to have an error rate ten times lower for the number of physical qubits per logical qubit to be reasonable, because at the *threshold* level, this number is infinite.

TLS: *Two-Level Systems* refers to a group of parasitic systems coupled to superconducting qubits. Their origin is not well known and is probably multiple (electrical fluctuators trapped at the substrate/air interfaces, substrate/ metal or metal/air, low-frequency magnetic fluctuators...).

Toffoli (gate): also called CCNOT, is a quantum gate operating on three qubits which modifies the value of the third qubit if that of the first two is 1.

Transduction: conversion of a signal of one type to a signal of another type. For example, to switch from the microwave regime (in GHz) to the optical photon regime.

Unitary transformation: reversible mathematical operation transforming the quantum state of a quantum system.

Transmon: type of superconducting qubit used in particular by IBM, Google and IQM.

TSV: *Through Silicon Vias*, vertical connectivity for integrated circuits that enables integrated circuit elements to be connected between their top and bottom layers. The technique can be used to extend this connectivity to chips made up of several integrated circuits superimposed. It is critical to the development of superconducting qubit chips.

TWPA: *Travelling-Wave Parametric Amplifier*. Amplifier operating at the quantum limit and at around 15 mK which amplifies the microwave pulse reflected by a superconducting qubit when it is read.

VQE: *Variational Quantum Eigensolver*: hybrid quantum algorithm used in chemical simulation, created in 2013. Its main contributor is Alan Aspuru-Guzik, a researcher who was part of the start-up Zapata Computing.

Wafers: semiconductor wafers (silicon or other) on which electronic chips are etched.

X: a single-qubit quantum gate that inverts its amplitude, going from  $|0\rangle$  to  $|1\rangle$  or from  $|1\rangle$  to  $|0\rangle$  in the base states.

Y: single-qubit quantum gate that rotates 180° around the axis Y in the Bloch sphere.

Z: single-qubit quantum gate that applies a sign change to the  $\beta$  component of the qubit vector, i.e. a phase inversion and a 180° rotation with respect to the Z axis.

## B. Members of the working group

- Frédéric BARBARESCO (Thales)
- Thierry BONHOMME (Académie des technologies, co-chair)
- Boris BOURDONCLE (Quandela, General Secretary)
- Philippe DULUC (Eviden)
- Marko ERMAN (Thales and Académie des technologies)
- Olivier EZRATTY (Epita and Quantum Energy Initiative)
- Philippe GRANGIER (CNRS Charles Fabry Laboratory)
- Daniel KAPLAN (French Academy of Technology)
- Catherine LAMBERT (Cerfacs, Chair)
- Jean-Claude LEHMANN (Académie des technologies)
- Anthony LEVERRIER (Inria Paris)
- Frédéric MAGNIEZ (CNRS IRIF)
- Mazyar MIRRAHIMI (Inria Paris)
- Jean-Philippe NOMINÉ (CEA DAM)
- Sophie PROUST (Inria, Digital Programme Agency and Académie des technologies)
- Gérard ROUCAIROL (Académie des technologies, co-chair)
- Claude WEISBUCH (CNRS, École polytechnique and University of Paris) of California in Santa Barbara (USA), Academy of Technologies)
- Lydia YAHIA CHERIF (Académie des technologies, referent)

N.B.: Academic ethics require t h a t any contributor to the drafting of a report brings to the group only his or her expertise, and refrains from promoting any personal or corporate interests. In the interests of full transparency, it should be pointed out that some members of the group, b y virtue of their professional position, are particularly committed to this discipline, in particular Boris BOURDONCLE, Frédéric BARBARESCO, Philippe DULUC and Marko ERMAN.

Additional contributions to the report from: Philippe CAMPAGNE-IBARCQ (Inria), Tom DARRAS (Welinq), Eleni DIAMANTI (CNRS), Julien LAURAT (Sorbonne University), Jean LAUTIER-GAUD (Welinq).

## C. Persons interviewed

- Cyril ALLOUCHE (Eviden)
- Frédéric BARBARESCO (Thales)
- Antoine BROWAEYS (Pasqal)
- Henri CALANDRA (TotalEnergies)
- Alain CHAMPENOIS (Quobly)
- Tom DARRAS (Welinq)
- Étienne DECOSSIN (EDF)
- Pierre DESJARDINS (C12)
- Matthieu DESJARDINS (C12 Quantum Electronics)
- Daniel ESTÈVE (CEA)
- Jérémie GUILLAUD (Alice&Bob)
- Félix GIVOIS (Genci)
- Anne-Lise GUILMIN (Eviden)
- Loïc HENRIET (Pasqal)
- Olivier HESS (Eviden)
- Iordanis KERENIDIS (CNRS)
- Anthony LEVERRIER (Inria)
- Christophe LABREUCHE (Thales)
- Stéphane LOUISE (CEA)
- Shane MANSFIELD (Quandela)

- Sabine MEHR (Genci)
- Tristan MEUNIER (CNRS)
- Grégoire MISGUICH (CEA)
- Damien NICOLAZIC (Eviden)
- Cécile PERRAULT (Alice&Bob)
- Jean-Philip PIQUEMAL (Sorbonne University)
- Pierre ROUCHON (PSL University)
- Félicien SCHOPFER (CNRS)
- Pascale SENELLART (Quandela)
- Emmanuelle VERGNAUD (Teratec)
- Daniel VERWAERDE (Teratec)
- Maud VINET (Quobly)
- Xavier WAINTAL (CEA)

This report was submitted for comment to the following people: Antoine BROWAEYS, Pierre DESJARDINS, Matthieu DESJARDINS, Jérémie GUILLAUD, Loïc HENRIET, Shane MANSFIELD, Tristan MEUNIER, Pascale SENELLART, Maud VINET.

This report reviews the construction and potential use of FTQC (*Fault Tolerant Quantum Computing*) computers to reliably perform complex calculations by overcoming the problems posed by the errors and noise inherent in quantum systems.

After recalling the reality of the quantum advantage and its needs, the report describes the use of error-correcting codes in the design of FTQCi computers. It then reports on the progress of the five most advanced physical technologies in the world for building such computers and the obstacles they will have to face in order to achieve the transition to scale necessary for the execution of useful applications. Finally, it discusses the technical and economic environment for quantum computers, how their performance can be compared and evaluated, and their future coexistence with other computing technologies (3D silicon, AI) or with supercomputers.

Académie des technologies Le Ponant - Building A 19, rue Leblanc 75015 PARIS +33(0)1 53 85 44 44 secretariat@academie-technologies.fr www.academie-technologies.fr

ISBN: 979-10-97579-63-0