

Marc Drudis,¹ Alberto Baiardi,¹ Mattia Chiurco,^{1,2} Francesco Tacchino,¹ Stefan Woerner,¹ and Christa Zoufal¹

¹IBM Quantum, IBM Research Europe — Zurich, Switzerland ²ETH Zurich, Switzerland

1 Introduction

Preparing structured quantum states is a core challenge in quantum computing, with applications ranging from ground-state estimation to quantum simulation. A particularly natural strategy is to mimic physical processes that generate such states, e.g., through real- or imaginary-time evolution under a target Hamiltonian. While Trotterized time evolution [1, 2, 3] provides a direct implementation of this idea, it typically requires deep quantum circuits to achieve sufficient accuracy. Variational quantum time evolution (VarQTE) [4, 5] offers a promising alternative by approximating dynamics within a parametrized ansatz using McLachlan’s variational principle [6], thereby reducing circuit depth. However, the practical deployment of VarQTE is limited by the cost of evaluating the quantum geometric tensor (QGT) and gradient terms, as well as by numerical instabilities arising from noisy estimates. Understanding how to efficiently compute these quantities is therefore central to making VarQTE viable.

In this work, we introduce **bowtie VarQTE**¹, a framework that reduces the gradient and QGT computation costs by exploiting the causal structure of quantum circuits. First, we introduce a method to compute the Quantum Geometric Tensor (QGT) of a parameterized quantum circuit at system sizes that lie beyond the reach of existing approaches. We do so by decomposing the problem into the independent classical simulation of smaller, exactly simulable sub-circuits, that exploit causal light-cones [8]. Supporting VarQTE dynamics with classical simulation techniques (where possible) focuses quantum resources on classically hard operations such as sampling, while ensuring numerical stability throughout the evolution. We estimate that, for the system sizes considered here, alternative approaches to approximating the QGT are either prohibitively slow or yield insufficiently accurate results. We benchmark our approach against Approximate Quantum Compiling (AQC) [9, 10, 11], a widely used state-preparation method that employs classical tensor-network techniques [12] to compile Trotter-based quantum circuits into short-depth approximations. Our method achieves similar state-preparation fidelities by directly following McLachlan’s variational principle [6], without requiring a classical representation of the target state (which can be a costly constraint in AQC). Finally, we explore the use of VarQTE, for both real-time (VarQRTE) and imaginary-time (VarQITE), in the context of Sampled-based Krylov Quantum Diagonalization (SKQD), for a 2D Heisenberg Hamiltonian of 63 qubits.

2 Methods

We consider an ansatz $U(\boldsymbol{\theta})$ and Hamiltonian $H = \sum_i c_i h_i$ that can both be decomposed into local terms. In VarQTE, the dynamics of the state are mapped to dynamics of the parameters by enforcing McLachlan’s variational principle [6]: $g\dot{\boldsymbol{\theta}} = \mathbf{b}$. The two relevant quantities are the gradient entries $b_i = \sum_j c_j \langle \partial_i \psi | h_j | \psi \rangle - \langle \partial_i \psi | \psi \rangle \langle \psi | h_j | \psi \rangle$ and the elements of the QGT $g_{i,j} = \text{Re}\{\langle \partial_i \psi | \partial_j \psi \rangle - \langle \partial_i \psi | \psi \rangle \langle \psi | \partial_j \psi \rangle\}$. **The central insight** of our method is that each parameter in the circuit and each term in the Hamiltonian induces an auxiliary state of the form $WPW^\dagger |0\rangle$, where P is a local Pauli operator and W is a *light-cone circuit* acting only on a limited number of qubits. The overlaps of these auxiliary states fully determine both the QGT and the gradient (see Eqs. 9–15 in [7]), and the locality of W means that the vast majority of entries can be simulated exactly on a classical device.

The method proceeds through three stages: (1) **Light-cone analysis** identifies the qubit support for each parameter’s light cone through circuit analysis. (2) **Hybrid simulation** simulates most circuits exactly via statevector methods while handling the more complex ones through measurements on the quantum computer or use of tensor networks. (3) **Overlap computation** efficiently computes the sparse overlap between all auxiliary quantum states to evaluate the QGT and gradient entries.

¹We refer the reader to Fig. 1 in [7] to understand where the name *bowtie* comes from.

It should be noted that step (2) (the computational bottleneck of the pipeline) can be trivially parallelized between exact statevector simulation, approximate classical methods such as tensor networks [12] and Pauli propagation [13], and quantum hardware. This flexible framework allows to concentrate most quantum resources into relevant tasks, such as sampling from the state preparation (which is not necessarily trivial with classical methods) or extending the dynamics by appending extra layers of Trotter steps after the variational circuit.

3 Results

We start by analyzing the cost of estimating variational quantities using different methods in literature. We provide an analysis of the resources needed for state of the art stochastic approximations of the QGT (such as SPSA [14]) to reach useful accuracies for VarQTE. Despite generous assumptions (such as no shot noise or approximation errors) we find that the number of function evaluations required to keep the approximation error below the threshold at which VarQTE still produces reliable quantum dynamics is so large that it would be prohibitively expensive in practice (whether executed on quantum hardware or classical simulators). In contrast, bowtie VarQTE computes the same QGT entries exactly, eliminating sampling noise and improving numerical stability while keeping classical resource usage in check.




Figure 1: Comparison of bowtie VarQTE with AQC. Bowtie VarQTE achieves comparable state preparation to AQC with a strict truncation threshold and using the same ansatz, and similar results to a second order Trotter circuit with twice the depth. We use $n = 50$ qubits.




Figure 2: SKQD results. On top, we have the overlap between our variational state and the ground state. At the bottom, we have the energy estimation; in (blue) we have the VarQTE energy and in (red) we have the energy obtained through sampling from the Krylov basis.

We benchmark bowtie VarQTE against AQC and Trotterized evolution. AQC is a circuit compilation method in which both the variational state for a given set of parameters and the target state (typically obtained from a much deeper Trotter circuit) are simulated with tensor networks. Gradient-based optimization is then used to minimize the difference between the two states, yielding a more hardware-efficient circuit that prepares the target state. As shown in Fig. 1, bowtie VarQTE achieves comparable fidelities to AQC throughout the time evolution and shows significant resource savings compared to conventional Trotter circuits.

Contrary to the common intuition in the literature (where the QGT is considered the expensive quantity due to its quadratic number of entries relative to the number of circuit parameters, versus a linear number of gradient entries), we find that estimating the gradient is significantly more expensive in practice. We demonstrate that computing a single QGT entry is mathematically equivalent to computing the gradient of each term in the Hamiltonian (see

Appendix C in [7]). Moreover, while the QGT entries of parameters at the beginning of the circuit are trivial to compute (owing to the small light-cones they span), the light-cones of the gradient terms are generally larger. The leading computational cost in our algorithm arises from simulating the auxiliary states in step (2), whereas computing their overlaps in step (3) is relatively inexpensive. Consequently, most of the total computation time is spent on the few but large light-cone circuits, which correspond to the terms in the Hamiltonian. A detailed breakdown of computational times and circuit sizes for each auxiliary circuit is provided in Fig. 5 [7]; the key results are summarized in Table 1.

Table 1: Range of light-cone sizes for the auxiliary states of several circuits and Hamiltonians.

Topology	System	Qubits	Layers	Parameter	Hamiltonian
1D chain	Heisenberg	35	4	1–18	10–20
1D chain	Heisenberg	50	3	1–14	8–16
Heavy-hex	Heisenberg	63	2.5	1–22	9–27

As a proof of principle, we study the application of bowtie VarQTE in the context of SKQD. The SKQD pipeline can be split into (1) finding a good initial state that maximizes the overlap with the ground state, (2) performing real-time evolution on this state, and (3) sampling from the collection of time-evolved states to build a relevant subspace for classical diagonalization. Bowtie VarQTE addresses each stage naturally. For a given quantum circuit, we can (1) approximate the ground state using VarQITE for imaginary-time evolution, (2) continue evolving the ansatz under real-time evolution using VarQRTE, and (3) focus our quantum resources on sampling from the resulting circuits (and potentially extensions via appended Trotter steps) to build a Krylov subspace. In Fig. 2, we show the results obtained for a Heisenberg Hamiltonian on a 63-qubit heavy-hex lattice. For more details on the system, refer to Eqs. 20–23 and Fig. 3 in [7].

4 Discussion and Conclusion

Bowtie VarQTE provides a resource-efficient state preparation primitive that fits naturally into hybrid quantum-classical workflows: the causal structure of parametrized circuits enables exact evaluation of variational quantities wherever light-cones admit exact simulation, ensuring numerical stability while concentrating quantum resources on classically hard tasks. Unlike methods such as AQC, it does not require a classical representation of the target state, broadening its applicability to regimes where such a representation is unavailable or costly. The framework is also modular: when light-cones grow beyond exact simulation, the simulation backend can be seamlessly replaced by tensor networks or quantum hardware, enabling a flexible hybrid strategy. The prepared circuits are directly deployable on hardware and serve as building blocks for downstream algorithms: imaginary-time evolution (VarQITE) yields warm-start states for ground-state search, while real-time evolution (VarQRTE) generates time-evolved ansätze, as demonstrated here within the SKQD pipeline, highlighting a particularly promising application direction. To the best of our knowledge, this constitutes the largest-scale demonstration of VarQTE in the literature: prior work was limited to approximately 20 qubits or relied on poor QGT approximations [14], whereas this work reaches 63 qubits with a reasonably deep circuit, and the light-cone methodology scales straightforwardly beyond 100 qubits. Importantly, this also provides the first rigorous noiseless baseline for VarQTE at scales where such a baseline was previously out of reach, separating algorithmic error from hardware noise (a distinction the field has so far lacked). We additionally release an open-source library [15] implementing the bowtie pipeline, enabling researchers to readily explore and extend VarQTE for their own systems.

References

- [1] Earl Campbell. Random Compiler for Fast Hamiltonian Simulation. *Physical Review Letters*, 123(7):070503, August 2019. Publisher: American Physical Society.
- [2] Masuo Suzuki. Generalized trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Communications in Mathematical Physics*, 51:183–190, 1976.
- [3] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.
- [4] Sam McArdle, Tyson Jones, Suguru Endo, Ying Li, Simon C. Benjamin, and Xiao Yuan. Variational ansatz-based quantum simulation of imaginary time evolution. *npj Quantum Information*, 5(1):75, 2019.
- [5] Xiao Yuan, Suguru Endo, Qi Zhao, Ying Li, and Simon Benjamin. Theory of variational quantum simulation, September 2019. arXiv:1812.08767 [quant-ph].
- [6] A.D. McLachlan. A variational solution of the time-dependent Schrödinger equation. *Molecular Physics*, 8(1), 1964.
- [7] Marc Drudis, Alberto Baiardi, Mattia Chiurco, Francesco Tacchino, Stefan Woerner, and Christa Zoufal. Bowtie varqte: A resource-efficient quantum state preparation primitive, 2026.
- [8] Andrew Eddins, Minh C. Tran, and Patrick Rall. Lightcone shading for classically accelerated quantum error mitigation, September 2024. arXiv:2409.04401 [quant-ph].
- [9] Niall Robertson, Albert Akhriev, Jiri Vala, and Sergiy Zhuk. Approximate Quantum Compiling for Quantum Simulation: A Tensor Network Based Approach. *ACM Transactions on Quantum Computing*, 6(3):1–15, 2025.
- [10] George Pennington, Kevin C. Smith, James R. Garrison, Lachlan P. Lindoy, Jason Crain, and Ben Jaderberg. Preparing 100-qubit symmetry-protected topological order on a digital quantum computer, 2026.
- [11] Arseny Kovyrshin, Dilhan Manawadu, Edoardo Altamura, George Pennington, and et al. Approximate quantum circuit compilation for proton-transfer kinetics on quantum processors. *Phys. Chem. Chem. Phys.*, 28:3035–3054, 2026.
- [12] David Elieser Deutsch. Quantum computational networks. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 425(1868):73–90, 09 1989.
- [13] Manuel S Rudolph, Tyson Jones, Yanting Teng, Armando Angrisani, and Zoe Holmes. Pauli propagation: A computational framework for simulating quantum systems. *arXiv preprint arXiv:2501.13101*, 2025.
- [14] Julien Gacon, Christa Zoufal, Giuseppe Carleo, and Stefan Woerner. Simultaneous Perturbation Stochastic Approximation of the Quantum Fisher Information. *Quantum*, 5:567, October 2021.
- [15] Marc Drudis. Bowtieqgt: Efficient quantum geometric tensor computation. <https://doi.org/10.5281/zenodo.20210713>, May 2026.