

NAPA: Intermediate-level Variational Native-pulse Ansatz for Variational Quantum Algorithms

Zhiding Liang, Jinglei Cheng, Hang Ren, Hanrui Wang, Fei Hua, Zhixin Song, Yongshan Ding,
Frederic T. Chong, Song Han, Xuehai Qian, Yiyu Shi

Abstract—Variational quantum algorithms (VQAs) have demonstrated great potentials in the Noisy Intermediate Scale Quantum (NISQ) era. In the workflow of VQA, the parameters of ansatz are iteratively updated to approximate the desired quantum states. We have seen various efforts to draft better ansatz with less gates. Some works consider the physical meaning of the underlying circuits, while others adopt the ideas of neural architecture search (NAS) for ansatz generator. However, these designs do not exploit the full advantages of VQAs. Because most techniques target gate ansatz, and the parameters are usually rotation angles of the gates. In quantum computers, the gate ansatz will eventually be transformed into control signals such as microwave pulses on superconducting qubits. These control pulses need elaborate calibrations to minimize the errors such as over-rotation and under-rotation. In the case of VQAs, this procedure will introduce redundancy, but the variational properties of VQAs can naturally handle problems of over-rotation and under-rotation by updating the amplitude and frequency parameters. Therefore, we propose NAPA, a native-pulse ansatz generator framework for VQAs. We generate native-pulse ansatz with trainable parameters for amplitudes and frequencies. In our proposed NAPA, we are tuning parametric pulses, which are natively supported on NISQ computers. Given the limited availability of gradient-based optimizers for pulse-level quantum programs, we choose to deploy non-gradient optimizers in our framework. To constrain the number of parameters sent to the optimizer, we adopt a progressive way to generate our native-pulse ansatz. Experiments are conducted on both simulators and quantum devices for Variational Quantum Eigensolver (VQE) tasks to evaluate our methods. When adopted on NISQ machines, NAPA obtained improved performance with decreased latency by an average of 86%. NAPA is able to achieve 96.482% and 99.336% accuracy for VQE tasks on H_2 and HeH^+ respectively. An average accuracy of 97.27% is achieved for medium-size quantum chemistry tasks on CO_2 , H_2O , and NaH . NAPA also demonstrates advantages on quantum optimization tasks even with considerable noises in NISQ machines.

I. INTRODUCTION

Operating on the principles of quantum mechanics, quantum computers have the potential to solve problems that are in-

Z. Liang and Y. Shi are with the Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN, 46556 USA.

J. Cheng and X. Qian are with the Department of Computer Science, Purdue University, West Lafayette, IN 47907 USA.

H. Ren is with the Department of Chemistry, University of California, Berkeley Berkeley, CA 94720 USA.

H. Wang and S. Han are with the Department of EECS, Massachusetts Institute of Technology, Cambridge, MA 02139 USA. F. Hua is with the Department of Computer Science, Rutgers University, New Brunswick, NJ 08901 USA.

Z. Song is with the Department of Physics, Georgia Institute of Technology, Atlanta, GA 30332 USA.

Y. Ding is with the Department of Computer Science, Yale University, New Haven, CT 06511 USA.

F. Chong is with the Department of Computer Science, University of Chicago, Chicago, Illinois 60637 USA.

Address comments to Zhiding Liang (email: zliang5@nd.edu).

tractable on classical computers [4]. As hardware technologies and quantum algorithms advance rapidly, today's quantum computers begin to demonstrate their potentials in solving problems of non-trivial size in areas such as quantum chemistry [7]. In 2019, Google claimed to have achieved quantum supremacy with the task of random circuit sampling on a 53-qubit quantum computer [6]. IBM introduced the Condor, a 1,121 superconducting qubit quantum processor at the end of 2023, together with a roadmap for modular quantum processors with more than 10,000 qubits in the coming years [1]. Recent studies have demonstrated promising approaches for constructing logical quantum processors using error correction techniques on reconfigurable atom arrays [9]. In the current NISQ era, however, emerging quantum devices are still prone to errors and sensitive to decoherence [19].

In superconducting quantum computers, qubits are sparse and not perfectly isolated from their environments. These quantum devices have not yet met the hardware requirements for error correction methods such as surface codes [16] due to the limited number of qubits and low gate fidelity. Without quantum error correction, current quantum computers can only handle small-scale circuits before running into irreversible errors, making practical-size algorithms infeasible. However, with elaborately designed noise-resilient algorithms, we can still expect to achieve quantum advantages in areas such as quantum chemistry [12], [49] much sooner than other applications like database search [28] and integer factorization [46].

Variational quantum algorithms [13], [35] have shown great noise resilience, and are considered as hybrid algorithms, where some parts are performed on a quantum device and others on a classical computer. Variational quantum eigensolver (VQE) [32], [50] is one of the most promising candidates in the variational computing paradigm. With VQE, we are able to estimate the ground state energy of a targeting quantum system by iteratively updating a parametrized quantum circuit. Quantum approximate optimization algorithm (QAOA) [29] and quantum neural network (QNN) [3] are also members of VQAs. QAOA attempts to solve combinatorial optimization problems, including maxcut problems, while QNN has exhibited exceptional capabilities in representing complex data [30]. These algorithms are among the most promising examples of NISQ algorithms since the number of required quantum gates remains moderate.

Exposing Native Pulse-Level Controls. The majority of existing quantum computers do not provide access for analog controls of qubits. Consequently, nearly all compilers implement a gate-based workflow [48], in which quantum algorithms are synthesized, compiled on classical computers, and finally executed on quantum computers. In the first place,

Method	Robustness		Parameters of Control		Applications		
	Noise-aware	System model	Amp.	Freq.	VQE	QAOA	QML
Gate-level (conventional)	✓	✓	✗	✗	✓	✓	✓
Pulse-level attempts [15], [38]	✗	✗	✓	✓	✓	✗	✗
NAPA (Proposed)	✓	✓	✓	✓	✓	✓	✓

TABLE I: Comparison between gate-level approaches, pulse-level approaches, and the proposed NAPA.

quantum circuits are generated or synthesized to implement certain functions of quantum algorithms. These gate circuits are usually not compatible with the underlying topology of quantum hardware. Therefore, SWAP gates need to be inserted to make quantum circuits executable on quantum computers. Then these circuits are decomposed into single-qubit and two-qubit gates that are natively supported by quantum computers. Finally, the circuits are dispatched to quantum backends, where they are “translated” into control signals on physical qubits such as transmons [44], trapped ions [10] and photons [33]. For superconducting quantum computers, the control signals are microwave pulses [5].

There is more flexibility if gate circuits are decomposed and controlled at pulse level, since we are dealing with a more fine-grained abstraction layer. The challenge now is to figure out how to effectively generate pulses for quantum algorithms. Existing techniques such as quantum optimal control (QOC) [15], [38] can be adopted to generate control pulses. QOC devises and implements the shapes of external controls on qubits to accomplish given tasks. As indicated in [45], QOC can handle quantum circuits of moderate size, but the scalability might be the issue. Despite various efforts [14] to optimize QOC and reduce its overhead, it is still computationally expensive. Besides, we need to consider the high-noise feature of NISQ devices, on which it is preferred to have a fixed set of allowed operations. On current quantum computers, a small group of gates are carefully calibrated regularly to maintain their accuracies. Overall, it is generally hard to take advantage of quantum pulses with NISQ machines.

Why Pulses for VQAs? Situations differ slightly in the case of VQAs due to its variational characteristics. During the “training” process, the parametric circuits of VQAs are updated iteratively. It is now unimportant whether the controls are accurately implemented on quantum hardware as long as the parametric circuits can reach the desired states. Instead of using gate-level compilation or QOC, we can implement quantum algorithms at the “native-pulse level”. At the native-pulse level, we can directly manipulate the native pulses that are supported by the quantum hardware. This paradigm change grants more fine-grained control and thus making it possible for better performance, scalability, and robustness. A recent work [24] has shed light on the feasibility and potentials for such a paradigm change. However, the research on pulse level optimization is still in its infancy. The capability of quantum pulses has not been fully explored, nor are they already robust or scalable on NISQ machines. As summarized in Table I, many critical issues remain unsolved.

To tackle these challenges, we propose NAPA, a native-pulse ansatz generator for VQAs. NAPA is the first to demonstrate the feasibility of native-pulse ansatz on NISQ machines.

Instead of using rotation gates, we directly use the pulses that are natively supported by the quantum processors. Compared with QOC, NAPA has fewer parameters and can be easily deployed onto NISQ machines, while QOC with realistic system models requires huge computation resources. On the other side, NAPA is superior to gate-based methods, since NAPA drops the abstraction layer of native gates and results in less circuit latency. We provide results from IBM’s superconducting quantum computers [36], while previous pulse-level works are only evaluated on simulators.

Contributions. The goal of this paper is to construct native-pulse ansatz for VQAs and demonstrate in both simulators and NISQ machines. The major contributions of NAPA include:

- **Native-pulse ansatz.** Our native-pulse ansatz is derived from native pulses that are extracted from quantum backends. In this way, we ensure that pulse ansatz is compatible with quantum hardware.
- **Progressive learning.** In NAPA, a non-gradient optimizer is employed. We provide a progressive way to “grow” our native-pulse ansatz in order to maintain a reasonable size for the parameters handled by the optimizer. New pulse blocks with zero amplitudes are appended at different “steps” of NAPA. This prevents the appended pulses from abruptly changing the ansatz circuits’ overall unitaries.
- **Results from NISQ machines.** Experiments are conducted on simulators and NISQ machines. The results show that the native-pulse ansatz outperforms the gate ansatz for VQA tasks in terms of accuracy and latency.
- **Exploration on frequency tuning.** We explore possible benefits of tuning pulse frequencies on transmons. Experimental results show that pulse frequency can be an extra degree of freedom for native-pulse ansatz.

Evaluation Highlights. Six NISQ machines are used to evaluate NAPA. We achieve latency reductions of up to **97.3%** compared to baselines. Accuracies up to **99.895%** are attained for small-size VQE tasks.

II. BACKGROUND

A. Variational Quantum Algorithms (VQAs)

The variational quantum eigensolver (VQE) is one of the most popular and promising VQAs. VQE is employed primarily to solve the ground and low-excited states of quantum systems. It has important applications in quantum many-body physics, quantum chemistry, and other fields [32], [42]. VQE has also been shown to be noise resistant in NISQ devices [50]. The Quantum Approximate Optimization Algorithm (QAOA) is a VQA that solves combinatorial optimization problems with sub-optimal solutions [29]. Quantum neural network (QNN) is a model of quantum machine learning (QML) that

use ansatz circuits to extract features from input data, followed by complex-valued linear transformations. QNN has great potential for applications in QML [8], quantum simulation, and optimization [39].

B. Quantum Optimal Control

Assuming we have a closed quantum system, the Hamiltonian of the system is provided by

$$H(t) = H_0 + \sum_{j=1} u_j(t) H(j) \quad (1)$$

where H_0 is the drift Hamiltonian, $H(j)$ is the control Hamiltonian, and u_j is time-dependent control signal. The Schrödinger's equation governs the system's dynamics:

$$\frac{d}{dt}|\psi\rangle = -iH(t)|\psi_0\rangle \quad (2)$$

where ψ_0 denotes the system's state at time $t = 0$. Quantum optimum control (QOC) can be used to calculate $u_j(t)$ which corresponds with control signals. With QOC, we can transform one quantum state into a desired state. Typically, a cost function is specified as the fidelity of generated pulses, which measures the difference between the simulated unitary matrix and the desired unitary matrix. Algorithms like GRAPE [17] and CRAB [11] have been developed to solve QOC problem.

C. Quantum Pulse Learning

Another method of generating pulses is to parameterize the quantum pulses and then optimize the parameters. We refer to such process as quantum pulse learning. To comprehend the trainable components of control pulses for superconducting quantum computers, the Hamiltonian of transmon can be defined as follows:

$$H = \sum_{i=0}^1 (U_i(t) + D_i(t)) \sigma_i^X + \sum_{i=0}^1 2\pi\nu_i(1 - \sigma_i^Z)/2 + \omega_B a_B a_B^\dagger + \sum_{i=0}^1 g_i \sigma_i^X (a_B + a_B^\dagger) \quad (3)$$

$D_i(t)$ and $U_i(t)$ are two major terms that are determined by the pulse learning, they represent modulated signals as shown in Equation 4. They are obtained by mixing a local oscillator with control signals. σ_X , σ_Y , and σ_Z are Pauli operators. ν_i is the frequency of qubit i , g_i is the coupling strength between qubits, ω_B is the frequency of control buses, a_B and $a_B^\dagger$ are the ladder operator for control buses.

$$\begin{aligned} D_i(t) &= \text{Re}(d_i(t)e^{i\omega_{d_i}t}) \\ U_i(t) &= \text{Re}[u_i(t)e^{i(\omega_{d_i}-\omega_{d_j})t}] \end{aligned} \quad (4)$$

where $d_i(t)$ and $u_i(t)$ are the signals of qubit i on drive channel and control channel. Since pulse learning adjust $d_i(t)$ and $u_i(t)$, $D_i(t)$ and $U_i(t)$ are changed accordingly. Consequently, the drive Hamiltonian is updated [34]. In this way, we can manipulate the quantum system with control signals.




Fig. 1: Comparison between compilation process for gate level and pulse level. Gate-level workflow consists of several layers and introduces redundancy, pulse-level workflow consists fewer layers that can provide reduction of circuit latency.

III. RELATED WORK

Pulse Learning Approach: Ctrl-VQE [38] tweaks the pulse shapes to perform state preparation. The methods are evaluated with Qutip [31] pulse simulator. And the results demonstrate that the total time required for state preparation is greatly reduced. VQP [34] uses pulses as basic components to build the QNN ansatz and exhibits latency advantages over gate-based QNN on a two-class image classification task. The experiments are conducted on Qiskit [36] pulse simulator. VQOC [18] presents a mathematical formalism of optimal control, which acts on pulse optimization for VQA tasks. Their method is similar to Ctrl-VQE [38], but they take advantages of neutral atom's properties. The evaluations are also performed on simulator. These previous works are exploratory and rely on classical simulations of small quantum systems [34], [38]. NAPA, on the other hand, provides results from NISQ machines. Moreover, Ctrl-VQE is only designed with the consideration for single qubit pulse, whereas, NAPA both consider single qubit pulse and two qubit pulse.

Ansatz Architecture Search: Hamiltonian simulation plays an important role in simulating quantum systems [32], [42]. VQAs can be deployed to perform Hamiltonian simulation. And previous works [32] pointed out that the choice of ansatz is important for VQAs. The conventional approaches of choosing ansatz depends heavily on the applications. [40], [42] are specially designed for VQE tasks. The unitary coupled-cluster singles and doubles (UCCSD) is still the golden standards for VQE ansatz. QAS [20] proposes a noise-aware scheme to search for ansatz structure. The robustness to noise is demonstrated on simulators. QuantumNAS [51] presents a comprehensive framework for noise-adaptive co-search of the ansatz. Hardware topology is considered during the search algorithm. QuantumNAS validates their methods with evaluations on NISQ machines.

Pulse Gate Compilation: [26] proposes a new compilation paradigm, based on the OpenPulse interface for IBM quantum computers. This work achieves lower error rates and shorter execution times in comparison with traditional gate-based




Fig. 2: Illustrations of the redundancy introduced by current gate-level compilation workflow. In the case of gate-level compilation, pulses with fixed parameters are inserted to first rotate the qubit to another axis. With pulse-level controls, we can avoid such redundancy and reduce latency without loss of capability to explore the Hilbert space.

compilation methods. Their technique is bootstrapped from existing gate calibrations, thus their pulses are in a simple form. [45] designs a general quantum compilation method to integrate multiple operations into larger units and then achieve high efficiency by optimizing this aggregation and creating customized control pulses.

IV. MOTIVATION

A. Deficiencies of Gate-level Compiler

Figure 1 illustrates the compilation of a quantum program from high-level programming languages to the physical signals on quantum machines. Redundancy is introduced by the existing implementation of rotation gates on quantum hardware. For example, we want the qubit to evolve to the point on the Bloch sphere as indicated in Figure 2. In order to reach such state, a basis transformation, a phase shift, and a second basis transformation must be performed. This example illustrates the redundancy when quantum circuits are compiled into native gates. On the other hand, a single Derivative Removal by Adiabatic Gate (DRAG) pulse is capable of implementing an arbitrary single-qubit rotation. We propose to bypass the abstract layer of native gates and use native pulses directly as parametric elements in the ansatz. In this way, this approach results in a pulse circuit with reduced latency and increased number of parameters, enhancing its flexibility.

B. Drawbacks of Optimal Control Pulses

In addition, attempts have been made to optimize the quantum pulse generator. For example, [25] proposes to use quantum optimal control to generate pulses for given unitaries. The applications of quantum optimal control are severely constrained by the excessive cost of pulse generators and the complexity of NISQ devices' system models. Besides, QOC is usually adopted when a target unitary or state transition is known. Hence, it is incompatible with VQAs, which lack target unitaries. In addition, QOC's gradient-based approach with back propagation cannot be used, because gradients for pulses are not accessible on NISQ machines.

C. Potentials of Pulse Ansatz

As the illustration in Figure 2, if we can directly control pulses, we can eliminate some redundancy. Using pulses rather than gates can result in a lack of calibrations, which may result in inaccurate operations. However, for VQAs, the parameterized pulse ansatz can be automatically adjusted for errors including under-rotation and over-rotation. Besides amplitudes of native pulses, the frequencies of qubit channels can also be tuned. In conclusion, pulse ansatz offers more degrees of freedom, enabling us to search for desired states with substantially shorter pulse latencies.

V. OVERVIEW OF NAPA FRAMEWORK

In the proposed framework for native-pulse ansatz, the pulse parameters determine the drive Hamiltonian in Equation 3. However, gradients are not available as we train the pulse ansatz on NISQ machines. Therefore, we have to use a non-gradient optimizer to train the parameters in the ansatz. To mitigate the drawbacks of non-gradient optimizer, we offer a progressive method to generate our native-pulse ansatz. Such progressive learning structure ensures that the optimizer's parameter dimensions do not exceed its capacity. This machine-in-loop training method also makes the framework noise-resilient. We illustrate the workflow of the proposed framework NAPA in Figure 3. Firstly, the configurations of the NISQ computers are extracted, which include information regarding qubit frequencies and the mapping between native gates and native pulses. After obtaining the native pulses from quantum backends, we can start to progressively construct our native pulse ansatz. Inspired by the hardware-efficient ansatz (HEA), single-qubit gates are placed for each qubit and two-qubit gates are applied to all available connections. The single-qubit pulses are derived from pulses that correspond with the Hadamard gate or Rx gate. And two-qubit pulses are derived from CX gate or CR pulse. CR pulses are preferred because they are the simplest pulses that enable the entanglement of two qubits. Now, we have two types of pulse layers. One consists of single-qubit native pulses on all qubits, whereas the other consists of two-qubit native pulses on available connections. The two types of layers are alternately inserted during the training process to help explore the Hilbert space. To train the native-pulse ansatz, we use non-gradient optimizers. The incrementally constructed ansatz prevents non-gradient optimizers from failing to work with huge dimensions.

Simulators and NISQ machines are employed to evaluate our methods. On the simulators, we present energy-distance curves for several molecules. The energy curves closely resemble those generated with the full configuration interaction (FCI) approach. The results show that our pulse ansatz can approximate the lowest energy states of molecules with much shorter pulse latencies. Our ansatz are also tested on NISQ machines, which have realistic noise and complex system model. In addition, we demonstrate that tuning frequencies on NISQ computers is beneficial.




Fig. 3: The overview of design and implementation of NAPA. The proposed pulse ansatz is composed of single-qubit native pulses (SNP) and cross-resonance pulses (CR). During the training process, the ansatz is “grown” after each step. In this progressive way, older parameters from the last step are unchanged as a “fixed list”, while newer parameters added in this step are updated by the optimizer through optimization iterations as a “partial list”.

VI. DESIGN AND IMPLEMENTATION DETAILS

A. Gate Ansatz versus Native-Pulse Ansatz

In the proposed NAPA, we replace the basic element of the variational quantum circuit with a native pulse and use quantum pulses to build a native-pulse ansatz. During the training process, parameters of native pulses are updated, thus obviating the requirement for decomposition into native gates. To create a native-pulse ansatz, we intentionally make sure that the employed pulses are parametric pulses supported by the NISQ device. Since we use native pulses as building blocks for our ansatz, we can explore more available parameters that are allowed to be tuned. In the case of gate-based ansatz, the parameters are limited to angles of rotation gates. These gates are decomposed into native gates before being implemented with microwave pulses and phase shifts. The microwave pulses are applied to qubits, while the phase shift is influencing on classical electronics. When the angles for the rotation gates change, it changes the phase-shift operations. However, the “redundant” microwave pulses remain fixed. Consequently, gate-based methods cannot take full advantage of parametric pulses. So, we choose to directly adjust the parameters of the microwave pulses that are acting on qubits. In this way, we introduce opportunities to contain more trainable pulses within the same circuit latency. That is, the latency cost for each parametric pulse is reduced, making it possible to generate pulse ansatz with more parameters and less latency at the same time. Another advantage of adjusting pulse parameters instead of gate angles is the mitigation of gate-to-pulse compilation noise. Moreover, for pulse-end cloud users, from scratch to build a pulse ansatz could greatly save the calibration cost which is pretty expensive. To experimentally realize a continuous parametric gate such as $RX(\pi/4)$, the amplitude of $X_{\pi/4}$ pulse is set to be half of that of a $X_{\pi/2}$ pulse, i.e., $A_{\pi/4} = A_{\pi/2}/2$. However, this step introduces noise due to the nonlinearity in NISQ devices [23]. NAPA can directly tune




Fig. 4: Reachable space comparison on Weyl Chamber for a) Single-CR. b) Multi-CRs. While a single-CR pulse can only cover one dimension on the Weyl Chamber, multi-CRs cover one of the surfaces of the Weyl Chamber.

the underlying pulse parameters and avoid such gate-to-pulse compilation noises.

B. Analysis the Power of Native-Pulse Ansatz

In the realm of quantum computation, the precision of operations, especially quantum gates, is quintessential.

Traditional Quantum Rotation Gates: Traditional quantum rotation gates such as RX and RY possess durations of $320dt$, where dt represents a standardized time unit equal to $0.22ns$. Predominantly, the training process for these gates is inclined towards modulating the angle of the quantum rotation gates. This modality implies that the underlying physical pulse remains unaffected, with adjustments predominantly focusing on the classical phase.

Single Qubit Native Pulse: Transitioning our attention to the single qubit native pulse (SNP), its duration is notably $160dt$. Instead of merely modifying the classical phase, one can directly train on the parameters of the physical pulse. This alteration to the pulse parameters allows for an effective




Fig. 5: Cross-resonance pulses have several components other than the intended ZX interaction. While this poses challenges for gate-level calibrations, it doesn't affect parametric pulses.

influence on the drive Hamiltonian, consequently crafting various operations. We propose that harnessing the power of the pulse ansatz in VQAs can transcend the traditional limitations of standard quantum gate operations.

Cross-Resonance: However, one needs to approach this perspective with caution. When deconstructing a quantum circuit, it becomes evident that the influence of a single qubit gate or pulse in terms of the overall circuit duration is relatively minor. The true juggernauts are the two-qubit gates or pulses, which indisputably play the role of a major term in shaping the circuit's dynamics. In the sphere of superconducting quantum machines, the cross-resonance pulses (CR) stand as the minimal two-qubit operation. Given its prominence and efficiency, our choice naturally gravitated towards employing cross resonance as our primary two-qubit operation in the formulation of the native-pulse ansatz.

Through rigorous testing via uniformly random sampling on the Weyl chamber, we derived the gate coverage for both single CR and multi CR as shown in Fig. 4. Remarkably, the gate coverage of the single CR parallels that of the CPhase gate. Moreover, when the angle is set to $\pi/2$, it is locally equivalent to either the CX or CZ gates. This local equivalence is particularly insightful, as gates that are locally equivalent can be interconverted through the addition of specific single-qubit operations. The analysis of multi-CR unveils a more

comprehensive coverage, particularly spanning the entire base of the Weyl chamber. This extensive coverage encompasses a range of standard gates such as the B gate, iSwap, and several others. Interpreting this, when adopting CR as the two-qubit native pulse, we harbor the potential to not only emulate a multitude of standard gates but also venture into the realm of some nonstandard gates. To further elucidate the intricacies of the CR pulse, we provide the equation to describe the effective Hamiltonian of the CR pulse as below:

$$H = Z \otimes A_2 + I \otimes B_2 \\ = a_x \hat{Z} \hat{X} + a_y \hat{Z} \hat{Y} + a_z \hat{Z} \hat{Z} + b_x \hat{I} \hat{X} + b_y \hat{I} \hat{Y} + b_z \hat{I} \hat{Z}. \quad (5)$$

For example, we start with an initial state t_0 which is located at the point I_0 at the weyl chamber in Fig. 4 (a). Now, based on the strength of the Hamiltonian terms we measured from the Hamiltonian tomography experiment with a given pulse duration t_1 , we can calculate the Hamiltonian based on Equation. 5, once we obtain the Hamiltonian, we can further refer to the equation:

$$U = e^{-iHt} \quad (6)$$

to obtain the unitary gate or state evaluation results in matrix representation (here we take t_0 equals 70.4 ns and approximate the result to three decimals):

$$\begin{bmatrix} 0.998 - 0.007j & -0.013 - 0.045j & 0 & 0 \\ 0.013 - 0.045j & 0.998 + 0.007j & 0 & 0 \\ 0 & 0 & 0.894 + 0.026j & 0.009 + 0.447j \\ 0 & 0 & -0.009 + 0.447j & 0.894 - 0.026j \end{bmatrix}$$

this indicates after implement the specific CR pulse with duration t_1 , the Hamiltonian evolve to the point H_1 at Fig. 4. Repeat the process with time evolution t_2 and t_3 , we can observe the evolution of the Hamiltonian on the Weyl chamber in Fig. 4. This inherent flexibility augments the potential of the pulse ansatz in VQAs. By allowing exploration across multiple directions in the Hilbert space, we inherently boost the search capability within the design space.

C. Parameters of Native-Pulse Ansatz

Tuning amplitudes. With native-pulse ansatz, we can tune parameters that are not accessible in gate-based ansatz. For example, the amplitudes of the pulses on the drive and control channels are invisible to gate-level users. Both SNP and CR are initialized with zeros as their amplitudes. The reason is explained in section VI-E.

We can see from Eq. 3 and 4 that when we change the parameters of the pulses, it changes the strength of the signal on the control channel of $d_i(t)$, which eventually changes the drive Hamiltonian. Experimental results presented in Table II confirm the advantages of pulse ansatz over gate ansatz. For experiments in Table II, the target is to solve the ground state energy of H_2 with VQE on the NISQ machine ibmq_jakarta. By comparing the results of SNP, CR, CR+SNP, and the two-gate ansatz, we demonstrate that the pulse ansatz provides a better energy value for the task in 31.2% less duration. The purpose of comparing the two-gate ansatz and the CR+SNP pulse circuit is to highlight the advantages of native-pulse

Operations	Circuit Level	Molecule Bond Length	Reference Energy	VQE (H_2) Result	Duration (on ibmq_jakarta)
SNP	Pulse Circuit	0.1Å	2.710H	4.380H	71.1ns
CR	Pulse Circuit	0.1Å	2.710H	2.927H	163.6ns
SNP	Pulse Circuit	0.75Å	-1.137H	-0.549H	71.1ns
CR	Pulse Circuit	0.75Å	-1.137H	-1.032H	163.6ns
CR + SNP	Pulse Circuit	0.75Å	-1.137H	-1.036H	234.7ns
Two Gate Ansatz	Gate Circuit	0.75Å	-1.137H	-0.534H	341.3ns

TABLE II: Comparison of trainability for different pulse circuits and gate circuits on ibmq_jakarta.




Fig. 6: The circuit is implemented by frequency-tunable pulses. The high probability of getting $|00\rangle$ state for the zero detuning indicates that parametric pulse operations are physically feasible and can be implemented with high fidelity. The changes of final state under different detuning indicate that the variation of pulse parameter indeed physically changes the corresponding quantum operations. Note that here we apply the measurement error mitigation technique [22] which reduces the impact of readout errors.

ansatz construction, as CX and CR are the simplest 2-qubit operations on the gate-level and pulse-level, respectively.

Tuning frequencies. Clifford and T gates are universal to perform arbitrary quantum operations [27], so gate calibrations mainly fine-tune a discrete set of $\pi/2$ and π pulses corresponding to Clifford operations [47]. In the progressive pulse learning protocol, pulse parameters are continuously varied, so it is necessary to verify that parametric pulses still produce physical quantum operations, i.e., pulses correspond to quantum operations and can be implemented with a high fidelity [52]. We demonstrate the performance of pulse implementation by running a pulse version of the gate sequence $(CX + H + H^\dagger + CX^\dagger)$ under different frequency detuning. If the pulse block $(CX + H)$ is physically feasible to realize, when pulse detuning is zero, the circuit should produce a final $|00\rangle$ state with high probability. Experiment results in Figure 6 show a near-to-one probability (00) readout result for the zero detuning case and thus prove the feasibility




Fig. 7: Pulse schedules of “CX+H+H+CX” circuit. D0 and D1 are the drive channels on the first and the second qubits. They control the transmission from input signals to gate operations. U0 is the control channel which provides supplementary control over the qubit to the drive channel. The two qubit basis gate CX is realized through two $\pi/4$ cross-resonance pulses plus one $\pi/2$ X pulse. Single qubit gates are first XY decomposed and then implemented by XY pulses. [37]




Fig. 8: Examples of a policy to grow the native-pulse ansatz. The policy grows the CR and SNP in different steps. CR is designed to induce local entanglement.

of parametric pulse operations. The high-fidelity results also indicate that NAPA does not introduce additional noise and can be authentically implemented.

In addition, the results with frequency detuning verify that the magnitude of the detuning range we choose ($\approx 2MHz$) is wide enough such that the variation of pulse parameters can lead to the obvious changes in the corresponding operation. The frequency could lead to leakage outside of the computational sub-space that may allow the optimizer to take short-cuts for better convergence [21]. In conclusion, frequency detuning with proper range is beneficial for VQAs.




Fig. 9: During the training phase, we can determine which part of the parametric pulse is trained. The complete list of parameters comprises a “fixed list” and a “partial” list. Pruning approaches were provided during training to further simplify the pulse ansatz.

D. Ansatz Construction

VQAs usually consist of parametric circuits with a fixed structure. The major focus in the current VQA research community is to find a way to guide the construction of an ansatz. Traditional ways are mostly physics or chemistry-inspired, which can be cost-inefficient. For instance, to achieve certain accuracy for algorithms with time-evolving blocks such as QAOA, a few trotter steps are sufficient [53]. However, the structured ansatz is not adaptive, and it may be wasteful when we aim to achieve a less demanding accuracy. In comparison, our progressive method generates pulse ansatz adaptively. The corresponding advantage is that the circuit depth is tailored for arbitrary desired accuracy and does not cause any experimental resource overhead.

E. Progressive Pulse Learning

Figure 8 gives an example of a policy to grow the native-pulse ansatz. This policy is inspired by the HEA proposed by [32]. Similar to HEA, single-qubit native pulses are applied to all qubits, while two-qubit native pulses are applied to available connections between qubits. In this way, we create a pulse ansatz that parallels the gate-level HEA structure. This approach allows for a similar arrangement of multiple qubits as in HEA. The difference between NAPA and HEA is whether a progressive policy is adopted. The progressive manner of NAPA can limit the number of parameters held by the optimizer. Figure 9 gives more details on our progressive policy. Importantly, the combination of SNPs and CRs is capable of covering the universal quantum gate set. SNPs and CRs are more essential components of the current basis gates used in quantum computing. For single-qubit operations, any arbitrary rotation can be accomplished using a single native pulse. Similarly, two-qubit operations are implemented through cross-resonance gates. For instance, CNOT gates can be realized using echoed cross-resonance gates. Thus, the combination of SNPs and CRs provides full coverage for the universal quantum gate set. A critical aspect to consider is the larger search space that comes with this approach. The challenge lies in whether we can still identify the desired solutions within this large search space. Experimental results have shown that with a well-designed optimization loop, it is possible to achieve results that surpass those of the gate-level circuits. This suggests that the pulse ansatz can offer sufficient entanglement, expressivity and good coverage of the

universal quantum gate set at the same time. In the process, each new step appends a new group of native pulses, which are subsequently updated through a series of iterations with the optimizer. These iterations only update the parameters within the newly appended pulses. At the same time, the pulses established in previous steps are kept in their current state and remain unchanged. Such progressive manner resembles the idea of **quantum speed limit**, which denotes the maximum rate of evolution of a quantum system, i.e., the minimum time required for a quantum system to evolve between two quantum states [43]. Therefore, the first step of NAPA might fails to reach a good enough quantum state. But it will gradually approximate the desired state with more steps. As for the condition to terminate the progressive growth, we propose to stop growing the ansatz when accuracy cannot benefit from the appended step.

Pruning. We propose to prune the pulse ansatz by removing pulses with parameters that are closest to zeros. The purpose of pruning is to further reduce the overall pulse latency and decoherence error. In addition, controls are less robust when their amplitudes are close to zeros. It is possible that pruning might force NAPA to search in a sub-optimal space, but the progressive way can compensate for possible accuracy degradation. Experimental results show that pruning can significantly reduce the overall latency by an average of 24.13% on small molecule VQE tasks such as H_2 , HeH^+ , and LiH , while accuracies experience minimal-to-none decreases.

Analysis of Scalability of NAPA. The pulse-level optimization algorithm is not compatible with gradient-based methods like parameter-shift rules at the gate level on real quantum machines. That is, we cannot obtain the gradients for pulse parameters from the backends of NISQ machines. We use COBYLA in NAPA, while the selection of non-gradient optimizer remains an open problem. Non-gradient optimizers have two common constraints: they can only handle a limited number of parameters and they do not guarantee convergence to the global optimum. Therefore, we propose training the ansatz in a “progressive” manner, which helps to limit the number of parameters being trained at the same time. Pulse-level ansatz has more parameters than gate-level ansatz but with a much shorter duration, which means pulse ansatz costs less on quantum resources but more on classical computational resources, the trade-off is worth especially in the NISQ era since the decoherence time is limited. A large number of parameters is the intrinsic problem for VQE algorithms, which is not the focus of this work, but we will think in-depth about some good solutions to this problem.

We summarize existing state of art work compared with NAPA in Table III. “Ctrl-VQE” [38], where pulse segments are built and the parameters are trained all together. When “Ctrl-VQE” is proposed, it only considers single-qubit pulse channel. Information of pulse phases and two-qubit pulses are also missing in the “Ctrl-VQE” framework, making it infeasible on real quantum devices. Considerations of two-qubit pulse channel and adaptivity of ansatz size provide NAPA with essentials advantages over “Ctrl-VQE”.

The scalability of parameterized pulses is jointly determined by two factors: the number of parameters and the perfor-

Model	Single Qubit Channel	Two Qubit Channel	Training Framework	Compatible with Real Machine
NAPA	Yes	Yes	Progressive	Yes
Ctrl-VQE [38]	Yes	No	Brute-Force	No

TABLE III: “Ctrl-VQE” only considers single-qubit pulse channel. Two-qubit pulses are missing in the “Ctrl-VQE” framework, making it inferior for certain tasks. Brute-force optimization makes “Ctrl-VQE” limited to small-size problems.

mance of the optimizer. Firstly, the number of parameters in parameterized pulses scales linearly with the number of pulses. This scaling is the same as the number of parameters in parameterized gates. One advantage of parameterized pulses is their capacity to integrate a larger number of parameters within the same circuit latency. However, it’s not always necessary to utilize all available parameters during the optimization process. The flexibility in the number of parameters means we can control the parameter count in pulses. This adaptability allows for different options: we can opt for a higher degree of parameterization to gain more flexibility or we can choose a more conservative approach, aligning the number of parameters with those in gate-level circuits to merely reduce circuit latency.

In gate-level VQE tasks, gradient-based optimizers are available with gradients derived through parameter-shift rules. At the pulse level, the concepts of gradient-based optimizers have recently been introduced [41]. To the best of our knowledge, these gradient-based pulse optimizers currently face to challenges to implement on real quantum devices. We would like to point out that the methods in NAPA are orthogonal to the choice of optimizers. NAPA can be implemented with gradient-based optimizers since it effectively mitigates the problem of high-dimensional parameter space. In conclusion, techniques in NAPA should work for both gradient-based and non-gradient optimizers.

VII. EVALUATION

A. Backend configuration

Our experiments are conducted both on simulators and NISQ machines. Simulators are used to validate that NAPA considers the system models of quantum backends. Simulations are run on a server with two Intel Xeon E5-2630 CPUs (8 cores/CPU), 64 GB DRAM, with CentOS 7.4 as the operating system. To confirm that our framework works on NISQ machines, we conduct experiments on six IBM’s quantum systems: *ibmq_cairo*, *ibmq_montreal*, *ibmq_toronto*, *ibmq_mumbai*, *ibmq_guadalupe* and *ibmq_jakarta*.

B. Compilation Setting

To illustrate our compilation settings and overhead, we show an example task of H_2 in the Table IV. In our progressive method, each step contains up to 50 iterations. In each iteration, quantum programs are measured with 1024 shots. The complexity of the task determines the number of steps that are needed. We also give the range of amplitudes and frequencies in the table. The hyperparameter Rhobeg of the optimizer COBYLA is set to 0.1.

Compilation Settings	Values
#Parameters	6*
#Measurement per iteration	1024
#Iteration per step	50
#Steps per task	2*
Range of Amplitudes	0 - 0.4
Range of Frequencies	-2MHz - 2MHz
Rhobeg	0.1

TABLE IV: Example of compilation settings for VQE task of the H_2 . * means the parameter is dependent on the task. Progressive learning framework determines the number of steps with respect to task definition.

C. Experiment Setup

We choose VQE problems as our evaluation tasks, which consist of several molecules including H_2 , HeH^+ , LiH , CO_2 , H_2O , and NaH , and a three-regular six-nodes graph Maxcut task. For molecular tasks, the ‘STO-3G’ basis set is employed to approximate the spin-orbitals of molecules, which are then transformed into fermionic terms. These terms undergo the Jordan-Wigner (JW) transformation to yield Pauli strings. VQE is then utilized to estimate the ground state energy based on these Pauli strings. In the case of the Maxcut problem, the problem is initially reformulated as a quadratic program, from which the corresponding Ising Hamiltonian is derived. This allows a VQE approach to solve the Maxcut problem using the deduced Ising Hamiltonian.

For the Maxcut problem, the primary metric of interest is the approximation ratio. It’s the value achieved by the VQE solution divided by the value of the objective function for the optimal solution. A higher approximation ratio indicates a closer approximation to the optimal solution, and thus better performance of the algorithm. In the case of molecule ground state energy estimation, the key metric is accuracy. This is typically assessed by comparing the energy computed by the VQE algorithm with the Full Configuration Interaction (FCI) results. FCI is a highly accurate quantum chemistry method that calculates the exact ground state energy of a quantum system within a given basis set. The closer the estimation result is to the FCI benchmark, the more accurate the VQE algorithm is considered to be for molecule tasks.

D. Native pulse

In the Table II, we compare the results of different ansatz types for VQE tasks. The results are collected from NISQ machines. To determine the capabilities of our native-pulse model, we employ the simplest native-pulse model and compare its performance to that of native gates. The Table II

Model		Cairo	Montreal	Toronto	NISQ machine Avg	Simulator	FCI
H_2	Step I	-1.093 (3.870%)	-1.087 (4.398%)	-1.073 (5.629%)	-1.084 (4.661%)	-1.121 (1.407%)	-1.137
	Step II	-1.107 (2.639%)	-1.110 (2.375%)	-1.073 (5.629%)	-1.097 (3.518%)	-1.123 (1.231%)	-1.137
	Inaccuracy Reduction	31.83%	46.00%	0.000%	24.52%	12.51%	-
$HeH+$	Step I	-2.813 (1.746%)	-2.845 (0.663%)	-2.820 (1.485%)	-2.826 (1.292%)	-2.855 (0.279%)	-2.863
	Step II	-2.833 (1.047%)	-2.866 (0.105%)	-2.834 (1.013%)	-2.844 (0.664%)	-2.856 (0.244%)	-2.863
	Inaccuracy Reduction	40.03%	84.16%	31.78%	48.61%	12.54%	-

TABLE V: Results of Estimated Energy for Molecules in Different Steps.

Model	Ansatz Level	Qubits	Duration	SNP Count	CR Count	Molecule	Energy	Reference Energy
Random Generated Ansatz	Gate Ansatz	2	682.7ns	16	2	H_2	-0.853	-1.137
RealAmplitude Ansatz [2]	Gate Ansatz	2	376.9ns	12	1	H_2	-0.974	-1.137
QuantumNAS [51]	Gate Ansatz	2	682.7ns	16	2	H_2	-1.033	-1.137
NAPA	Pulse Ansatz	2	71.1ns	3	0	H_2	-1.100	-1.137
RealAmplitude Ansatz	Gate Ansatz	2	753.8ns	24	2	$HeH+$	-2.691	-2.863
NAPA	Pulse Ansatz	2	199.1ns	1	1	$HeH+$	-2.866	-2.863
QuantumNAS	Gate Ansatz	6	7296.0ns	40	12	LiH	-6.914	-7.882
NAPA	Pulse Ansatz	4	199.1ns	4	2	LiH	-7.590	-7.882

TABLE VI: Comparison of duration, pulse counts, and estimated energy of gate ansatz and the native-pulse ansatz generated by NAPA on NISQ machines. SNP indicates Single-Qubit Native Pulse, and CR indicates Cross-Resonance Pulse.

demonstrates that our native-pulse approach can deliver superior VQE results in less durations. Pulse ansatz with only two native pulses can produce $-1.036H$ energy, whereas the two-gate ansatz can only achieve $-0.534H$ energy. The two-gate ansatz consists of both single-qubit and double-qubit gates. The duration of the pulse ansatz is approximately 30% less than that of a gate ansatz. The duration numbers are derived from the pulses that are returned by quantum backends. In conclusion, we are able to produce superior VQE results with significantly less latency on NISQ machines.

E. Simulation results

To confirm that our framework can generate accurate energy curves for given molecules, we use simulators to evaluate our methodologies. The specifications of “fake” backends are collected by the simulators and used as Hamiltonian configurations to run the pulses. We use pulse simulator that is provided by IBM qiskit toolkit. It simulates continuous time Hamiltonian dynamics of a quantum system, with controls specified by pulses. The simulation results in Figure 10 indicate that our methods can produce comparable results to those of the FCI methods. When attempting to compute the ground state energy of $HeH+$ molecules, we discover that the simulated results are only 0.244% off from the FCI results. In the case of H_2 molecules, our simulated results deviate from FCI data by 1.23%. As shown in Table V, the insertion of a native pulse in STEP two improved the performance of the native-pulse ansatz by approximately 12% on simulators. Overall, simulation results validates the methodology of NAPA.

F. NISQ machine results

Our techniques are also evaluated on several NISQ computers. As presented in Figure 11 and Table VI, a toy model of




Fig. 10: Simulation results of $HeH+$ molecule. The pulse ansatz curve almost matches the FCI value. The lowest energy point reaches an accuracy of 99.756%.

native-pulse ansatz is capable of producing promising results. In the instance of the $HeH+$ molecule, we attain an average accuracy of 99.336%, with 99.895% being the highest achievable accuracy. The absolute difference in energy is 0.003H. The accuracy is close to the requirement of computational chemistry (0.0016H), which is the chemical accuracy constant for computational chemistry. It qualifies the typical minimum energy gap that can be verified through experiments. As for the H_2 molecule, we attain an average accuracy of 96.482% and a maximum accuracy of 97.625% with native-pulse ansatz. For LiH , a bigger molecule, we obtain an energy accuracy of 96.295%. These accuracy figures are derived from NISQ machines with gate and measurement errors that exceed 1%. We observe that *ibmq_montreal* tends to return the best results,




Fig. 11: Evaluation of native-pulse ansatz on NISQ machines for H_2 , HeH^+ , and LiH VQE tasks. NISQ machines include *ibmq_cairo*, *ibmq_montreal*, *ibmq_toronto*, *ibmq_mumbai*, *ibmq_guadalupe*, and *ibmq_jakarta*. Our toy ansatz can achieve great accuracy on all the NISQ machines. For *ibmq_montreal* where we obtain the best results, the average CX error is 1.518% as the average readout error is 3.457%. NAPA is proven to be robust and error-resilient on NISQ machines.




Fig. 12: Influence of frequency shift on the estimated energies. Even with a narrow frequency shift range, we notice obvious changes in the estimated energies. The experimental results show that for the H_2 molecule's ground state energy estimation task, as we vary the detuning frequency, the energy estimation scans over a broad range, which validates our methods to tune frequency in pulses.

while *ibmq_mumbai* typically returns the worse results. Then, we confirm that *ibmq_mumbai* has larger error rates for gates and measurements on average than *ibmq_montreal*.

The results collected from NISQ computers demonstrate that our approaches are highly error-tolerant. In terms of total duration time, our native-pulse approaches have significant advantages over the current gate ansatz generator. To ensure equality, the gate ansatz of the baselines is implemented on the same NISQ computer, and their duration is determined with the acquired pulse schedules. We are able to demonstrate a 97.3% reduction in ansatz duration compared to QuantumNAS, and our estimated energy numbers for LiH are lower. With a duration decrease of 89.6%, our method can obtain better energy values while dealing with H2 molecules. QuantumNAS does not report the energy numbers for the HeH+ molecule. Thus we compare our techniques to the Real Amplitude Ansatz, which shows that we are able to reduce duration by 73.6% while maintaining similar performance.

In Fig. 12, we detune the pulse frequency and use a few steps to estimate the energy of the H2 molecule. The experimental results show that as we vary the detuning, the energy estimation scans over a broad range, which ensures that the pulse parameter variation in the learning step is wide

enough to cover the ground energy. Moreover, the evident response of energy estimation with pulse parameterization shows that our ansatz structures are valid in covering all spin-orbital configurations of the electronic structure. The estimated energy can then progressively approximate the desired ground energy through pulse parameter learning.

Figure 13 depicts the relationships between the computed energy and the number of iterations. Despite the fact that it is not visible in Figure 13, we validate that our progressive approaches work on VQE tasks. Since a non-gradient optimizer is deployed, we can observe several peaks of the curve where the non-gradient optimizer attempts to update the pulse ansatz's parameters but obtains poorer results. As indicated in the Table V, we are able to reduce the deviation from FCI values by around 40%, when native-pulses "grow" progressively. Only in one instance where the pulse ansatz were conducted on *ibmq_toronto* were the results not improved. Taking into account this failure, we find a 24% improvement for H2 molecule experiments and a 48% improvement for HeH+ molecule experiments. We also evaluate NAPA on larger chemistry tasks as well as optimization tasks as shown in Fig. 14. For CO_2 , H_2O , and NaH , the pulse ansatz generated by NAPA has a duration of 1031.1 ns on




Fig. 13: Energy trends with #iterations for a H_2 VQE task. The data is collected from *ibmq_montreal*. NAPA with non-gradient optimizer for pulse-level optimization. The vertical dotted line in the middle separates the steps in progressive learning. We can see several “peaks” on the curves because non-gradient optimizer attempts might be made in a bad “direction”. The accuracy achieved in Step I, encompassing optimization iterations one through 26, is 95.974%. Subsequently, Step II, spanning iterations 27 to 76, further enhances the accuracy to 96.746%.




Fig. 14: Evaluation of native-pulse ansatz on NISQ machines for CO_2 , H_2O , and NaH 's VQE tasks as well as three-regular six-nodes graph maxcut task. NISQ machines include *ibmq_montreal* and *ibmq_jakarta*. “Gate level” in the Maxcut task indicates the “TwoLocal” ansatz from Qiskit circuit library which consists of Ry and CZ gates.

ibmq_montreal and the estimated ground state energies are -183.856H (99.70% accuracy), -73.574H (97.96% accuracy), and -150.984H (94.15% accuracy), respectively. For a three-regular six-node graph, the experiment is conducted on NISQ machine *ibmq_jakarta*, NAPA produces the pulse ansatz with resulting improved approximation ratio of 3.4% compared to a gate-level ansatz of similar structure. We use the “TwoLocal” ansatz which is a type of HEA that consists of Ry and CZ gates as our gate-level baseline. Results demonstrate the feasibility of NAPA on different real-world applications.

VIII. CONCLUSION

NAPA is a framework for constructing native-pulse ansatz for VQAs. As a result of removing an abstraction layer of native gates, native pulses provide huge latency advantages.

Then, we employ progressive learning to “grow” our ansatz. Thus, our pulse ansatz is better able to explore the Hilbert space, while the optimizer is still able to handle the problem. Extensive experiments are conducted on NISQ machines, and the results indicate that an average of 86% reduction in latency is obtained with up to 99.895% accuracy on small molecule VQE tasks. Experiments on larger-size molecule VQE tasks achieve an average accuracy of 97.27%. The results of NAPA on maxcut problem demonstrate the feasibility of NAPA on different problems. How to explore the applications of NAPA on different types of VQAs remains an open question. Also, potentials of NAPA on quantum algorithms other than VQAs are to be investigated. Overall, our experimental results show that pulse-level methods like NAPA can greatly reduce the duration of quantum circuits, thus enhance the capabilities of quantum computers.

IX. ACKNOWLEDGMENTS

This work is funded in part by EPiQC, an NSF Expedition in Computing, under award CCF-1730449; in part by STAQ under award NSF Phy-1818914/232580; in part by the US Department of Energy Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing Program; and in part by the NSF Quantum Leap Challenge Institute for Hybrid Quantum Architectures and Networks (NSF Award 2016136), in part based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, and in part by the Army Research Office under Grant Number W911NF-23-1-0077. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

FTC is the Chief Scientist for Quantum Software at Infleqtion and an advisor to Quantum Circuits, Inc.

REFERENCES

- [1] “Ibm debuts next-generation quantum processor & ibm quantum system two, extends roadmap to advance era of quantum utility,” <https://newsroom.ibm.com/2023-12-04-IBM-Debuts-Next-Generation-Quantum-Processor-IBM-Quantum-System-Two,-Extends-Roadmap-to-Advance-Era-of-Quantum-Utility>.
- [2] “Qiskit,” <https://qiskit.org/>.
- [3] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, “The power of quantum neural networks,” *Nature Computational Science*, vol. 1, no. 6, pp. 403–409, 2021.
- [4] D. Aharonov and M. Ben-Or, “Fault-tolerant quantum computation with constant error rate,” *SIAM Journal on Computing*, 2008.
- [5] T. Alexander, N. Kanazawa, D. J. Egger, L. Capelluto, C. J. Wood, A. Javadi-Abhari, and D. C. McKay, “Qiskit pulse: Programming quantum computers through the cloud with pulses,” *Quantum Science and Technology*, vol. 5, no. 4, p. 044006, 2020.
- [6] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell et al., “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [7] S. Axelrod, E. Shakhnovich, and R. Gómez-Bombarelli, “Excited state non-adiabatic dynamics of large photoswitchable molecules using a chemically transferable machine learning potential,” *Nature communications*, vol. 13, no. 1, pp. 1–11, 2022.

- [8] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [9] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter et al., "Logical quantum processor based on reconfigurable atom arrays," *Nature*, pp. 1–3, 2023.
- [10] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, "Trapped-ion quantum computing: Progress and challenges," *Applied Physics Reviews*, vol. 6, no. 2, p. 021314, 2019.
- [11] T. Caneva, T. Calarco, and S. Montangero, "Chopped random-basis quantum optimization," *Physical Review A*, vol. 84, no. 2, p. 022326, 2011.
- [12] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. Sawaya et al., "Quantum chemistry in the age of quantum computing," *Chemical reviews*, vol. 119, no. 19, pp. 10856–10915, 2019.
- [13] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio et al., "Variational quantum algorithms," *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, 2021.
- [14] J. Cheng, H. Deng, and X. Qia, "Accqoc: Accelerating quantum optimal control based pulse generation," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2020, pp. 543–555.
- [15] A. Choquette, A. Di Paolo, P. K. Barkoutsos, D. Sénéchal, I. Tavernelli, and A. Blais, "Quantum-optimal-control-inspired ansatz for variational quantum algorithms," *Physical Review Research*, vol. 3, no. 2, p. 023092, 2021.
- [16] P. Das, C. A. Pattison, S. Manne, D. M. Carmean, K. M. Svore, M. Qureshi, and N. Delfosse, "Afs: Accurate, fast, and scalable error-decoding for fault-tolerant quantum computers," in *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 2022, pp. 259–273.
- [17] P. de Fouquieres, S. G. Schirmer, S. J. Glaser, and I. Kuprov, "Second order gradient ascent pulse engineering," *Journal of Magnetic Resonance*, vol. 212, no. 2, pp. 412–417, 2011.
- [18] R. de Keijzer, O. Tse, and S. Kokkelmans, "Pulse based variational quantum optimal control for hybrid quantum computing," *arXiv preprint arXiv:2202.08908*, 2022.
- [19] Y. Ding and F. T. Chong, "Quantum computer systems: Research for noisy intermediate-scale quantum computers," *Synthesis Lectures on Computer Architecture*, vol. 15, no. 2, pp. 1–227, 2020.
- [20] Y. Du, T. Huang, S. You, M.-H. Hsieh, and D. Tao, "Quantum circuit architecture search for variational quantum algorithms," *npj Quantum Information*, vol. 8, no. 1, pp. 1–8, 2022.
- [21] D. J. Egger, C. Capecci, B. Pokharel, P. K. Barkoutsos, L. E. Fischer, L. Guidoni, and I. Tavernelli, "A study of the pulse-based variational quantum eigensolver on cross-resonance based hardware," *arXiv preprint arXiv:2303.02410*, 2023.
- [22] L. Funcke, T. Hartung, K. Jansen, S. Kühn, P. Stornati, and X. Wang, "Measurement error mitigation in quantum computers through classical bit-flip correction," 2020. [Online]. Available: <https://arxiv.org/abs/2007.03663>
- [23] J. Gil-Lopez, M. Santandrea, G. Roland, B. Brecht, C. Eigner, R. Ricken, V. Quiring, and C. Silberhorn, "Improved non-linear devices for quantum applications," *New Journal of Physics*, vol. 23, no. 6, p. 063082, jun 2021. [Online]. Available: <https://doi.org/10.1088%2F1367-2630%2F20063082>
- [24] P. Gokhale, O. Angiuli, Y. Ding, K. Gui, T. Tomesh, M. Suchara, M. Martonosi, and F. T. Chong, "Optimization of simultaneous measurement for variational quantum eigensolver applications," in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 2020, pp. 379–390.
- [25] P. Gokhale, Y. Ding, T. Propson, C. Winkler, N. Leung, Y. Shi, D. I. Schuster, H. Hoffmann, and F. T. Chong, "Partial compilation of variational algorithms for noisy intermediate-scale quantum machines," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2019, pp. 266–278.
- [26] P. Gokhale, A. Javadi-Abhari, N. Earnest, Y. Shi, and F. T. Chong, "Optimized quantum compilation for near-term algorithms with open-pulse," in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2020, pp. 186–200.
- [27] D. Gottesman, "Theory of fault-tolerant quantum computation," *Phys. Rev. A*, vol. 57, pp. 127–137, Jan 1998. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.57.127>
- [28] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, 1996, pp. 212–219.
- [29] G. G. Guerreschi and A. Y. Matsuura, "Qaoa for max-cut requires hundreds of qubits for quantum speed-up," *Scientific reports*, vol. 9, no. 1, pp. 1–7, 2019.
- [30] J. Herrmann, S. M. Lima, A. Remm, P. Zapletal, N. A. McMahon, C. Scarato, F. Swiadek, C. K. Andersen, C. Hellings, S. Krinner et al., "Realizing quantum convolutional neural networks on a superconducting quantum processor to recognize quantum phases," *Nature Communications*, vol. 13, no. 1, pp. 1–7, 2022.
- [31] J. R. Johansson, P. D. Nation, and F. Nori, "Qutip: An open-source python framework for the dynamics of open quantum systems," *Computer Physics Communications*, 2012.
- [32] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, 2017.
- [33] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, "Linear optical quantum computing with photonic qubits," *Reviews of modern physics*, vol. 79, no. 1, p. 135, 2007.
- [34] Z. Liang, H. Wang, J. Cheng, Y. Ding, H. Ren, X. Qian, S. Han, W. Jiang, and Y. Shi, "Variational quantum pulse learning," *arXiv preprint arXiv:2203.17267*, 2022.
- [35] M. Lubasch, J. Joo, P. Moinier, M. Kiffner, and D. Jaksch, "Variational quantum algorithms for nonlinear problems," *Physical Review A*, vol. 101, no. 1, p. 010301, 2020.
- [36] D. C. McKay, T. Alexander, L. Bello, M. J. Biercuk, L. Bishop, J. Chen, J. M. Chow, A. D. Córcoles, D. Egger, S. Filipp et al., "Qiskit backend specifications for openqasm and openpulse experiments," *arXiv preprint arXiv:1809.03452*, 2018.
- [37] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, "Efficient z gates for quantum computing," *Phys. Rev. A*, vol. 96, p. 022330, Aug 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.96.022330>
- [38] O. R. Meitei, B. T. Gard, G. S. Barron, D. P. Pappas, S. E. Economou, E. Barnes, and N. J. Mayhall, "Gate-free state preparation for fast variational quantum eigensolver simulations," *npj Quantum Information*, vol. 7, no. 1, pp. 1–11, 2021.
- [39] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn et al., "Quantum optimization using variational algorithms on near-term quantum devices," *Quantum Science and Technology*, vol. 3, no. 3, p. 030503, 2018.
- [40] P. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding et al., "Scalable quantum simulation of molecular energies," *Physical Review X*, vol. 6, no. 3, p. 031007, 2016.
- [41] Y. Peng, J. Young, P. Liu, and X. Wu, "Simuq: A domain-specific language for quantum simulation with analog compilation," *arXiv preprint arXiv:2303.02775*, 2023.
- [42] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, "A variational eigenvalue solver on a photonic quantum processor," *Nature communications*, vol. 5, no. 1, pp. 1–7, 2014.
- [43] D. P. Pires, M. Cianciaruso, L. C. Céleri, G. Adesso, and D. O. Soares-Pinto, "Generalized geometric quantum speed limits," *Physical Review X*, vol. 6, no. 2, p. 021031, 2016.
- [44] A. P. Place, L. V. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman et al., "New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds," *Nature communications*, vol. 12, no. 1, pp. 1–6, 2021.
- [45] Y. Shi, N. Leung, P. Gokhale, Z. Rossi, D. I. Schuster, H. Hoffmann, and F. T. Chong, "Optimized compilation of aggregated instructions for realistic quantum computers," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2019, pp. 1031–1044.
- [46] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.
- [47] L. S. Theis, F. Motzoi, S. Machnes, and F. K. Wilhelm, "Counteracting systems of diabaticities using DRAG controls: The status after 10 years," *EPL (Europhysics Letters)*, vol. 123, no. 6, p. 60001, oct 2018. [Online]. Available: <https://doi.org/10.1209%2F0295-5075%2F123%2F60001>
- [48] T. Tomesh, P. Gokhale, V. Omole, G. S. Ravi, K. N. Smith, J. Viszlai, X.-C. Wu, N. Hardavellas, M. R. Martonosi, and F. T. Chong, "Supermarq: A scalable quantum benchmark suite," in *2022 IEEE*

International Symposium on High-Performance Computer Architecture (HPCA). IEEE, 2022, pp. 587–603.


- [49] T. Tomesh, K. Gui, P. Gokhale, Y. Shi, F. T. Chong, M. Martonosi, and M. Suchara, "Optimized quantum program execution ordering to mitigate errors in simulations of quantum systems," in 2021 International Conference on Rebooting Computing (ICRC). IEEE, 2021, pp. 1–13.
- [50] D. Wang, O. Higgott, and S. Brierley, "Accelerated variational quantum eigensolver," *Physical review letters*, vol. 122, no. 14, p. 140504, 2019.
- [51] H. Wang, Y. Ding, J. Gu, Y. Lin, D. Z. Pan, F. T. Chong, and S. Han, "Quantumnas: Noise-adaptive search for robust quantum circuits," in 2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA). IEEE, 2022, pp. 692–708.
- [52] C. Weedbrook, S. Pirandola, R. García-Patrón, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd, "Gaussian quantum information," *Rev. Mod. Phys.*, vol. 84, pp. 621–669, May 2012. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.84.621>
- [53] X. Yang, X. Nie, Y. Ji, T. Xin, D. Lu, and J. Li, "Improved quantum computing with the higher-order trotter decomposition," 2022. [Online]. Available: <https://arxiv.org/abs/2205.02520>



Zhiding Liang is a third-year Ph.D. student at the University of Notre Dame, Department of Computer Science and Engineering, advised by Prof. Yiyu Shi. His research focuses on quantum pulse control, hardware/software codesign for quantum computing, and VQAs. His research has been recognized as the DAC Young Fellow twice, both in 2020 and 2021, and highlighted by IBM Qiskit, and appears in top conferences such as DAC, ICCAD, and QCE. He is the recipient of the Edison Innovation Fellowship.




Jinglei Cheng is a final-year Ph.D. student at Purdue University, Department of Computer Science. He is advised by Prof. Xuehai Qian in the ALCHEM Group. He received Bachelor of Engineering Degree in the Department of Microelectronic Science and Engineering at Tsinghua University in 2018. His research focuses on VQAs and quantum computer architecture. And his research papers appear in conferences including MICRO, ISCA, DAC and QCE.



Hang Ren is currently a Ph.D. student at the University of California, Berkeley. He received Bachelor's degree from Nankai University. His research is about quantum error mitigation and shadow tomography.



Hanrui Wang Hanrui Wang is a final-year Ph.D. student at MIT EECS advised by Prof. Song Han. His research focuses on quantum computer architecture, ML for quantum and efficient ML. His research has been recognized by ACM student research competition 1st place award, best poster award at NSF AI Institute and appears in top conferences such as MICRO, HPCA, DAC, ICCAD and NeurIPS. He is the recipient of Qualcomm Fellowship, Unitary Fund, and Nvidia Fellowship Finalist.



Fei Hua is a 6th-year Ph.D. at the Computer Science Department at Rutgers. His major interests are quantum program compilation, quantum error mitigation and correction. He got his M.S. degree in Computer Science from Washington University at St. Louis. He did his internships at IBM Research in 2022, and at Pacific Northwest National Lab (PNNL) in 2023. He is one of the lead developers for the open-source compiler QASM Transpiler.



Zhixin Song is currently a Ph.D. student at the Georgia Institute of Technology. He received Bachelor's degree from the Ohio State University in 2019. His research interests include near-term quantum algorithms for solving differential equations and computational fluid dynamics (CFD), cooperating HPC systems and quantum hardware.



Yongshan Ding is an Assistant Professor of Computer Science at Yale University. Ding completed his Ph.D. from the University of Chicago. He is a recipient of the William Rainey Harper Dissertation Fellowship, one of UChicago's highest honors, and the Siebel Scholarship. Before that, he received his B.Sc. degrees in Computer Science and Physics from Carnegie Mellon University.




Frederic T. Chong (IEEE Fellow) is the Seymour Goodman Professor in the Department of Computer Science at the University of Chicago. He is also Lead Principal Investigator for the EPiQC Project, an NSF Expedition in Computing. Chong received his Ph.D. from MIT in 1996 and was a faculty member and Chancellor's fellow at UC Davis from 1997-2005. He was a Professor of Computer Science, Director of Computer Engineering, and Director of the Greenscale Center for Energy-Efficient Computing at UCSB from 2005-2015.



Song Han is an associate professor at MIT EECS. He received his PhD degree from Stanford University. He proposed the "Deep Compression" technique including pruning and quantization that is widely used for efficient AI computing, and "Efficient Inference Engine" that first brought weight sparsity to modern AI chips, which influenced NVIDIA's Ampere GPU Architecture with Sparse Tensor Core. He pioneered the TinyML research that brings deep learning to IoT devices.



Xuehai Qian is an associate professor of the Department of Computer Science at Purdue University. His research interests include domain-specific systems and architectures, performance tuning and resource management of cloud systems, and parallel computer architectures. He got his Ph.D from the University of Illinois Urbana Champaign. He is the recipient of W.J Poppelbaum Memorial Award at UIUC, NSF CRII and CAREER Award, and the inaugural AC-SIC Rising Star Award.



Yiyu Shi is currently a professor in the Department of Computer Science and Engineering at the University of Notre Dame, the site director of National Science Foundation IUCRC Alternative and Sustainable Intelligent Computing, and the director of the Sustainable Computing Lab (SCL). He received his B.S. in Electronic Engineering from Tsinghua University, Beijing, China in 2005, the M.S and Ph.D. degree in Electrical Engineering from the University of California, Los Angeles in 2007 and 2009 respectively. His current research interests focus on hardware intelligence and biomedical applications.