# Tangling with the Entangled – The impacts of quantum computing on safety

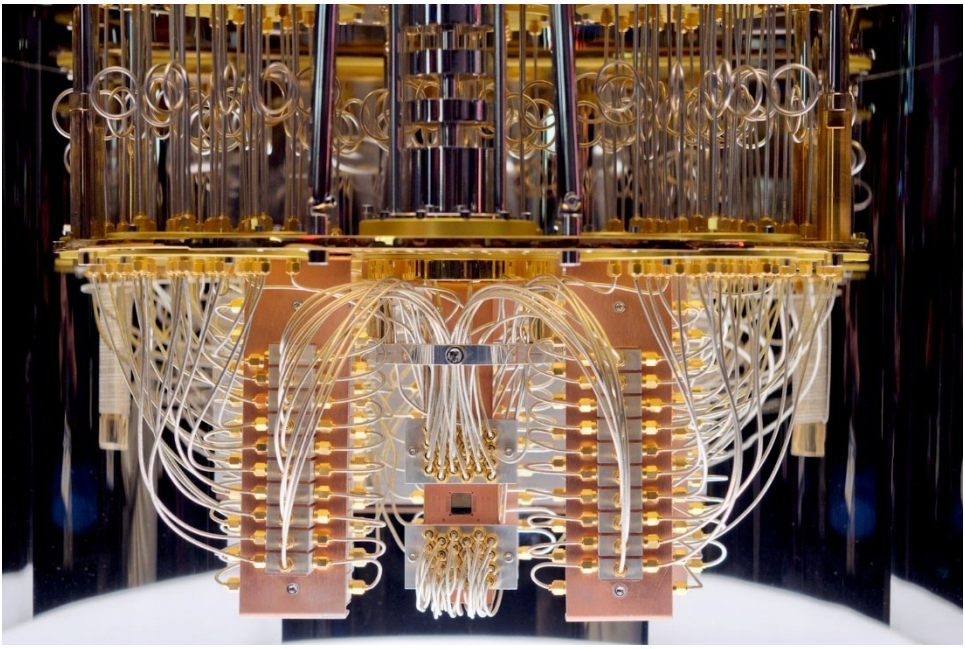

**Quantum computers are a new way of processing information that promise to solve a host of problems facing humanity, but they also carry risks. James Cruise and Bob Oates discuss the quantum future of safety.**

The safety community has long known that when new technologies emerge, they bring both new risks and new opportunities. On the horizon is quantum computing: complex, difficult to understand, but incredibly powerful.

While widespread access to this technology is likely to be at least a decade away, the rapid advances in the size and quality of quantum computers mean that now is the time to be preparing for how this new technology will interact with safety.

But if we are to prepare, we must first understand what the impact of quantum computers will be, in the world of safety. What is quantum computing? What opportunities do quantum computers present to enhance safety? What new risks do quantum computers introduce? Most importantly, what challenges must be addressed before quantum computers can be trusted to make safety critical design decisions?

We'll address these questions in turn, providing a high-level introduction to the technology, and going on to present the good (safety enhancements), the bad (risks that they undermine safety), and the ugly (problems that have not yet been well characterised.)

## Quantum Computing

The proponents of quantum computing view the emergence of quantum computers as a revolutionary step for society akin to the way steam power drove the industrial revolution and our ability to manipulate electrons drove the computing revolution. The ability to control and interact with the quantum mechanical world promises an alternative computing paradigm that natively processes information by representing data and transformations using structures that are more analogous to mathematical concepts such as matrices and linear algebra than the Boolean algebra used by classical computers. This new way of representing and reasoning about information makes some problems, thought to be intractable for even the most powerful supercomputers, solvable for the first time in human history.

A common challenge for quantum computing is that in this early stage of development, many of the discussions about its potential are driven by quantum physicists, who have their own terminology and concepts, such as the Bloch sphere pictured as a representation of a quantum bit (or "qubit.") But just as modern programmers do not need a working knowledge of the semiconductor junctions that enable transistors, it is useful for many to think about what these devices can do, rather than focus on how they do it.

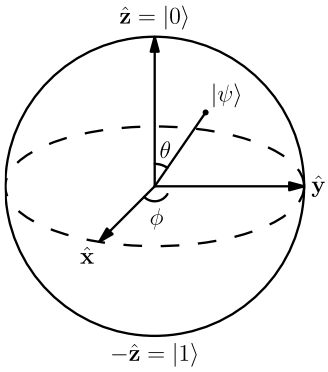


There are a wide variety of fantastic references for interested readers to look at if you want to explore the "how" of quantum computers. We suggest Nielson and Chuang's, "Quantum Computation and Quantum Information" [1] as a good starting point.

For our discussion it is sufficient to know that the output of a quantum computer will still be a sequence of numbers, the output of a chain of operations. As for classical computers, this output will just be a binary string. However, unlike classical computing while carrying out the computation, quantum computing considers all binary strings by associating an 'amplitude' with each one, this describes the probability for each string that this will be returned as the output string. Computation is then carried out by manipulating these amplitudes.

For example, consider a computation on 3 qubits. There are 8 possible binary strings that they could represent. For argument's sake we will assume that each string is equally likely at the beginning (not necessarily true in all cases!)

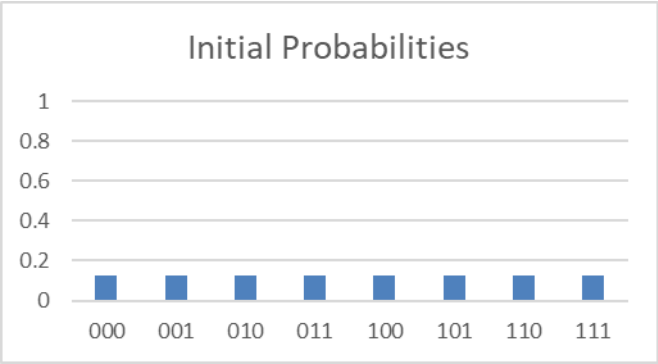


When a quantum computer processes those bits, it uses gates to manipulate those probabilities.

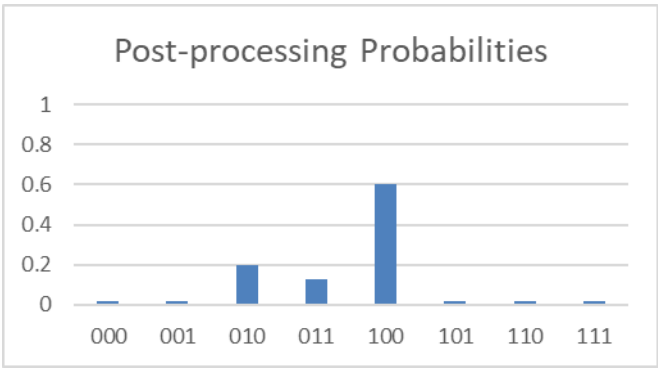


At the end of the calculation we "measure" the qubits which produces a random binary string using the assigned distribution. In our example we are more likely to get the answer "100" but "010" and "011" are also reasonably likely to occur, and all the strings have a chance.

This departure from traditional binary representation allows some computational problems to be solved more efficiently. In some cases, this means a quantum computer offers solutions that are polynomial faster and even exponentially faster for specific examples. Problems that have already been shown to benefit from this approach include:

- Finding the best entry in an unstructured list – Grover's algorithm
- Simulation of quantum mechanical systems – Hamiltonian evolution
- Factoring large numbers – Shor's algorithm

These abstract problems have huge real-world ramifications. Grover's algorithm can be applied to combinatorial optimisation or satisfiability problems ultimately offering more environmentally friendly logistics, efficiency savings for organisations, and better utilisation of resources. Hamiltonian evolution can be used to design new materials with beneficial properties for sectors such as aerospace, classical computing, and communications. The most explored application of Shor's algorithm is in the field of cryptanalysis (more on that later!)

A small number of subroutines have also been identified with the potential of creating other beneficial algorithms. These include "the quantum Fourier transform", and amplitude and phase estimation algorithms.

It's important to temper the enthusiasm for quantum computers with a dose of reality. They will not be useful for all problems, and they have a several inherent drawbacks that means that rather than replacing traditional computing, quantum computers will likely augment traditional computing, allowing us to take advantage of the best paradigm for the job on a case-by-case basis.

**"It's important to temper the enthusiasm for quantum computers with a dose of reality ... the primary challenges facing quantum computers can be summed up as them being error prone, slow, and complex."**

The primary challenges facing quantum computers can be summed up as them being error prone, slow, and complex.

Error-prone because unlike classical computing where the data representation using voltages has huge redundancy and hence natural error resilience, in qubit representation there is no such redundancy. This means any slight interference or misalignment can lead to an error that needs to be dealt with. Further, the probabilistic nature of quantum computers means that each algorithm really needs to be run several times before you can be confident in an answer. In short, if you want to add two numbers together, you're always going to be better off with a classical computer!

Slow because there are hard physical limits to the speed of quantum computers that are dictated by the physical properties of the systems used to implement them. For trapped-ion quantum computers this limit is the rotational speed of an ion, which fixes quantum computers to be in the region of 1000 times slower than a modern GPU. Now consider that you have to run error correction and programs multiple times on top of this, all compounding the slowdown in comparison to classical computers.

Complex because the laws of quantum physics introduce counter-intuitive properties into the system. For example, a property called "the no cloning theorem" means that you cannot just copy data from one place to another preventing branching calculations. Entanglement means that changing the copy also changes the original. This means that for computational tasks that rely on iterating around a single data set (for example, training AI) you are forced to reload the data into the computer every iteration and you cannot use checkpointing to protect your calculation.

These limitations ultimately prevent quantum computers from solving the data deluge that we're experiencing with modern computing. Quantum computers will be limited to high computation, low data applications, with classical, and maybe even newer paradigms, picking up the slack everywhere else.

## The Good – Enhancing safety with quantum computers

With these limitations in mind, we can explore the areas where quantum computers could potentially benefit the safety community. Here we present two scenarios, one rooted in a well-understood and researched application domain for quantum computers, and a second, in a more speculative space that is still being characterised by researchers.

Our first scenario builds on one of the most promising commercial uses for quantum computing: chemistry and materials science. Consider a world where chemical experiments can be analysed entirely virtually, without the need for human operators to go anywhere near volatile or harmful byproducts. Automated systems, capable of exploring a previously inconceivable range of possible permutations in order to develop exotic materials with properties that enable lighter, stronger, more conductive building blocks for the systems of tomorrow. As this technology becomes more common-place, bespoke materials, with exactly the right properties for their application can be specified, requested, and formulated to resist environmental wear, avoid the creation of stress fractures, or prevent catastrophic collapse.

The virtualisation of chemistry not only makes experiments more efficient, but also enables properties to be explored that can't be directly observed during live experiments, such as reaction rate. These insights will allow for the creation of high-fidelity digital twins that can pre-empt maintenance needs and schedule corrective actions for potential failures before they happen.

Our second scenario targets one of the most challenging problems in software safety today: software verification for complex systems. As software becomes more complex there is an explosion in the space of input parameters and possible outputs, that makes testing every scenario impossible. This is made more complex still when software systems interact, generating even more permutations. One approach to give confidence that software will act as intended is formal verification. Formal verification can be viewed as a logical satisfiability problem whereby a number of logical clauses are created, to represent undesirable states within the software. For example, one clause may represent the activation of a function that should not be reachable without certain safeguards being active (such as allowing a dangerous piece of machinery to operate whilst its protective inspection cage is open.) We can now search the software for input parameters that cause undesirable conditions. If the resulting clauses produce an unsatisfiable problem, then there exists no combination of events in the software's execution that can result in the undesirable condition. If a solution is found, then that solution represents a way for the software to behave in an unsafe manner. Currently, a key problem with this approach is that solving large-scale satisfiability problems is computationally expensive. However, quantum computers have the potential to tackle much larger satisfiability problems, and thus formally verify much more complex software.

**"Every discipline, safety included, is likely to be challenged by quantum computers to change the way they operate in response to revolutionary new possibilities."**

Of course, both these scenarios raise new questions about how we assure the output of our new quantum algorithms. How much testing needs to be done to gain confidence that our highly specified super-material doesn't have a new failure mode that wasn't tested for? How much tool validation is required of a quantum verification to provide the confidence to stop running traditional tests? The former question starts to blend the problems of software validation with the problems of mechanical engineering, whereas the latter can be viewed as an extreme case of tool qualification.

Every discipline, safety included, is likely to be challenged by quantum computers to change the way they operate in response to revolutionary new possibilities. In fact, quantum computing is in such an exciting period of discovery that the most important contributions to safety have probably not been found yet. We need to be ready to embrace those new opportunities as they arise. It is noteworthy that we have constrained ourselves in this article to only explore the benefits of quantum *computers*, not the wider benefits from the broader family of quantum technologies such as quantum sensors (like the quantum magnetometer pictured), which promise new opportunities for non-invasive medical monitoring, and industrial system monitoring.

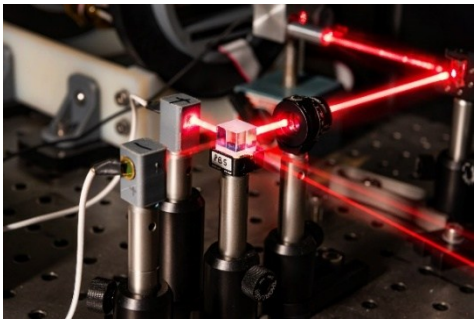

## The Bad – The risks to safety posed by quantum computers

The most talked about quantum computing risk is the threat that they pose to cryptography.

Modern cryptography relies on a family of mathematical problems being "computationally infeasible" to solve with traditional computers. Computational infeasibility implies that a traditional computer would take in the order of millions of years to find a solution. Sadly, several of these cryptographic mathematical problems could be solved in a matter of hours by quantum computers with the right specification. This raises the spectre that an adversary with access to a quantum computer could launch sophisticated cyber-attacks against systems whose safety and security relies on cryptographic techniques to prevent malicious harm.

The attacks that quantum computers could enable include reading confidential information in transit, sending fake messages/commands that look like they come from trusted devices or people, and signing malicious software so that it looks like it was produced by a trusted source. In terms of the SCSC's Data Safety Guidance [2] HAZOP guidewords: integrity, intended destination/usage, traceability, and fidelity can all be compromised by an attacker with a quantum computer, if vulnerable cryptography was the sole protection in place.

The complexity of quantum computing has made it difficult for many people to tell vulnerable algorithms and quantum resistant algorithms apart. Judging if an algorithm is vulnerable requires an understanding of the underlying mathematical problem that a cryptographic control relies on and, where appropriate, the size of the key that it uses. Looking at encryption algorithms specifically, an important variable to consider is if the algorithm of interest is asymmetric or symmetric. For asymmetric encryption, the key used to encrypt the message (the public key) is different to the one needed to decrypt the message (the private key.) For symmetric encryption the same key both encrypts and decrypts messages.

Quantum computers pose a much more pressing danger to asymmetric algorithms than symmetric algorithms. Asymmetric algorithms are vulnerable to quantum computers calculating the private key from the public key, as the mathematical relationship between the two is commonly a prime number factorisation or a discrete logarithm problem (both of which can be solved by variants of Shor's algorithm.) In comparison, quantum computers only offer a modest speed up for attacks against symmetric algorithms, by optimising the search process for identifying the key used (using Grover's algorithm.)

The limited scope of the attacks enabled by quantum computers offers little comfort for the millions of safety systems that rely on asymmetric encryption, including https connections for data transfer, digital signatures for validating firmware, and most remote network access software for allowing operators to interact with difficult-to-access systems.

In addition to the confusion about what algorithms are vulnerable, nobody knows precisely *when* quantum computers will have the power to target cryptographic algorithms. The specification of a quantum computer that is "cryptographically relevant" is not, as is commonly reported, simply a case of how many qubits the computer is built from. The connectivity between the qubits, their error rate, and the underlying materials that the qubits are built from, all have an effect. The specification required is also a function of the underlying mathematical problem. Predicting the future is always a dangerous thing to do in technology, but the general trend in the literature is that asymmetric encryption that relies on "the elliptic curve discrete logarithm problem" (for example TLS 1.3 which is used by modern https connections) is likely to be the first to be challenged by quantum computers. Followed by encryption that relies on the "prime number factorisation problem" (for example RSA, which is used by older https connections.) This is in part because of the smaller key sizes used by elliptic curve algorithms. The uncertainty around when algorithms will become vulnerable makes prioritising which systems to protect extremely difficult, especially for industries with long-lifetime assets (e.g. digital control systems in industries such as energy and aerospace) and sectors with long-term confidentiality goals such as healthcare and law enforcement.

But engineers should not be complacent about when to address these risks. Whilst most quantum-enabled attacks are several years away, nation states are already reportedly harvesting data encrypted by vulnerable algorithms, in the hope that it will still be useful when they acquire the power to read it. This means that from an intended destination/usage perspective, attacks have already started.

Lengthening the keys used, and increasing the frequency of key rotation may offer short-term protection against some quantum-enabled attacks, depending on context. Architectural changes may address other vulnerable systems, by no longer relying on asymmetric cryptography to perform key exchanges. But wholesale conversion of systems from asymmetric to symmetric cryptography simply isn't practical, as that transformation introduces new, often intractable, key management challenges. But the future is not entirely bleak. Significant efforts have been made around the world to identify new cryptographic techniques that are resistant to both traditional and quantum computers. Post-quantum cryptography has made great strides, with three specifications being officially approved by the US National Institute of Standards and Technology (1 key encapsulation mechanism and 2 digital signatures.)

Many large companies such as Microsoft, Amazon, and Google have already begun moving towards those algorithms under the hood, but for many organisations upgrading will require managed change, working with their software supply chain, and ensuring that they are ready for the emergence of quantum computers. As with any large-scale digital transformation there are going to be challenges along the way. The new algorithms have different computational and communication characteristics which have already exposed latent errors in networking infrastructure for early adopters. The extra bandwidth consumed by the new algorithms, and the much longer packets and signatures are going to have far-reaching consequences for digital infrastructure. In addition, despite the rigorous testing, there is no guarantee that an algorithm, or the implementation of that algorithm, will remain secure in the

face of future developments. In this regard, the fact that NIST has only approved a single key encapsulation mechanism so far means there is a single point of failure in the security for early adopters who have decided to follow their advice. Though it should be noted that there are several other algorithms that are already available but are still in various stages of being assessed.

## The Ugly – Challenges yet to be faced

We've made the argument that quantum computers bring both opportunities to make the world a safer place, and a more dangerous one. But there are several challenges that need to be addressed before quantum computers are ready to be applied. Here we present two of those challenges that the safety community may be able to offer insights into.

As we've frequently alluded to, quantum computers are inherently error prone. There is a very active research space looking for approaches to perform automated error correction on quantum computers. Large organisations such as Google have made progress [3] but the overheads from quantum error correction are large. Some people predict that we will need redundancy in the order of 100-1000 times, i.e. up to 1000 physical devices implementing qubits to realise a single, reliable logical qubit. If we wait for quantum computers to reach the sizes where this level of redundancy is possible it is quite possible that we will miss opportunities to make the world a safer place. Some researchers are now investigating what we can achieve with devices that have inherent error rates. There's potentially a lot to learn from both communications and safety science in this space. Communications specialists build reliable systems from components with well-characterised noise properties every day. Similarly, safety engineers build trustworthy systems from components that have known failure conditions.

Our second challenge is that of quantum verification. Classical software needs to be analysed and tested to make sure that it has been specified and implemented correctly. Likewise, quantum algorithms, represented by 'quantum circuits' also need verification and validation. The image below is an example of a quantum circuit representation for an error correction function. In this example, five physical qubits are used to represent the same "logical" qubit. Errors that form in individual physical qubits can be detected and corrected to make the overall system more robust.

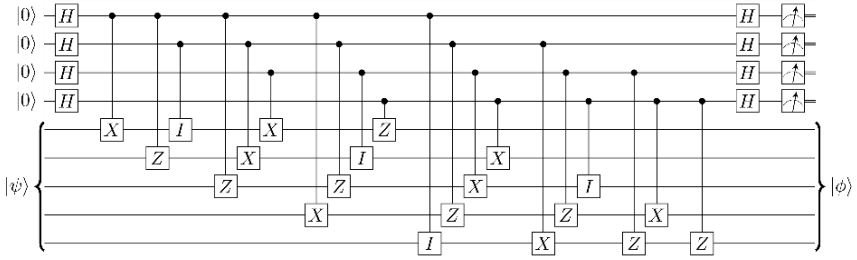


From a hardware perspective the current architecture for quantum computers makes this a much more challenging problem. Quantum gates, the implementation of operations on qubits, are realised by applying physical processes to qubits, for example irradiating a qubit with a microwave pulse. Therefore, each quantum gate is implemented separately for each qubit. In a classical computing architecture, we can be confident that if the adder within the ALU of a computer works for one pair of registers, that it will work for any other pairing. But

for quantum computers, there can be latent errors in quantum gates that will lie undetected until that specific hardware qubit has that specific operation performed on it.

From a software perspective there is little knowledge about how simple test cases can be used as building blocks to test the validity or otherwise of more complex systems, a problem that is not dissimilar to modular safety cases.

We suggest that there may be lessons to be learned from the way AI components will be ultimately integrated into safety systems. Design patterns that build trustworthy systems when one component is known to have a significant error rate but provides useful input nonetheless may be the key to integrating quantum computers into safety critical workflows.

## Conclusion / Call to arms

It's fair to say that quantum computers are an emerging technology. In fact, we believe that they're roughly where classical computing was in the 1940s. The uses of this technology in a few decades time are likely to surprise all of us!

Quantum computers are not a competitor to classical computers, but another device to be integrated by the dynamic and flexible silicon computer, alongside other quantum technologies such as sensors and timers.

We are likely to see impacts within the safety community as we seek to enhance safety with this exciting new technology and seek to protect systems from the harms that it makes possible. The safety community should be ready for a quantum computing revolution and be prepared to contribute their knowledge and expertise to this entangled web.

**Bob Oates, Cambridge Consultants**
Dr Bob Oates is a specialist in the cyber security of safety critical systems and the assurance of AI-enabled systems. He has worked for over ten years on a wide variety of problems in the defence, maritime, aerospace and energy sectors. He advises critical national infrastructure organisations on how to prepare for threat actors equipped with quantum computers.

**James Cruise, Cambridge Consultants**
Dr James Cruise is an expert in quantum computing and quantum algorithms. He is a mathematician by training but has worked in quantum computing for over ten years with range of organisations including government, startups, and consulting. He has particular interest and history in helping organisations understand the specifics and practicalities of how quantum computing will change their business including from a security perspective.

## References

[1] Michael A. Nielson and Isaac L. Chuang, "Quantum Computation and Quantum Information", Cambridge University Press, 2010, 2nd Edition, ISBN: 9781107002173

[2] Data Safety Guidance (version 3.7), SCSC-127J, Jan 2025, https://scsc.uk/scsc-127J

[3] Google Quantum AI and Collaborators, "Quantum error correction below the surface code threshold", Nature, 2024, DOI: https://doi.org/10.1038/s41586-024-08449-y