Terra Quantum

A New Path to Quantum Utility

Quantum Software Powered by Tensor Networks

Business Whitepaper


Tensor Networks are the early
gift of Quantum Computing
for unlocking utility today.

Table of Contents

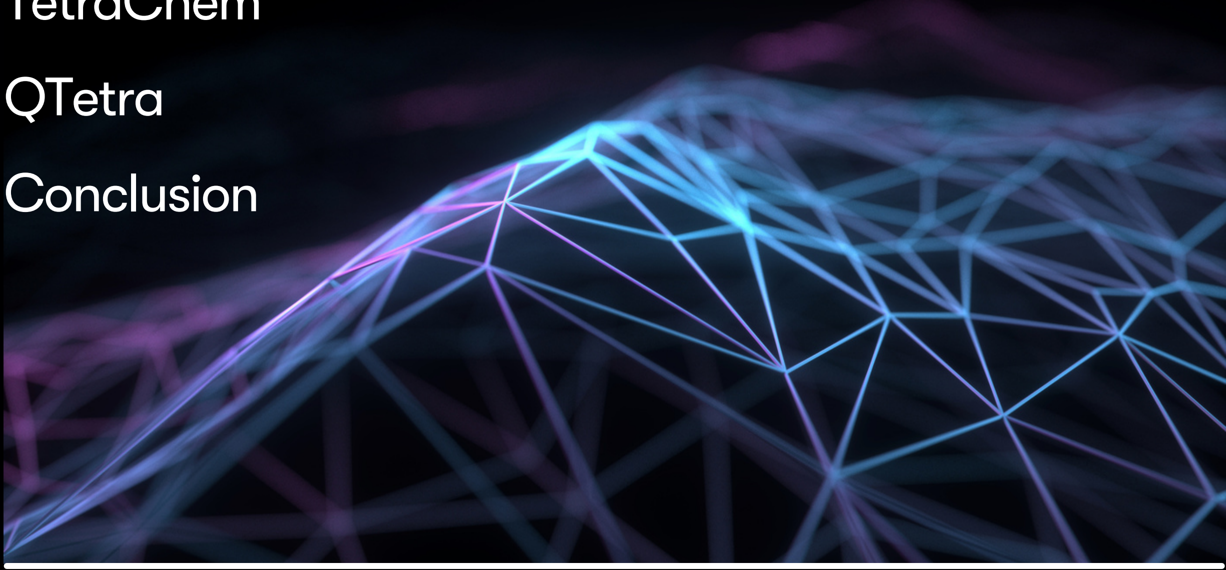

Executive Summary

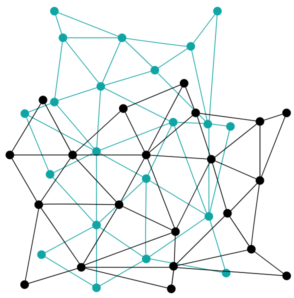

QUANTUM SOFTWARE POWERED BY TENSOR NETWORKS

It is well known that quantum physics is profoundly changing the future of high-performance computers. What is only appreciated by a far smaller group of experts is that it is also influencing today's software landscape.

Tensor network algorithms are sweeping through not only the field of many-body quantum physics where they originated, but also the quantum computing industry and, more broadly, all the business and technical fields that use tensors. This includes areas such as machine learning where important tools (e.g., TensorFlow and Tensor Processing Units) are named after tensors.

What is driving this adoption? Tensor network algorithms allow classical software to manipulate mathematical objects called tensors more efficiently than ever before, driving impressive performance gains.

At Terra Quantum, we offer our clients access to tensor network tools to add value to their businesses today. Yet what excites us most is our plan to implement these algorithms on quantum computers as well, enabling a hybrid quantum-classical approach to computation. There are strong, technical arguments that this will deliver awesome levels of performance.

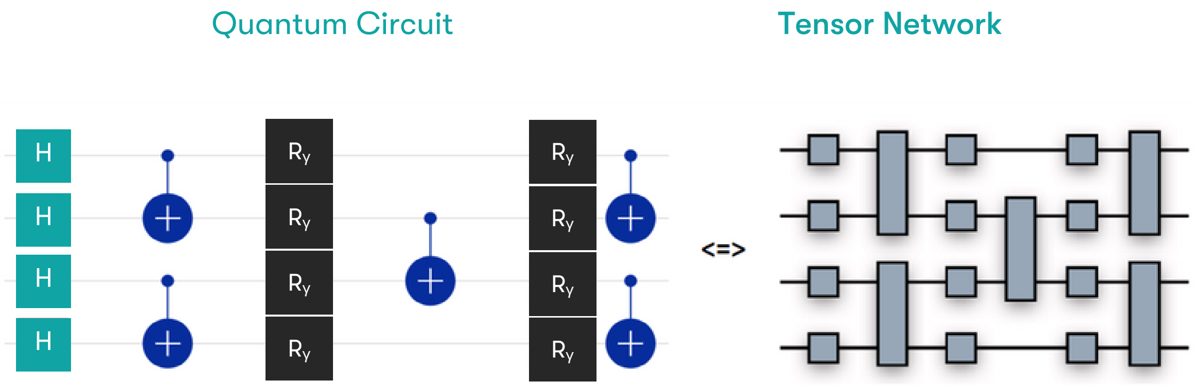

Fig. 1: Vizualisation of a Quantum Circuit simulated via a Tensor Network

The Gift of Quantum Computing

PERFORMANCE ENHANCEMENT

The classical computing industry has changed the world and consequently is worth trillions. Quantum computers are much less mature but if they fulfil their potential, they could be just as revolutionary. The moment when they first exceed the capabilities of a classical computer is dubbed "quantum advantage."

Despite the considerable head start of classical technology, the first claims of quantum advantage have already happened. In 2019, to worldwide acclaim, 53 qubits were subject to a total of over 1000 quantum logic gates before each qubit was measured. It was thought that a classical supercomputer would never be able to simulate this procedure.

Recall, however, that quantum computing hardware is not alone in its journey from the academic labs to the wider world.

Tensor network algorithms, also developed by quantum physicists, are making great strides. These algorithms, run on rather ordinary classical computers, were able to simulate the quantum advantage experiments.

Over the intervening years, quantum hardware has improved but tensor network techniques have grown ever more sophisticated. This has led to alternating claims and counterclaims as to whether quantum advantage has been achieved.

Regardless, tensor network algorithms are extremely powerful. They also have a wide variety of useful purposes. **Hence, the "gift" of quantum computing research may not actually be the hardware but the software powered by tensor networks.**

What are Tensors?

Since all forms of data can be represented as numbers, the job of a computer (at a certain level of abstraction) can be considered as storing and manipulating lots of numbers. Different arrangements of numbers go by different names: a line of numbers is a vector, a 2D grid of numbers is a matrix, and a 3D grid is a tensor.

To extend this idea, we need to change our thinking. Fortunately, new names aren't needed – they are all tensors from this point!

Imagine a so-called "human computer," a person hired to do numerical calculations prior to the invention of electronic computers. Given the task of handling various numbers in an orderly fashion, they could invent a system whereby indices label, say, the row, column, page, filing cabinet, etc., where the number is written down and stored. In this way, the indices do not indicate spatial dimensions but rather spatial information.

When we return to electronic computers, we can retain this perspective because ultimately, the indices indicate where the computer stores the numbers in its memory. Programmers, who are rarely concerned with memory allocation, focus on the idea that an appropriately long list of indices can refer to the stored numbers.

This level of abstraction may seem to leave us with no visual tools at our disposal. However, the representation of tensors shown in Fig. 2 fills that void. This is a tensor network diagram.

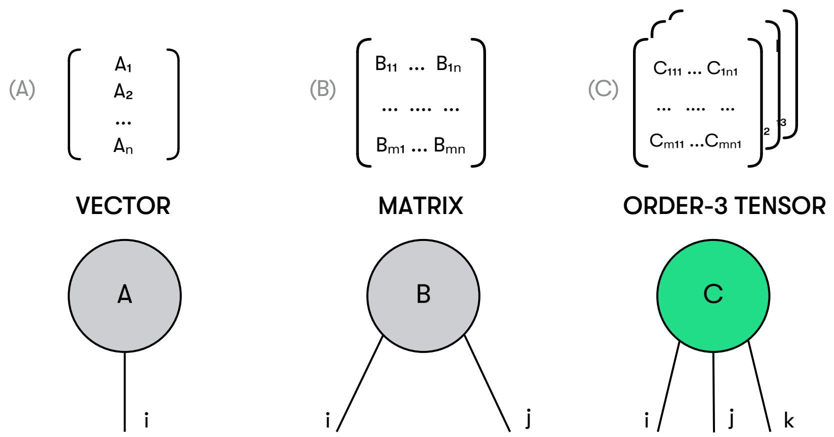

Fig. 2: Data shown as grids of numbers (top) and in the tensor network format (bottom)

Terra Quantum's Hybrid Approach

Quantum algorithms running on quantum computers constitute a unique technology. Despite this, they share enough similarities with tensor network algorithms that the two are often treated as rivals. At Terra Quantum, we view them as partners.

The logic behind this partnership is threefold:

- Firstly, tensor network algorithms running on ordinary classical computers make it easier to test out and develop quantum algorithms. No supercomputers are required.
- Secondly, tensor networks deliver value right now. We regularly productionize these algorithms for our clients by running them in QMware's high-performance cloud computing centers.
- Thirdly, they will only become more performant when implemented on, or in combination with, quantum computers.

As we explain in our technical white paper and our research papers, this is because when the so-called "rank" of a tensor increases, the algorithms get slower on classical computers. Quantum computers aren't affected by this.

To illustrate these similarities and differences, consider the following analogy. Imagine that people found the answers to their problems by jumping high in the air, plucking a delicate object - the answer - out of the sky before carefully bringing it back down to the ground with them as they land. The answers to the questions that involve more intricate correlations live higher up. These are the extremely valuable answers that everyone wants to find.

A tensor network performs every step of this problem-solving procedure very well – much better than alternative classical algorithms. Quantum computers, on the other hand, can jump much higher into the air. Unfortunately, they tend to lose their balance and have a hard time landing. This means they may not find the optimal answer to the problem and when they crash-land, only some fragments of their answer survive intact. [To picture it more clearly, imagine that your optimal answer in the sky is a great work of literature but the wreckage of the crash-landing consists of torn pages. The book's insights were in the connections (or correlations) between the sentences but all you can recover is how frequently various phrases appeared in the text.]

As classical hardware improves, tensor networks will be able to jump higher but still nowhere near as high as quantum computers. At the same time, the fidelity of quantum computers should improve. This will enable them to jump in a much more controlled and accurate way. Ingenious methods are being developed to prevent quantum computers from losing too much of the information contained in the answer when they land.

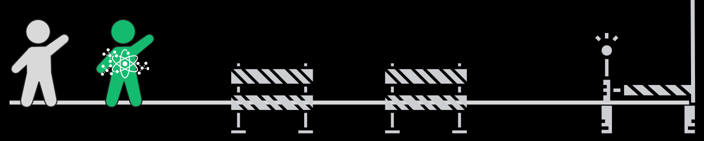

TetraBox

CREATING OUR ALGORITHMIC TOOLSET

TetraBox is the basis for all our applications. It is a library containing all the most important algorithms based on tensor trains – the most popular type of tensor network. Indeed, **Te**nsor **tra**ins are why we call all our products **Tetra**.

The algorithms include basic operations with tensor trains: TT-cross approximation, tensor train optimization (see TetraOpt), the AMEn algorithm and others.

The library is written in C++ with a Python interface, making it very simple to use. Due to the CUDA software layer, it has a fast and efficient implementation on both CPUs and GPUs.

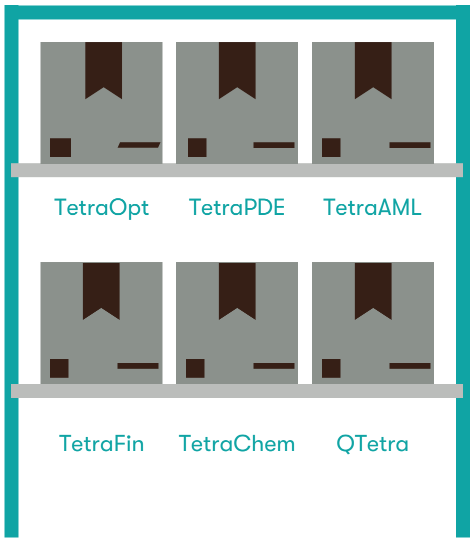

TetraOpt

OPTIMIZING BLACK-BOX FUNCTIONS

Optimization is the process of improving systems and operations to make them more efficient, productive, and profitable. Instead of attempting this through intuition alone, major companies are formulating these goals as mathematical optimization problems that are amenable to algorithmic solutions: banks perform collateral optimization; manufacturers optimize the designs of important components; and logistics companies are constantly solving complex routing and scheduling problems.

Black-box optimization is an especially challenging mathematical problem because you don't have an analytical formula, you don't have the derivatives, and you only have a small budget of function evaluations. They occur in many, and surprisingly different, scenarios.

Pharmaceutical companies employ black-box optimizers to assess the probability that a small drug will bind to its molecular target. At the same time, tech giants building massive neural networks in their data centers will perform black-box hyperparameter optimization.

TetraOpt is our exciting, new black-box optimizer. It comes with a long list of advantages: it is exponentially more efficient than simple approaches like grid search; it is universal (in the sense that it can tackle any optimization problem); it is simple and intuitive to use; it supports high levels of parallelization; and it is compatible with quantum implementations. TetraOpt has a shorter runtime than its rivals and can get to lower cost functions.

TetraPDE

SIMULATING WITH A SPEED-UP

The laws of physics, chemistry and biology are written in the form of partial differential equations (PDEs). These equations model the complex relationships between variables and the way they change from place to place and moment to moment. They are used throughout the manufacturing sector, in computer aided design (CAD) software, and in chemical engineering. Developing better methods to computationally solve PDEs could therefore lead to better fuel cells and biomedical devices, to name just two examples.

Analysing, designing, and optimizing industrial components is a problem with very large memory and runtime requirements.

This is because space is typically discretized into many small subdomains and the resulting equations are repeatedly solved at every timestep. Smaller subdomains and shorter timesteps lead to more reliable simulation results but the necessary resources can become prohibitive.

With TetraPDE, we can employ these Finite Element Methods on structured grids with exponential speed-ups over rival methods. This happens because the runtime of TetraPDE depends on the grid size in a fundamentally different way. TetraPDE is already designing new industrial chemical mixers by numerically solving difficult fluid flow PDEs known as the incompressible Navier-Stokes equations.

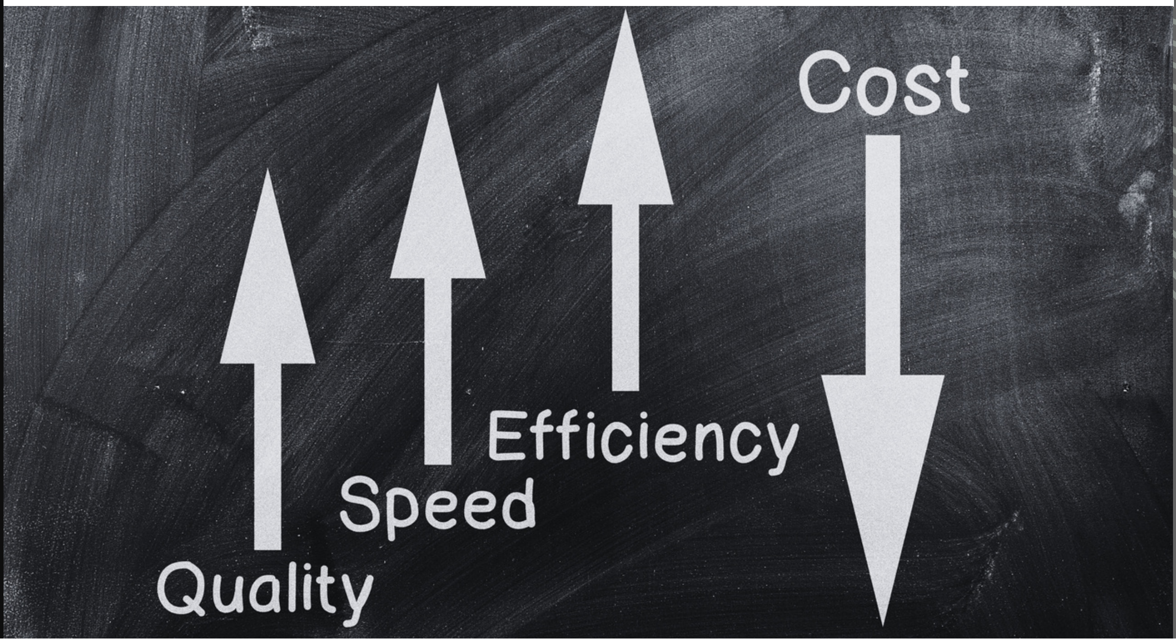

TetraAML

COMPRESSING AND ACCELERATING MODELS

The goal of artificial intelligence and machine learning is to build computers, that through experience and access to data, automatically improve their predictions and decisions. Neural networks are the dominant paradigm and have been described as the new electricity due to their ability to uncover highly complex relationships, solve unseen problem instances across industries, and create huge economic value.

The main challenge associated with the boom in neural networks is the massive (and growing) amount of compute they require.

The most reliable way to create a neural network model that will exceed the current state of the art is to make it larger and train it for longer. This constant one-upmanship needed to get to the top - and to stay there - is only becoming more exorbitant.

Rest assured - Terra Quantum has a solution. TetraAML is our toolbox to help you quickly and cheaply train neural networks and then fine-tune them to perform specific tasks for your business. It uses the best proprietary and open source methods to compress and accelerate your models, which facilitates their efficient deployment and enables real-time performance on any target hardware, whether in the cloud, on premises or even on edge devices (such as cameras and smartphones).

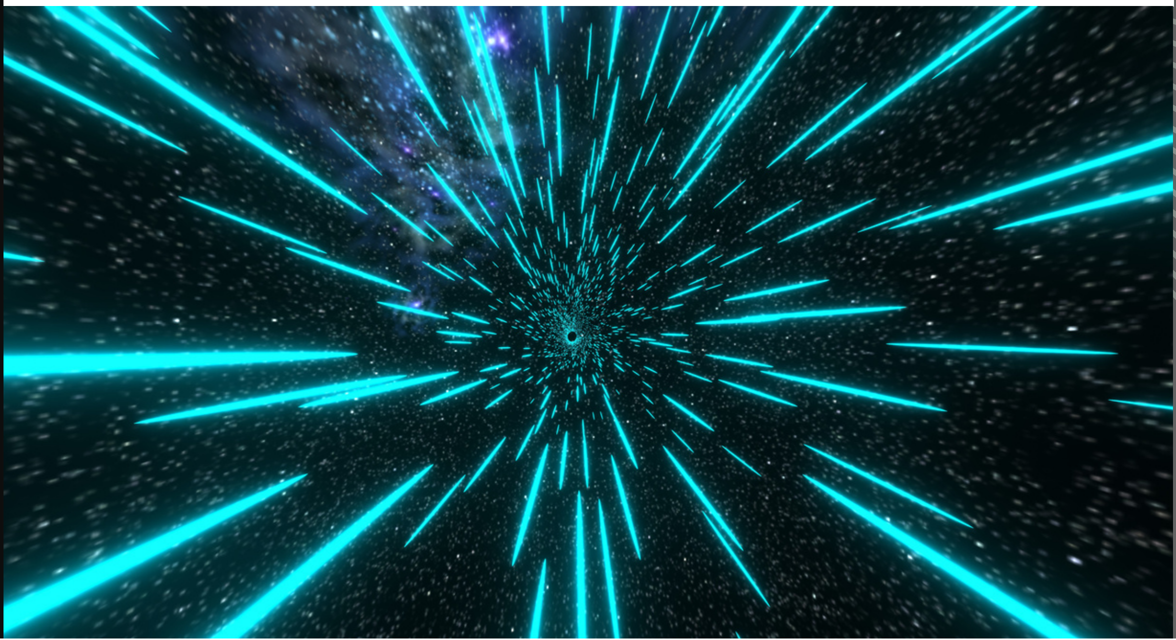

TetraFin

ENHANCING OPTIONS PRICING

How do traders and risk managers maximize returns in volatile markets? How can they separate themselves from the crowd – from their competition – while giving investors the confidence that their funds are safe?

Options are financial instruments that enable trading desks to implement complex strategies. Crucially, the options need to be priced accurately and frequently. This is vital for assessing the risks and identifying market imbalances.

Monte Carlo methods for option pricing are common because even though they are computationally expensive, they are effective.

The calculations are run many times per day, so even tiny improvements to the underlying methods can greatly improve the financial performance of the entire firm.

The approach we have developed in TetraFin involves calculating high dimensional integrals and solving partial differential equations. It demonstrates a much better complexity scaling $O(N^{-2})$ in comparison to classical and quantum Monte Carlo simulations, which scale as $O(N^{-1/2})$ and $O(N^{-1})$, respectively. Hence, this is not a tiny tweak to the existing approaches. It could be game-changing.

TetraChem

SIMULATING COMPLEX CHEMICAL SYSTEMS

The aim of quantum chemistry is to obtain detailed and useful information about the electronic structure of chemical systems. The fields of quantum chemistry and quantum computation are both progressing rapidly but despite this, practically relevant chemical systems cannot yet be studied with quantum computers.

Simulating chemically interesting quantum systems is so challenging that while there are clear incentives to use high performance computing environments for quantum chemistry calculations, it is debatable whether classical coprocessors, such as graphical processing units (GPUs), are being used efficiently.

Tensor network algorithms – running on ordinary, classical central processing units (CPUs) - are taking great strides in being able to straightforwardly handle the complexities of simulating many-body quantum systems. Based on this success and the fact that GPUs are so technologically mature and well-understood, we have made convincing predictions about how fast tensor network simulations will be able to run on GPUs.

As we are bringing these plans to fruition, we are also researching our next step, the move to quantum processors (or QPUs). Although quantum hardware is in its infancy, tensor network algorithms were developed by quantum physicists for quantum many-body physics, and it is becoming clear that they are naturally well-suited to run on a QPU. That is, they are not only quantum-inspired but also quantum-compatible.

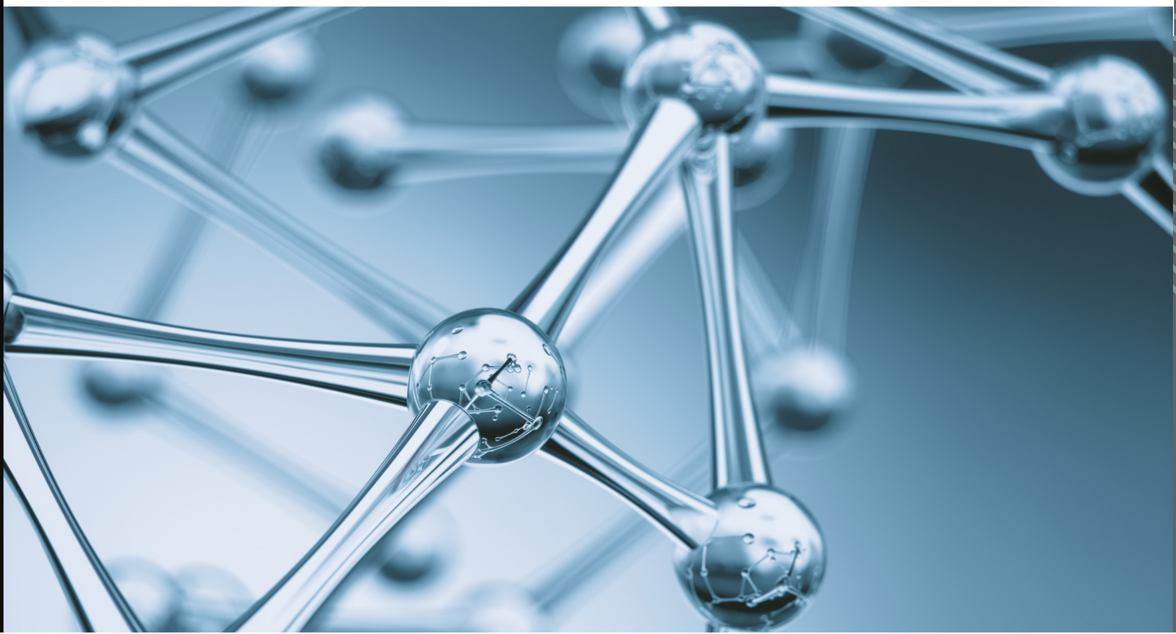

QTetra

REALIZING QUANTUM CIRCUITS ON QUANTUM COMPUTERS

Quantum algorithms consist of three main steps:

- Encoding the data (also known as quantum state preparation),
- Operating on the data vector (mainly with matrix-vector multiplications)
- Reading the results (also known as quantum state tomography).

Implementing these steps for arbitrary vectors and matrices requires exponentially large quantum circuits. Without an astute approach, this would make large-scale quantum computing unachievable!

Fortunately, tensors networks save the day. They represent data in a very similar way to quantum circuits. When we only consider vectors and matrices that are effectively represented as tensor networks, then all these steps can be performed efficiently: that is, with linear, not exponential, complexity in the number of qubits.

QTetra is where we do this. It is the quantum extension of our tensor train algorithms. All the basic sub-routines necessary for realizing quantum circuits on quantum computers are kept here.

Conclusion

TENSOR NETWORKS UNLOCK QUANTUM UTILITY TODAY

In this whitepaper, we have presented our vision for the business impacts of tensor networks, starting with the context and a broad overview before providing concrete descriptions of our main algorithm libraries. Combined with the other pillars of our business, extending across the full scope of quantum technologies – from high performance compute to unbreakable security – there will be tremendous added value.

To go further into the technical details, please see our associated research papers.

Tensor networks are the gift of quantum computing – an early gift that is here to unlock quantum utility today!

Terra Quantum AG
www.terraquantum.swiss
Kornhausstrasse 25
9000 St. Gallen, Switzerland

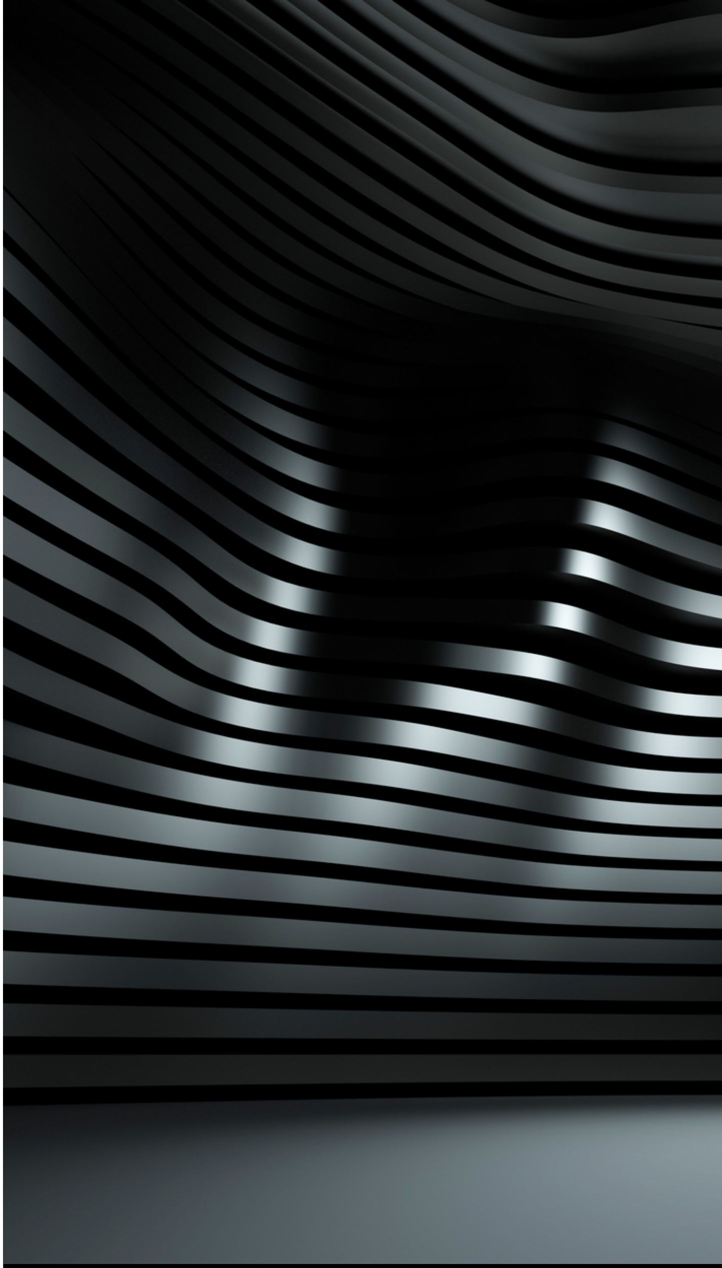

Contact

Vishal Shete
Managing Director UK &
Global Director of Commercialization
vs@terraquantum.swiss

Karan Pinto
Director of Growth
karan@terraquantum.swiss

Romi Sumaria
Head of Business Development
& Partnerships
romi@terraquantum.swiss

Terra Quantum

Take the leap towards your next level of performance. Tap into breakthrough algorithms and the power of quantum computing powered by tensor networks, unlocking crucial competitive advantage.

A New Path to Quantum Utility. A Business Whitepaper by Terra Quantum