Variational Quantum Circuit Decoupling

Ximing Wang, Chengran Yang, and Mile Gu

Decoupling systems into independently evolving components has a long history of simplifying seemingly complex systems. They enable a better understanding of the underlying dynamics and causal structures while providing more efficient means to simulate such processes on a computer. Here we outline a variational decoupling algorithm for decoupling unitary quantum dynamics – allowing us to decompose a given n-qubit unitary gate into multiple independently evolving subcomponents. We apply this approach to quantum circuit synthesis - the task of discovering quantum circuit implementations of target unitary dynamics. Our numerical studies illustrate significant benefits, showing that variational decoupling enables us to synthesize general 2 and 4-qubit gates to fidelity that conventional variational circuits cannot reach.

When studying a system of coupled harmonic oscillators, a key approach is to decouple its dynamics into various normal modes. We can then study each mode individually, enabling a deeper understanding of underlying dynamics. In the era of computer simulation, this divide-and-conquer approach has further operational importance, enabling leverage of parallel processing [1, 2], and often drastically reducing the computational costs of simulation or optimization. Decoupling complex systems into simpler components has seen success in diverse settings, from modeling the motion of human hands to hydraulic simulation [3, 4].

In quantum systems, decoupling can also be immensely beneficial. Many quantum computing and simulation algorithms require the approximate synthesis of desired n-qubit operations using shallow circuits of elementary quantum gates. Yet, the exponential growth of Hilbert space makes this a highly intractable task [5–7]. Meanwhile, quantum systems offer entirely new possibilities for decoupling. Even in strongly interacting circuits, separable dynamics may exist on a non-local basis [8]. The mutual information between decoupled quantum subsystems is generally much lower than classical counterparts [9].

Here, we propose embedding a divide-and-conquer approach to enhance existing variational circuit algorithms [5, 6]. The procedure involves designing a variational decoupling algorithm that first breaks a many-qubit interaction into smaller components (see figure 1). We can then optimize these smaller components individually or break them down further via recursive applications of variational decoupling. We demonstrate the method for compiling circuits approximating arbitrary two and four-qubit gates, where it identifies circuits of much higher fidelities than direct variational methods.

Framework – Consider being given black-box access to some unknown n-qubit unitary quantum process U. Our goal is to approximate this process with a sequence of elementary quantum gates - subject perhaps to constraints on circuit complexity or depth. That is, discover some circuit W - programmable on quantum computers - that implements U. Such a problem arises naturally in tasks such as circuit compilation or quantum simulation [5, 6],

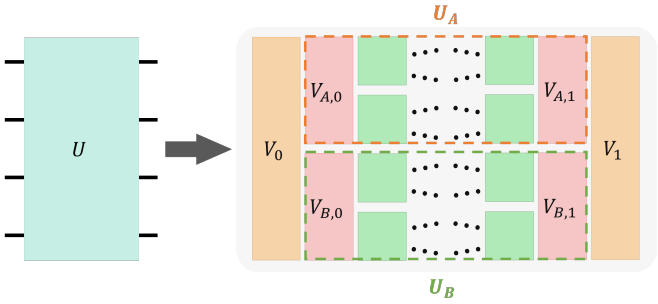

FIG. 1. **Variational Quantum Circuit Decoupling.** Given a target complex unitary quantum operation U, the decoupling of U involves identifying shallow pre and post-processing operation V_0 and V_1, such that $U \approx V_1(U_A \otimes U_B)V_0$, for some operators U_A and U_B act locally on subsystems A and B respectively. This procedure can be repeated recursively, breaking down U_A and U_B into smaller components, until all subsystems are sufficiently small.

where U represents a complex quantum circuit we wish to simplify; some unknown quantum device we wish to reverse-engineer or some natural process we wish to simulate.

Systematic solutions to this problem are clearly intractable. Process tomography to determine the matrix elements of U alone would scale exponentially with n, as would any general method of decomposing U into elementary gates. This has motivated variational quantum algorithms (VQAs) for circuit discovery [5, 6]. Such approaches typically involve two components: (a) An ansatz, a family of quantum circuits $W(\boldsymbol{\theta})$ parameterized by some k-dimensional vector $\boldsymbol{\theta}$, and (b) A cost function C that that measures the performance loss of each candidate circuit in approximating W. The goal is to identify the parameter set $\boldsymbol{\theta}$ so that the corresponding circuit $W(\boldsymbol{\theta})$ has a minimal cost. Meanwhile, the cost function is generally chosen to be normalised (takes a maximum value of 1) with the following properties:

- **Measurable and Gradient Measurable**, such that $C(W)$ and each partial derivative $\frac{\partial}{\partial\theta_i}C$ can be efficiently estimated.

- **Faithful**, such that C is non-negative, and $C(W) = 0$ if and only if no general n-qubit quantum circuit are preferred over W.

This then allows gradient descent-based methods to find $\boldsymbol{\theta}$ within the parameterized circuit to minimize the cost. However, the rate at which variational circuits converge can become severely limited due to barren plateaus - especially in scenarios where circuits being optimized over had no constraints [11–13].

Variational Decoupling – We adopt a divide-and-conquer approach. Instead of attempting to learn a circuit decomposition for the entirety of U, we break the circuit down into a series of smaller circuits as illustrated in figure 1. Specifically we first divide n-qubit system into two subsystems A and B. We then identify pre-processing unitary V_0 and post-processing unitary V_1 such that there exists U_A on A and U_B on B in which $V_1^\dagger U V_0^\dagger$ most closely approximates $U_A \otimes U_B$.

We document below a variational means to learn gate decomposition for V_0 and V_1 without needing any knowledge of unitary operators U_A and U_B. This decoupling can proceed recursively until all localized operators are sufficiently small for conventional variational methods. Finally, various sub-problems are recombined to give a full circuit decomposition of U.

The key behind the variational decoupling algorithm is to find an appropriate cost function. First, let $V_0(\boldsymbol{\theta}_0)$ and $V_1(\boldsymbol{\theta}_1)$ be circuit paramatrizations of V_0 and V_1 with respective parameters $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$. Meanwhile set $W = V_1^\dagger U V_0^\dagger$ as the resulting quantum process formed by pre and postprocessing of U. Since we can always engineer W given U, V_0, and V_1, we can define our cost function concerning W alone for convenience. Here we desire a normalised cost function that (i) is faithful, such that $C(W) = 0$ when $U = V_1 U V_0^\dagger = U_A \otimes U_B$ or $(U_A \otimes U_B)S$, and (ii) exhibits both self and gradient measurability. Here we propose the cost function as the decoupling cost C_D defined as

$$C_D(W) = \iint D[W(|\psi\rangle\langle\psi|_A \otimes |\phi\rangle\langle\phi|_B)W^\dagger] \mathrm{d}\,|\psi\rangle\,\mathrm{d}\,|\phi\rangle\,,$$

where $D(\rho_{A,B}) = \frac{4^m}{(4^m-1)} \frac{L(\rho_A)+L(\rho_B)}{2}$, $L(\rho) = 1 - \mathrm{Tr}[\rho^2]$ is the linear entropy and the integration is take over Haar random initial states $|\psi\rangle_A$ on A and $|\Phi\rangle_B$ on B. Here m is either the number of qubits in A or B, whichever is smaller. We now illustrate the following:

Theorem 1. *The decoupling cost C_D is efficiently measurable, efficiently differentiable, and faithful.*

To show C_D is *efficiently measurable* and *efficiently differentiable*, we developed a method to evaluate it using a circuit with twice the size of W without the need to prepare Harr random states, as illustrated in figure 2. We show that this state can be prepared with a constant depth of quantum circuit with gate determinations efficiently achievable through a classical process. The gradient of C_D can be measured by the same circuit with a generalised parameter-shift rule [14–17].

To establish *faithfulness*, first observe that the linear entropy is used to quantify the entanglement between two subsystems [18]. $D(W(\rho_A \otimes \rho_B)W^\dagger)$ quantifies the amount of entanglement W generates between A and B when applied to a product state. Consequently, our cost function C_D represents the expected amount of entanglement generated by W when averaged across all pure product inputs. As such, $C_D(W) = 0$ if and only if the operator W is equivalent to the identity or the swap gate up to local unitary gates.

We can then proceed with variational decoupling by choosing appropriate antsazes for $V_0(\boldsymbol{\theta}_0)$ and $V_1(\boldsymbol{\theta}_1)$; such that W is parameterized by the joint vector $\boldsymbol{\theta}^* := \boldsymbol{\theta}_0^* \oplus \boldsymbol{\theta}_1^*$. Note that, just as in standard circuit compilation, we may also choose to limit the gate depth of V_0 and V_1 should we wish for shallow circuit approximations of U.

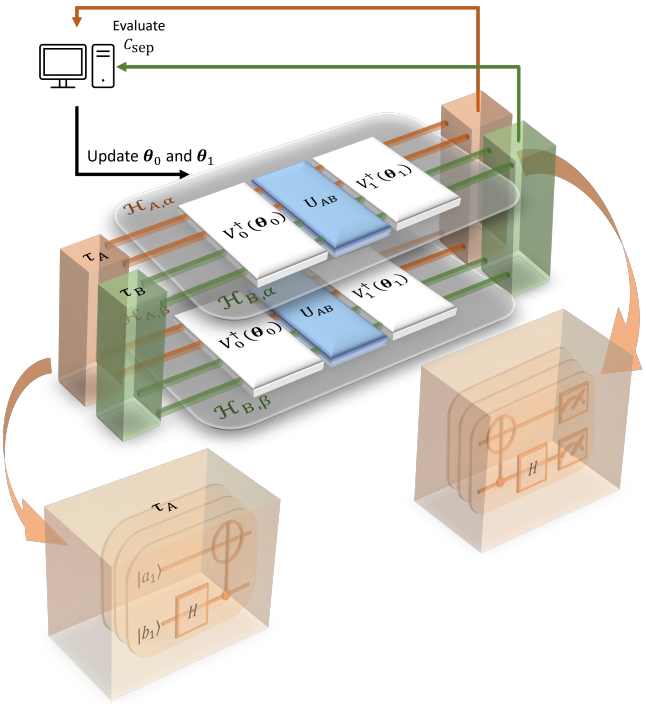

FIG. 2. *The schematic diagram of the variational decoupling algorithm.* The qubits in the quantum computer are virtually divided into 4 subsystems: $\mathcal{X}_{A,\alpha}$, $\mathcal{X}_{B,\alpha}$, $\mathcal{X}_{A,\beta}$, and $\mathcal{X}_{B,\beta}$. Two copies of the evolution $W := V_1^\dagger U V_0^\dagger$ are applied to each copy of the system $\mathcal{X}_{A,i} \otimes \mathcal{X}_{B,i}$. Let d_k be the dimension of the system $\mathcal{X}_{k,i}$. The input state is initialized as $\tau_A \otimes \tau_B$, where τ_k is the projector state to the symmetric subspace of $\mathcal{X}_{k,\alpha} \otimes \mathcal{X}_{k,\beta}$. The cost function C_D is measured at the end with a destructive swap test and minimized by tuning the parameters $\boldsymbol{\theta}$ for the next iteration.

Performance Comparisons – We demonstrate the improved efficacy of our decoupling method in two distinct scenarios:

- *The exact compilation of a general* 2*-qubit gate* - where the ansatz is shown figure 3. Here, the circuit family is universal, and can theoretically synthesize any 2-qubit unitary operator [19]. Our decoupling algorithm breaks this into two stages. The first (coloured orange) decouples U into two single-qubit circuits. The second stage (red) then optimises these circuits.

- *A shallow circuit approximation of a general* 4*-qubit gate*, where the structure of the ansatz follows the figure 4. For both V_0 and V_1, we utilize four layers of decoupling circuit, where each layer comprises a shallow circuit and single qubit and control gates. The circuit compilation is done via three stages. Stage one (orange) decouples the circuit into two generic 2-qubit gates. The second stage (red) decouples these 2-qubit gates further into 4 single qubit interactions. Finally, stage 3 then optimises each of the four single-qubit gates.

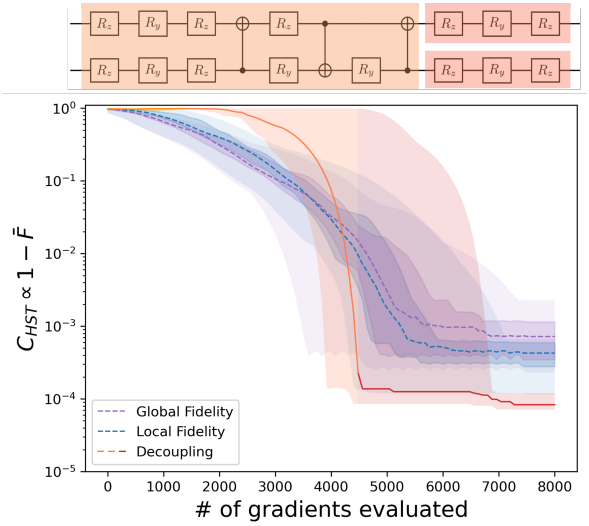

FIG. 3. *Variational Compilation of a general 2-qubit unitary with ADAM optimization.* The fidelity $\bar{F}$ is compared between the direct compilation (purple for global cost functions and blue for local cost functions) and our decoupling method (solid lines). For each method, the training is repeated 20 times with different initializations. Our method divides the task into two steps, each represented by a different color. The first phase (yellow) decouples the unitary to single-qubit unitary operators, and the second phase (red) maximizes the fidelities. The phase transition region is colored orange, where some of the processes are in phase one and the rest are in phase two. The data of each method are divided into quartiles, and the lines represent the median of each process. The second and third quartiles are shaded with the corresponding color.

In each setting, we select target U at random according to the Harr measure. We benchmark our method against standard variational approaches, where all circuit parameters are optimized by gradient descent to maximize the gate fidelity $\bar{F}$ simultaneously. In the standard approach, this optimization is done by minimizing a cost function C_{HST} or the localized cost function C_{LHST} [5].

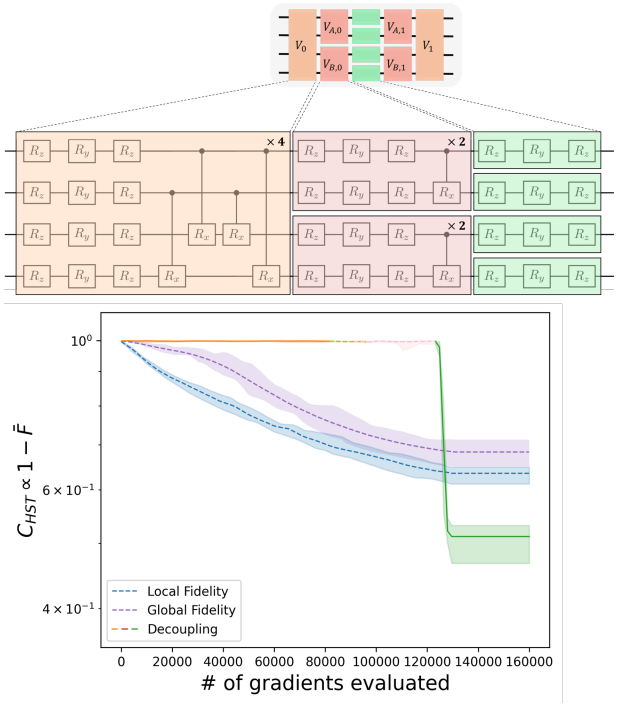

FIG. 4. *Variational Compilation of generic 4-qubit unitary with ADAM optimization.* Similar to the 2-qubit case, the data of each process is collected from 20 independent training processes, while the regions between the first and last quartiles are shaded with the corresponding colors as in FIG 3. Here, the optimisation is divided into three phases. The circuits are divided into 3 types: 4-qubit ansatzes (orange), 2-qubit ansatzes (red), and the middle single qubit gates (green). The ansatzes are designed to explore stronger expressive power with a limited number of gates, which can be generalized to any number of qubits [20].

Discussion – Here, a variational mean to decouple complex quantum circuits is proposed. There are several natural avenues for future research. The first is in studies of noise resilience – whereby the dynamics of processes we wish to discover or the gate-set we choose to implement are not pure. The second is the extension of our divide-and-conquer approach to variational settings beyond circuit synthesis. In particular, the decoupling method could be generalized to the digital-analog quantum computing paradigm [21], where part of the quantum logical circuit may be replaced by tunable hardware-dependent evolutions. This combination of digital and analog quantum processes enlarges the set of efficiently implementable ansatzes for quantum compilation.

[1] Keshav Pingali, Donald Nguyen, Milind Kulkarni, Martin Burtscher, M Amber Hassaan, Rashid Kaleem, Tsung-Hsien Lee, Andrew Lenharth, Roman Manevich, Mario Mendez-Lojo, Dimitrios Prountzos, and Xin Sui. The Tao of Parallelism in Algorithms. In *32nd ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 12–25, 2011.

[2] P. W. Trinder, K. Hammond, H.-W. Loidl, and S. L. Peyton Jones. Algorithm + strategy = parallelism. *J. Funct. Prog.*, 8(1):23–60, January 1998.

[3] Ying Wu and Thomas S Huang. Capturing articulated human hand motion: A divide-and-conquer approach. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 1, pages 606–611. IEEE, 1999.

[4] Peng Wang, Bo Yu, Dongxu Han, Dongliang Sun, and Yue Xiang. Fast method for the hydraulic simulation of natural gas pipeline networks based on the divide-and-conquer approach. *Journal of Natural Gas Science and Engineering*, 50:55–63, 2018.

[5] Sumeet Khatri, Ryan LaRose, Alexander Poremba, Lukasz Cincio, Andrew T. Sornborger, and Patrick J. Coles. Quantum-assisted quantum compiling. *Quantum*, 3:140, May 2019.

[6] Nikita A Nemkov, Evgeniy O Kiktenko, Ilia A Luchnikov, and Aleksey K Fedorov. Efficient variational synthesis of quantum circuits with coherent multi-start optimization. *Quantum*, 7:993, 2023.

[7] Marc G. Davis, Ethan Smith, Ana Tudor, Koushik Sen, Irfan Siddiqi, and Costin Iancu. Towards Optimal Topology Aware Quantum Circuit Synthesis. In *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 223–234, October 2020.

[8] Giuseppe De Tomasi, Daniel Hetterich, Pablo Sala, and Frank Pollmann. Dynamics of strongly interacting systems: From Fock-space fragmentation to Many-Body Localization. *Phys. Rev. B*, 100(21):214313, December 2019.

[9] Max Tegmark. Consciousness as a state of matter. *Chaos, Solitons & Fractals*, 76:238–270, 2015.

[10] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization, January 2017.

[11] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nat Commun*, 9(1):4812, November 2018.

[12] Samson Wang, Enrico Fontana, M. Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J. Coles. Noise-induced barren plateaus in variational quantum algorithms. *Nat Commun*, 12(1):6961, November 2021.

[13] M. Cerezo and Patrick J. Coles. Higher Order Derivatives of Quantum Neural Networks with Barren Plateaus. *Quantum Sci. Technol.*, 6(3):035006, July 2021.

[14] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. Quantum circuit learning. *Phys. Rev. A*, 98(3):032309, September 2018.

[15] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Phys. Rev. A*, 99(3):032331, March 2019.

[16] Jun Li, Xiaodong Yang, Xinhua Peng, and Chang-Pu Sun. Hybrid Quantum-Classical Approach to Quantum Optimal Control. *Phys. Rev. Lett.*, 118(15):150503, April 2017.

[17] Dawei Lu, Keren Li, Jun Li, Hemant Katiyar, Annie Jihyun Park, Guanru Feng, Tao Xin, Hang Li, Guilu Long, Aharon Brodutch, Jonathan Baugh, Bei Zeng, and Raymond Laflamme. Enhancing quantum control by bootstrapping a quantum processor of 12 qubits. *npj Quantum Inf*, 3(1):1–7, October 2017.

[18] E. Santos and M. Ferrero. Linear entropy and Bell inequalities. *Phys. Rev. A*, 62(2):024101, July 2000.

[19] Vivek V. Shende, Igor L. Markov, and Stephen S. Bullock. Minimal Universal Two-qubit Quantum Circuits. *Phys. Rev. A*, 69(6):062321, June 2004.

[20] Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, and Dacheng Tao. Expressive power of parametrized quantum circuits. *Phys. Rev. Research*, 2(3):033125, July 2020.

[21] Lucas Lamata, Adrian Parra-Rodriguez, Mikel Sanz, and Enrique Solano. Digital-analog quantum simulations with superconducting circuits. *Advances in Physics: X*, 3(1):1457981, January 2018.

[22] Leonardo Banchi and Gavin E. Crooks. Measuring Analytic Gradients of General Quantum Evolution with the Stochastic Parameter Shift Rule. *Quantum*, 5:386, January 2021.

[23] The $\alpha = 0.01$, $\delta = 0.0005$, $\beta_1 = 0.8$, $\beta_2 = 0.9$.