Biannual Highlights – H1 2024

August 2024

Kipu Quantum: application- & hardware-specific algorithms for early industrial usefulness, in 18 to 36 months from now

Kipu Quantum – We stand at the forefront of quantum computing algorithms

We are specializing in application- and hardware-specific quantum software solutions that drastically reduce the requirements for solving industry use cases.

These solutions enable even the currently small and noisy quantum processors to yield significantly improved results, while approaching industrial usefulness.

Experience useful quantum computing now – With Kipu!

To solve significant industry problems with quantum computers, using more than 50-70 qubits in an algorithm is crucial. Despite having processors with over 100 noisy qubits, this remains out of reach. Current qubit specifications and gates are inadequate for greedy quantum algorithms.

However, we anticipate overcoming this limitation in the next 18-36 months with advancements in quantum computing paradigms and hardware. This is particularly promising for optimization problems in industry.

Kipu's Tech Edge – Our algorithmic compression achieves a massive reduction in the required circuit depth

We outperform competing state-of-the-art quantum algorithms through our digital, analog and digital-analog compression techniques, dramatically reducing the necessary circuit depth by orders of magnitude.

Our technology seamlessly integrates with leading hardware concepts such as superconducting circuits, ion traps, and neutral atoms.

Presenting our Quarterly Highlights!

Kipu Quantum constantly moves the needle on our path to making quantum computers useful. In the first half of 2024, we realized several breakthroughs, which we briefly discuss in this document.

Please reach out should you want to learn more!

Biannual highlights featured in here – the Kipu team reached several new records on our path towards early industrial usefulness

- Presented largest MIS optimization ever solved on a gate based quantum computer.
- IonQ leap-froged Quantinuum's superior fidelity processor, thanks to Kipu algorithms.
- New record in portfolio optimization achieved using non-hybrid methods.
- Slashed required parameters to train ML models for image processing by 1,000x.
- Improved digital algorithm makes large-scale optimization beyond 100 qubits possible on IBM hardware.
- A new hardware-specific quantum algorithm for neutral atoms solves MIS problems beyond 100 qubits, 200 x faster.
- Up to 1,000 x faster time-to-solution for BASF’s logistic problems was achieved on IBM and IonQ hardware.
- Reduction of circuit depth by 50 x allows accurate chemical modeling to be done on small & noisy quantum hardwar.

We present the largest MWIS optimization ever tackled on gate-based quantum hardware, using all 36 qubits of IonQ Forte

Experimental results for 36 qubit Maximum Weighted Independent Set on IonQ Forte with Kipu

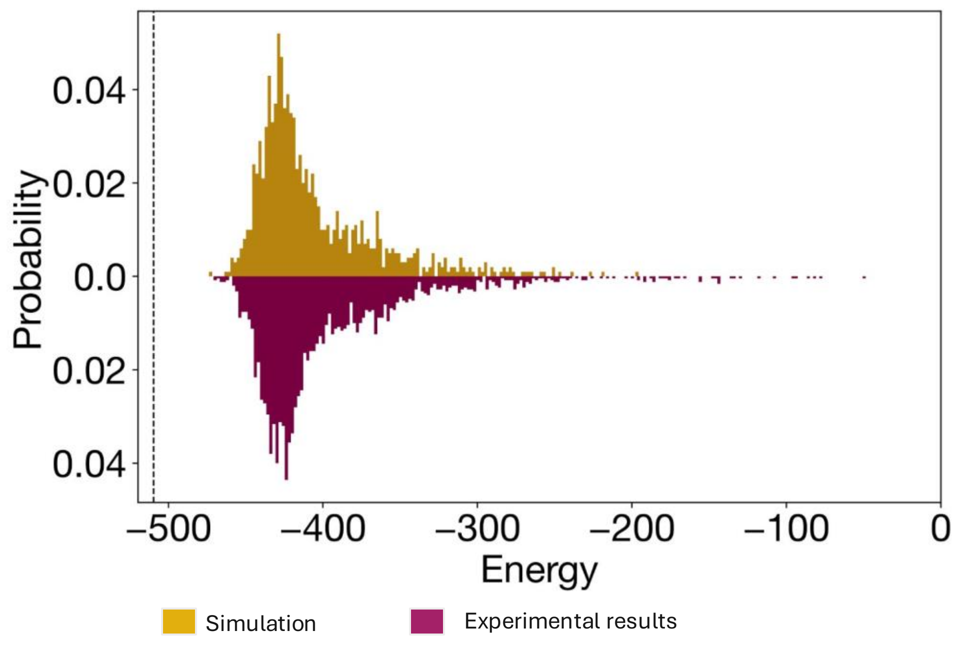

- Using many physical qubits in algorithms is challenging, because of the accumulation of noise. However, Kipu's algorithms are more robust against noise, because they require orders of magnitude fewer steps.
- This allowed us to use all 36 qubits of IonQ Forte, obtaining the largest maximal (weighted) independent set (MWIS) optimization tackled on a gate-based quantum computer.[2] MWIS can be applied to job shop scheduling, network optimization in energy and telco applications, as well as target interaction in drug discovery.
- While this problem size is no „quantum advantage" yet, we consider this an important milestone, because the algorithm will also allow to solve intractable problems on larger ion traps, i.e. beyond ~60 qubits[3].

1 We show the energy distribution of simulated vs experimental results. We consider a randomly generated weighted MIS problem with 36 qubits and experimentally showed how our circuits suit current quantum hardware, such that the results agree with the expected outcome. We achieved an approximation ratio of 0.92 with merely two Trotter steps.
2 Due to our compression, we only required 64 2-qubit gates without hybrid workflows like QAOA, using every single available qubit of IonQ Forte.
3 Further improvement expected by adding error mitigation & hybrid workflows.

Hardware performance is a key factor to getting good results with quantum, but Kipu's compression leads to decisive leaps

A-B test showing the impact of applying Kipu's digital compression on IonQ vs Quantinuum[1]

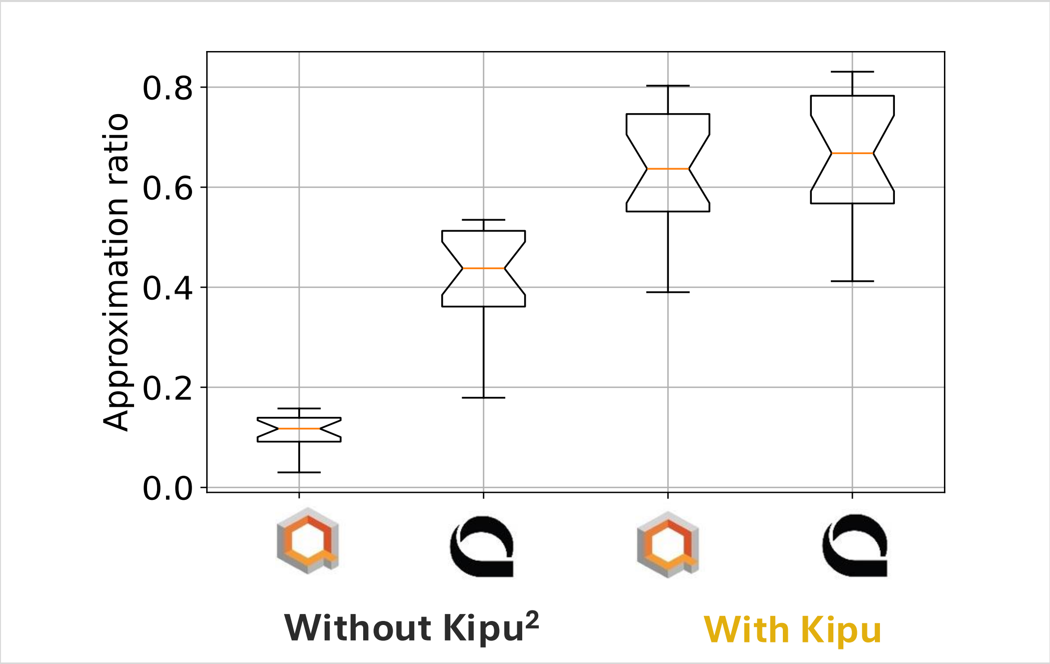

- IonQ and Quantinuum are the leading players in ion-trap hardware, in terms of qubit quality, Quantinuum has the superior machine, which typically translates into better results in direct benchmarks
- However, when combining Kipu's digital compression technology with IonQ's fidelities (99.26%, 36 qubits), IonQs hardware yields a ~6x increased approximation ratio from 0.10 to 0.65, outperforming what Quantinuum's higher fidelities (99.8%, 32 qubits) without Kipu compression can deliver with AR = 0.45.[3]
- Due to our compression, we need about 10x less gates and simultaneously improve the solution quality, allowing IonQ hardware to catch up with Quantinuum.

1 1 We built a simple noisy trapped ion emulator (all-to-all connectivity) using the specified average one- and two-qubit gate fidelities of IonQ Forte and Quantinuum H2. This model is not intended to fully model hardware. However it allows a fair comparison of running algorithms with different gate fidelities. We ran 20 different portfolio instances with 13 assets/qubits each.
2 Reference was digitized quantum annealing (DQA).
3 As measured using the approximation ratio, i.e. how close the solution is to the global optimium (1 = 100% as a target).

Kipu technologies can also be applied to portfolio optimization - we set a new record and beat QAOA by 100x in time-to-solution

Experimental results using non-hybrid Kipu algorithm vs. simulated results with a hybrid QAOA[1]

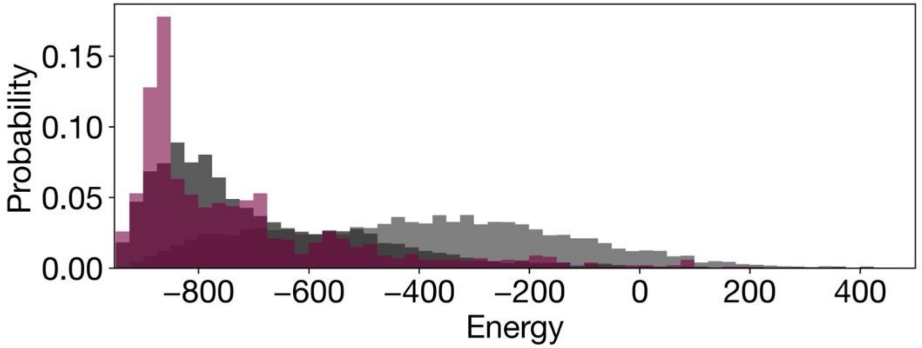

Kipu algorithm – On hardware requires 122 2-qubit gates, approximation ratio (AR) of 0.75.

Simulated QAOA with one layer – Requires 190 2-qubit gates for AR of 0.42.

Simulated QAOA with 5 layers – Requires 950 2-qubit gates for AR of 0.72 – *cannot be implemented though because of the large number of gates*

- To reduce the impact of hardware noise, hybrid instead of pure quantum algorithms are often used. This comes at a price – Iterating between classical and quantum takes long and is expensive, and the scalability beyond 20+ qubits is a concern. Often, they are also too greedy and cannot be implemented.
- A pure quantum algorithms by Kipu outperforms hybrids like QAOA – We find a better solutions, about 100x faster and are compressed enough to allow implementation on noisy hardware
- This is one of the largest and highest quality demonstrations of portfolio optimization done on a quantum computer[2], and to our knowledge the largest pure quantum implementation for this problem[3]

1 We show the energy distribution of experimental results of Kipu's pure quantum algorithm done on IonQ Forte, vs. simulated results for hybrid methods QAOA with 1 and 5 layers, respectively.
2 We tested 20 asset portfolio optimization instance from real data.
3 For the largest hybrid (QAOA) implementation refer to this work by J.P. Morgans quantum team: https://www.nature.com/articles/s41534-023-00787-5

Digital-analog quantum compression improves ML models for image classification, better prediction with 1000x fewer parameters

Quantum-based machine learning models vs. larger classical CNNs for identification of breast cancer

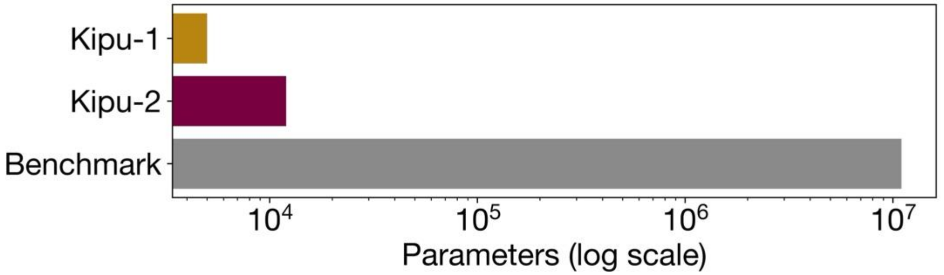

Kipu-1 – CNN using 5,000 parameters, with 4 qubits

Kipu-2 – CNN using 12,000 parameters, with 9 qubits

Benchmark – CNN using 11 million parameters, no quantum[1]

arxiv [Click here for the paper!]

- A critical pain point in AI/ML, especially in convolutional neural networks (CNNs) is the need for a large number of (expensive) parameters to train models, e.g., in image classification.
- Bringing quantum computing based on digital-analog compression into the workflow reduces the number of parameter by 1,000x at an improved performance.
- The resulting CNNs outperforms classical CNN on 4 and 9 qubits – the performance KPIs AUC2 of 91.8 and 92.6% respectively exceeds the 90.1 and 91.9% obtained with ResNEt-18 and Google AutoML. The same trend is found for the accuracy ACC with 87.2% and 89.1% vs 86.3% and 86.1% respectively.
- Even higher quality images can be achieved and require ~150 qubits on neutral atoms.[3]

1 1 Tested against ResNet-18 & Google AutoML.
2 AUC = area under the receiver operating characteristic curve; ACC = Accuracy. Test case was breast cancer; method also applicable for other diseases such as pneumonia, as well as image based data sets in finance (fraud detection), digital farming, object identification in AGVs.
3 Validation on 4-9 noiseless qubits was done using simulated neutral atom devices.

Improved digital algorithm makes large-scale optimization beyond 100 qubits possible on IBM hardware

We proved that Kipu can solve combinatorial optimization beyond 100 qubits on IBM

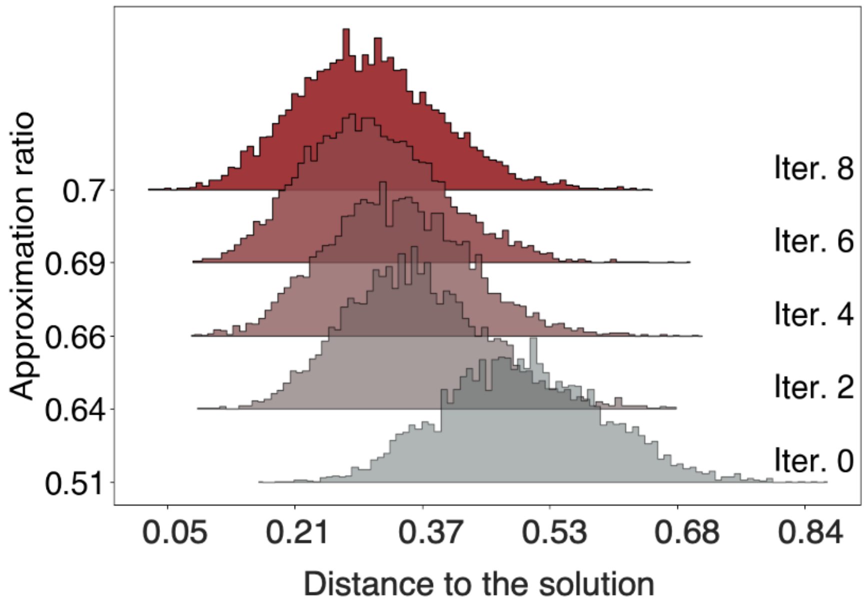

arXiv [Click here for the paper!]

- Our quantum algorithms outperform non-scalable ones like QAOA. We achieve 200 x higher success rate, better answers and scaling to larger problems.[1]
- We used 100 qubits of IBM Brisbane to solve a spin glass problem, which maps to industry problems like scheduling. This is relevant for logistics and manufacturing.
- This is a key milestone on our path to quantum advantage – solving 100+ qubit problems is a prerequisite. Based on IBMs roadmaps for 2024 and 2025, we expect to solve an intractable problem soon, by boosting classical heuristics and hybrid solvers!

1 Apples-for-apples comparison done on 20 qubits; approximation ratio or AR as a metric for how close we are to a global optimum is improved by 30%

Brand-new analog algorithm opens a path to near-term usefulness on neutral atom hardware for graph problems

Fast, high quality, large scale optimization on neutral atoms is possible with Kipu

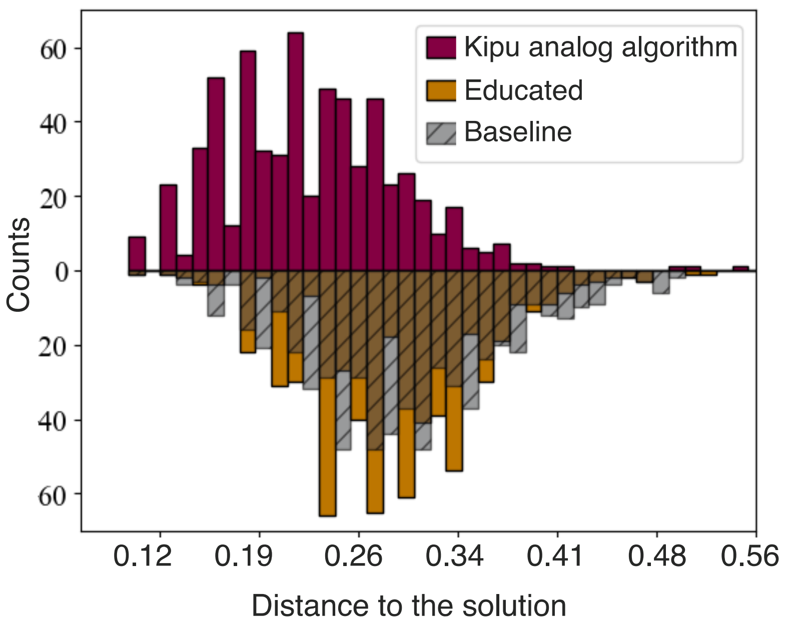

Click here for the paper.

Click here for a practical example with customer MásOrange

- Neutral atom hardware can solve use cases that can be mapped against a MIS[2], such as network optimization.
- On a 100+ qubit problem, we proved that our analog algorithm improves what neutral atoms can do:
  - Accelerate to 200 x time-to-solution – getting an answer fast is crucial for practical applications
  - Get 20 % better answer[3]
  - No classical optimization needed[4] – this assures a path to run 500+ qubits.
- Based on QuEra's and Pasqal's roadmaps beyond 2024, we expect to see neutral atom QPUs with hundreds of qubits, this will approach quantum advantage when combined with our algorithms!

1 Figure shows the energy distribution from experiments performed on QuEra's cloud accessed Aquila device solving MIS on a 100-qubit graph for T = 1µs; 2 MIS Maximum independent set; 3 measured in approximation ratio or AR, a metric for how close we are to a global optimum; 4 such as costly hybrid methods or Bayesian optimization protocols

With BASF, Kipu algorithms gave 1,000 x improvement & allowed solving non-implementable problems on real hardwarde

Drastic boost by combining Kipu algorithms vs. existing problem decomposition methods

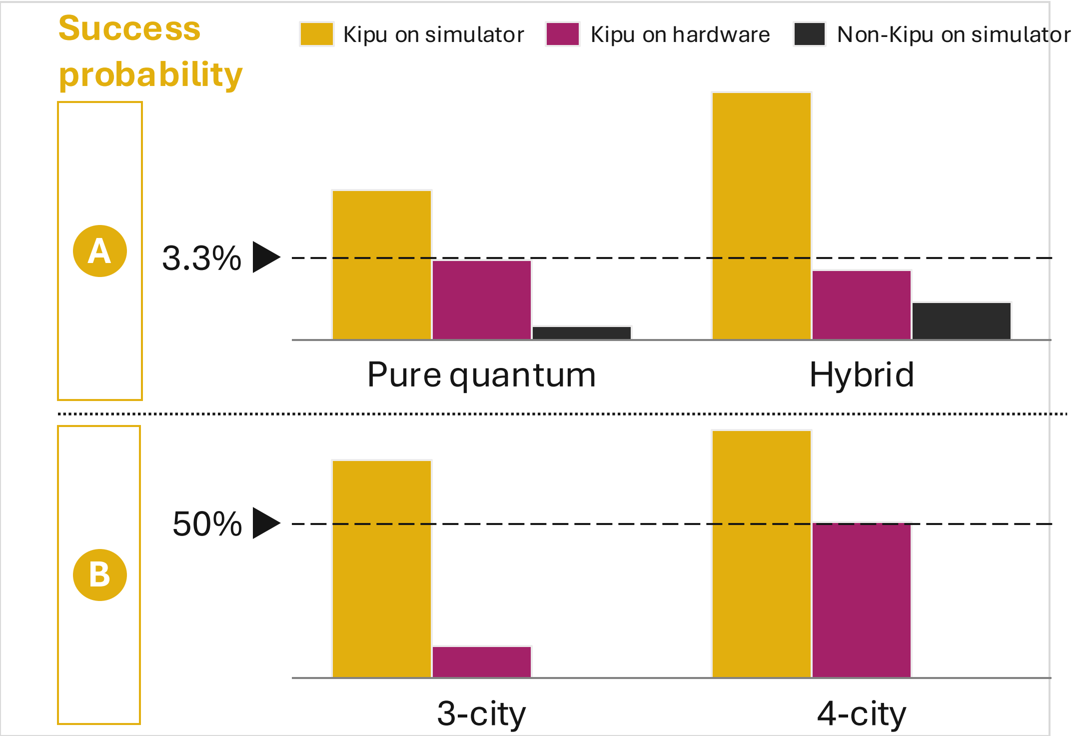

arXiv [Click here for the paper.]

- Many challenges in manufacturing and logistics companies can be mapped on traveling salesman problem (TSP) or job-shop scheduling problem (JSSP).
- At scale, these problems are hard, even for super computers. However, quantum often fail due to short coherence times of the QPUs or give suboptimal solutions.
- With our customer BASF, we recently demonstrated:
  - Up to 1,000 x higher success probability compared to standard algorithms such as QAOA.
  - We ran problems on real hardware that were previously not implementable within the limits of today's hardware.
- A near-term path to quantum advantage lies in quantum solutions that combine our algorithms & near-term hardware to boost classical heuristics and hybrid solvers.

(A) We solve a JSSP subproblem. We show that Kipu outperforms QAOA in terms of success probability; (B) 3-city TSP instances are solved on IBM and IonQ

Kipu brings chemical modeling on near-term quantum computers, by reducing required circuit depth & time by 50 x

With Kipu, modelling chemical reactions may be realized on quantum computers by 2025

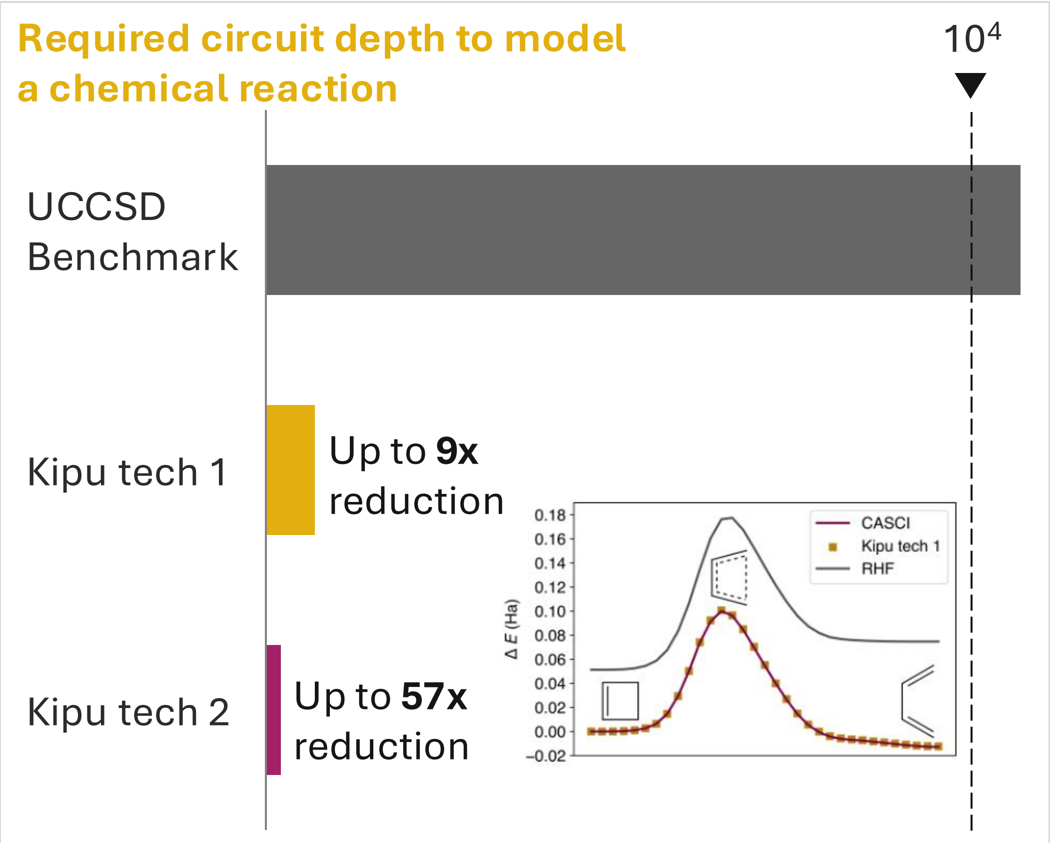

- Predicting the properties of molecules & chemical reactions[1] is considered a very hard application for quantum computers.
- Even on fault tolerant hardware, current algorithms will require huge circuit depths, and consequently would take weeks, i.e., too long for real-time applications.
- We show for a use case[2] that Kipu's algorithms require less quantum resources than UCCSD while achieving:
  - Chemical accuracy in comparison to exact methods
  - Up to a 50-fold reduction in the number of CX gates and up to 57-fold reduction in circuit depth
- Implications are significant – We reduce the required number of gates to an extent that NISQ hardware can deal with as soon as 2025, and provide a way to shorten the time-to-solution by a factor of 50 as well!

1 Underlying use case is atomistic modelling / electronic structure modelling; 2 Geometries as used in Nature Physics 19, 1787 (2023), link

Application- and hardware-specific algorithms for early industrial usefulness

Ways to engage with the Kipu Quantum team

- Test our algorithm libraries
- Work with us on a showcase
- Develop a concrete pilot with us
- Embark on an 18 – 36 months journey to demonstrate first usefulness

Join us on our journey towards useful quantum computing!