



Navigating the Dynamic Noise Landscape of Variational Quantum Algorithms with QISMET

Gokul Subramanian Ravi*
University of Chicago, USA

Kaitlin Smith
University of Chicago, USA

Jonathan M. Baker
University of Chicago, USA

Tejas Kannan
University of Chicago, USA

Nathan Earnest
IBM Quantum, USA

Ali Javadi-Abhari
IBM Quantum, USA

Henry Hoffmann
University of Chicago, USA

Frederic T. Chong
University of Chicago, USA

ABSTRACT

In the Noisy Intermediate Scale Quantum (NISQ) era, the dynamic nature of quantum systems causes noise sources to constantly vary over time. Transient errors from the dynamic NISQ noise landscape are challenging to comprehend and are especially detrimental to classes of applications that are iterative and/or long-running, and therefore their timely mitigation is important for quantum advantage in real-world applications.

The most popular examples of iterative long-running quantum applications are variational quantum algorithms (VQAs). Iteratively, VQA's classical optimizer evaluates circuit candidates on an objective function and picks the best circuits towards achieving the application's target. Noise fluctuation can cause a significant transient impact on the objective function estimation of the VQA iterations' tuning candidates. This can severely affect VQA tuning and, by extension, its accuracy and convergence.

This paper proposes QISMET: Quantum Iteration Skipping to Mitigate Error Transients, to navigate the dynamic noise landscape of VQAs. QISMET actively avoids instances of high fluctuating noise which are predicted to have a significant transient error impact on specific VQA iterations. To achieve this, QISMET estimates transient error in VQA iterations and designs a controller to keep the VQA tuning faithful to the transient-free scenario. By doing so, QISMET efficiently mitigates a large portion of the transient noise impact on VQAs and is able to improve the fidelity by 1.3x-3x over a traditional VQA baseline, with 1.6-2.4x improvement over alternative approaches, across different applications and machines.

CCS CONCEPTS

• Computer systems organization → Quantum computing.

*Correspondence: gravi@uchicago.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASPLOS '23, March 25–29, 2023, Vancouver, BC, Canada

© 2023 Association for Computing Machinery.

ACM ISBN 978-1-4503-9916-6/23/03...\$15.00

<https://doi.org/10.1145/3575693.3575739>

KEYWORDS

quantum computing, variational quantum algorithms, error mitigation, transient error, superconducting qubits, noisy intermediate-scale quantum, variational quantum eigensolver

ACM Reference Format:

Gokul Subramanian Ravi, Kaitlin Smith, Jonathan M. Baker, Tejas Kannan, Nathan Earnest, Ali Javadi-Abhari, Henry Hoffmann, and Frederic T. Chong. 2023. Navigating the Dynamic Noise Landscape of Variational Quantum Algorithms with QISMET. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '23)*, March 25–29, 2023, Vancouver, BC, Canada. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3575693.3575739>

1 INTRODUCTION

Quantum computers leverage superposition, interference, and entanglement to give them significant computing advantage in chemistry [17], optimization [26], machine learning [6] and other domains of critical interest. In near-term quantum computing, called Noisy Intermediate-Scale Quantum (NISQ), we expect to work with machines which comprise 100-1000s of imperfect qubits [33].

Today, noise prevents quantum computers from surpassing the capabilities of classical computers in almost all applications. NISQ devices suffer from high error rates in the form of state preparation and measurement (SPAM) errors, gate errors, qubit decoherence, crosstalk, etc. These errors stem from multiple noise sources such as the imperfect classical control of the device, thermal fluctuations, destructive qubit coupling, imperfect insulation of the qubits, quasi-particles, and other external stimuli [9, 19, 24, 27, 36]. The dynamic nature of quantum systems causes these noise sources to constantly vary over time, with some noise sources experiencing greater transient fluctuations than others. Additionally, because current fabrication techniques lack the precision to make homogeneous quantum device batches, the transient noise properties are unique to each device.

Effectively understanding and mitigating different forms of error is critical for quantum advantage in real-world applications. Multiple error mitigation techniques have been explored on NISQ devices in the recent past [8, 10, 13, 22, 28, 29, 31, 37–41, 44]. While these techniques have the potential to greatly improve execution fidelity, they almost always view machine noise as static, or at least expect the noise to be stable for sufficiently long periods of time such that their characteristics can be adequately captured, appropriately characterized, and then effectively mitigated. Unfortunately, this is




Figure 1: Optimal VQA convergence is obtained in the ideal noise-free scenario (orange). In reality, VQA is affected by static and transient noise, and estimates can be much worse than ideal (red). QISMET (green) attempts to avoid significant transient error and thereby reaches close to the otherwise unrealistic blue line with only static noise. Prior error mitigation proposals can help bring the blue (and green) closer to the orange.

often insufficient—in this work we show that the transient errors from the dynamic noise landscape are challenging to comprehend and are especially detrimental to classes of applications that are iterative and/or long-running, and therefore their timely mitigation is important for significant fidelity improvements. Here, we use the term ‘transient error’ to refer to a temporary change in a quantum circuit’s output distribution, produced due to sudden non-insignificant shift in device characteristics of one or more qubits in the quantum circuit [9, 19, 36]. The effect of these transient errors is seen even with the statistical robustness offered by executing multiple circuit shots (which are required to capture probabilistic output distributions).

The most popular examples of iterative long-running quantum applications are variational quantum algorithms (VQAs). VQAs are considered one of the most promising quantum tasks for near-term quantum advantage in the NISQ era. They have wide application in approximation [26], chemistry [30] etc, which are usually designed as minimization problems. VQAs are hybrid quantum-classical algorithms and they iteratively run a parameterized quantum circuit (QC) on the quantum machine. The QC parameters are optimized each iteration by a classical tuner/optimizer to try and approach the global minimum of the variational objective function. The classical optimizer inherently attempts to adjust the QC to the noise characteristics of the quantum device, and hence, *in theory*, gives VQAs the potential for quantum advantage even on noisy machines.

Transient errors can be severely detrimental to VQAs. Iteratively, the VQA tuner estimates gradients of some form [21] across multiple circuit candidates to choose the best set of VQA parameters for further tuning. The VQA tuner works under the underlying assumption that the noise landscape of the device is unchanged during this gradient estimation process across the candidates. This is often not the case. Noise fluctuation can cause a significant transient impact on the objective function estimation of one or more

circuit candidates, and this can severely affect measured gradients and, therefore, the accuracy and convergence of VQA.

This discrepancy between VQA’s ideal and reality is illustrated in Fig. 1. The orange line shows the scenario with no noise, which is unrealistic (unless noise is entirely mitigated). The goal is to follow this line as closely as possible during VQA evaluation. The blue line shows VQA estimations if only static noise was present (this could potentially encompass dynamic noise that varies at very coarse temporal granularity). This is again unrealistic, but most prior error mitigation and optimization proposals have focused primarily on this setting, attempting to lower this blue line to bring it closer to orange, achieving considerable benefits. The reality is represented by the red line which shows that VQA estimations are affected by both static as well as transient noise - traditional VQA optimizations can be less effective in dealing with transients.

This paper proposes *QISMET: Quantum Iteration Skipping to Mitigate Error Transients*, to navigate the dynamic noise landscape of VQA. QISMET tries to eliminate the effect of VQA transient errors to the maximum extent possible—lowering the red line in Fig. 1 to bring it close to the blue line (shown in green). Other error mitigation techniques can be added orthogonally to bring the blue line (and, by extension, QISMET’s green line) closer to the orange ideal scenario.

QISMET is built with three key insights: ① QISMET identifies a key capability in VQAs: a previous VQA iteration acts as an optimal reference circuit to estimate the effect of transient errors on a current VQA iteration. To utilize this, the current execution block reruns the reference circuit from the previous iteration, and the difference between the VQA outcomes is estimated. ② Using traditional per-iteration gradient calculation, combined with QISMET’s novel transient error estimation, QISMET is able to predict an estimate of the transient-free gradient for every iteration. ③ Finally, QISMET designs a controller to decide if a particular iteration is acceptable to VQA or not. A VQA iteration is accepted only if the direction of the observed VQA gradient from the machine loosely matches the direction of the predicted transient-free gradient (relative to the previous iteration). In other words, a new set of VQA parameters is accepted only if the machine estimation of the parameters being good or bad, relative to the current parameters, matches the predicted transient-free estimation of the same (within some margin of error).

Contributions and Results:

① To the best of our knowledge, this work is among the first to study the effects of transient errors on VQAs.

② This paper proposes QISMET, which actively avoids instances of high fluctuating noise which have a significant transient error impact on specific VQA iterations and can severely impact VQA accuracy and convergence.

③ To achieve this, QISMET estimates transient error in VQA iterations and designs a controller to keep the VQA iteration gradients faithful to the transient-free scenario.

④ By doing so, QISMET efficiently mitigates a large portion of the effects of transient noise on VQAs and is able to improve the fidelity of VQAs by 1.3x-3x over a traditional VQA baseline, across different applications and machines, with 1.6-2.4x improvement over alternative approaches.




Figure 2: VQA: a hybrid algorithm that alternates between classical optimization and quantum execution.

⑤ To diligently analyze the effects of transients, this work also builds transient error noise models for target VQA applications from observing real device transients. These are then integrated with the Qiskit simulator.

2 BACKGROUND

Quantum Information: A qubit can exist in a linear superposition of the basis states $|0\rangle$ and $|1\rangle$, taking the form of $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. Together with other QC capabilities, this gives rise to the potential for quantum computers to exponentially outperform classical computers in certain applications. Qubits are manipulated through gates that modify their probability amplitudes, i.e. the coefficients of $|0\rangle$ and $|1\rangle$. There are many non-trivial single-qubit gates that can modify qubit state. Pairs of qubits can be manipulated via multi-qubit interactions such as the two-qubit controlled-X, or CX gate. Together, these enable universal quantum computation.

Noise in the NISQ era: NISQ devices are error-prone and up to around 100 qubits in size today [33]. These devices are extremely sensitive to external influences and require precise control, and as a result, some of the biggest challenges that limit scalability include limited coherence, gate errors, readout errors, and connectivity. We will dive deeper into different noise sources and errors in Section 3. Multiple forms of error mitigation strategies have been proposed to correct different forms of quantum errors. These include, but are not limited to, noise aware compilation [28, 40]; correcting measurement errors [8, 39]; scheduling for crosstalk [10, 29]; extrapolating for zero noise [20, 22, 23, 41]; decoherence mitigation through dynamical decoupling [7, 16, 18, 32, 38], spin-echo correction [15], gate scheduling [37]; etc. In addition, some of these can be used in conjunction to achieve better fidelity [34]. Many of these techniques make deployment decisions that are dependent on noise characteristics of qubits and gates which have to be captured through expensive full machine characterization, and are thus captured only at some coarse granularity of time (eg., roughly once a day on IBMQ machines). Therefore, they are suboptimal in the presence of fine-granularity dynamic fluctuations.

Variational Quantum Algorithms: Variational algorithms expect to have innate error resilience due to hybrid alternation with a noise-robust classical optimizer [25, 30]. An overview of this process is illustrated in Fig. 2. There are multiple applications in the VQA domain such as the Quantum Approximate Optimization

Algorithm (QAOA) [11] and the Variational Quantum Eigensolver (VQE) [30]. Our applications in this work target VQE but QISMET is broadly applicable across all VQAs. An important application of VQE is the ground state energy estimation of a molecule, a task that is exponentially difficult in general for a classical computer [14]. Estimating the molecular ground state has important applications in chemistry, such as determining reaction rates and molecular geometry. At a high level, VQE can be conceptualized as a guess-check-repeat algorithm. The check stage involves the preparation of a quantum state corresponding to the guess. This preparation stage is done in polynomial time on a quantum computer, but would incur exponential cost in general on a classical computer. This contrast gives rise to a potential quantum speedup for VQE [14]. The quantum circuit used in each iteration of VQE (and VQA in general) is termed an ansatz which describes the range of valid physical systems that can be explored and thus determines the optimization surface. Traditionally, the ansatz is parameterized by 1-qubit rotation gates. The VQA problem is represented as a Hamiltonian and is a linear combination of multiple Pauli terms. The lowest eigenvalue of the Hamiltonian corresponds to the system's ground state energy [25]. Every iteration, the VQA objective function calculates the expectation value of this Hamiltonian. This objective function is derived from ansatz measurements over different bases. An illustration of this is shown in Fig. 8.

Classical Tuning for VQAs: The classical tuner/optimizer variationally updates the parameterized circuit until the measured objective converges to a minimum. For simple VQA problems and in the presence of minimal noise, the optimization surface is convex around the global minimum and smooth [21]. As noise increases, in line with NISQ machines, the optimization surface can potentially become non-convex and non-smooth. The surface contour worsens as the problem complexity increases because of increase in circuit depth, number of parameters, and entanglement spread [21]. In this work we predominantly focus on the popular Simultaneous Perturbation Stochastic Approximation (SPSA) [4] optimizer which has shown some robustness to noise. SPSA is a method of stochastic gradient approximation, which only requires two measurements of the objective function per iteration, regardless of the dimension of the optimization problem.

3 MOTIVATION: TRANSIENT NOISE AND ITS IMPACT

3.1 Device Level

Non-uniformity among superconducting (SC) transmon qubits degrades the capacity of many near-term QCs to execute meaningful quantum workloads. At the physical level, every superconducting transmon qubit is unique because of the unavoidable process variation during device fabrication. These unintended differences across a quantum chip are called defects. One of the most severe defects that stochastically appears during transmon manufacture is a two-level system (TLS) [27]. The name TLS might be confusing because a qubit, when working as intended, is also a system with two levels. However, the difference is that a TLS defect is inopportunistically placed such that it destructively couples to qubits during computation, significantly reducing the duration of time in which quantum state information can be maintained [24].




Figure 3: Transient fluctuations in T1 times observed over 65 hours [9].

The Josephson Junction (JJ), or two superconductors separated by a thin metal-oxide insulator, is the critical component of the transmon qubit [12]. JJs are delicate components that must be fabricated with ultimate precision to achieve properties that enable the superconducting circuit to hold quantum state. Unfortunately, microscopic defects like TLS, taking the form of impurities inside materials or irregularities in atomic crystal structure, can appear in areas within the oxide layers or on the surface of the chip. TLS defects take on their own charge properties, and because each defect is unique, TLS is difficult to characterize. If TLS is close in proximity to the active components of the transmon, parasitic coupling causes qubit energy to be absorbed, causing fluctuation in qubit parameters such as T1 (amplitude) and T2 (phase) coherence times. Since transmon quantum systems are dynamic with evolving electric fields and states within JJs, the coupling strength of TLS that is near-resonant to a qubit varies over time, causing the qubit fluctuations to themselves be transient in nature [9, 36]. Apart from TLS, qubit parameters are also affected (potentially to a lesser degree) by thermal fluctuations, magnetic flux, quasi-particles, etc. [9, 19].

Fig.3 shows transient fluctuations in T1 times observed over 65 hours on a transmon qubit [9] for reasons described above. It should be noted that not all fluctuations are detrimental to a target circuit. We expect circuits to suffer significant fidelity impact only from fluctuations that result in very low T1 times. These are the outliers that are circled in the figure. Our expectation that impactful transients are an exception rather than the norm is evident in our analysis in the rest of this section and in our quantitative results. Furthermore, other fluctuating qubit parameters can also have significant impact; T1 times are just one example.

3.2 Circuit Level

Fluctuations in T1 times and other qubit/gate-level characteristics impact the execution of quantum circuits. The severity of impact could depend on:

- Ⓐ Width of the QC: More number of qubits increases the probability of high impact transients on some (or at least one) of the qubits,
- Ⓑ Depth of the QC: i) Deeper circuits are closer to the decoherence limit; therefore, a transient decrease in T1/T2 times can severely reduce circuit fidelity, and ii) deeper circuits usually imply more CX gates, which provide more sites for a substantial impact of error fluctuation,




Figure 4: Impact of transient errors on circuits. The top two figures show circuit fidelity data collected over a 45-hour period. Each data point is the average circuit fidelity obtained over a batch of 140 (same) circuits over a one-hour period. The zoomed figure shows the variation in fidelity over one of the batches.

© Circuit State: Effect of errors is state dependent. For example, a circuit that carries a superposition of states with a high proportion of 0s is less affected by decoherence.

Examples of the impact of transients on circuit fidelity are shown in Fig.4. The two primary figures show circuit fidelity data collected over a 45-hour period. Each data point is the average circuit fidelity obtained over a batch of 140 circuits over a one-hour period. The top circuit is a 4-qubit circuit with a depth of 6 CX gates. The average fidelity is around 83% and the total variation in fidelity is only around 5%. The bottom circuit is 8-qubits with a depth of approximately 50 CX gates. The average fidelity is around 25% and the variation is a concerning 35%. The zoomed figure at the bottom shows the fidelity variation over one of the batches, nearly 100% variation across the batch. The dashed line indicates a 90% threshold - in our applications we find that avoiding the extremes beyond this 90% is the most beneficial trade-off. More on this in Section 7.7.

3.3 Application Level (VQA)

Finally, the impact of transient errors on quantum circuits have a cumulative effective on long-running applications such as VQAs whose outcomes are influenced by per-iteration circuit fidelity. Discrepancies in gradient evaluations caused by transient noise fluctuations can lead the classical tuner to proceed along unfavorable directions in the VQA space, severely impeding convergence and accuracy. Section 4.1 discusses this in more detail.

A severe case of transient error affecting VQA is shown in Fig.5. This is obtained from an experiment run on IBMQ Jakarta over




Figure 5: Extreme impact of transient errors on VQA tuning (multiple sharp spikes), from an experiment run on IBMQ Jakarta over a 24-hour period.

roughly a 24-hour period. While we would ideally expect this VQA to somewhat monotonically converge to the minimum, multiple sharp spikes caused by transient fluctuations are seen. The spikes shown here are not the result of the optimizer jumping out of local minimas - such spikes are smoother and of lower amplitude. Although the detrimental impact of some transient fluctuations might be negated over a few iterations (for instance, in the case of the first circled spike), others fluctuations have a more lasting impact (second circled spike). The end expectation value reached at the end of 500 iterations is no better than that at 100 iterations, clearly showcasing the detrimental impact of transient noise. While such severe impacts might not always occur, milder scenarios are very common and are still detrimental, especially for long running use cases.

4 QISMET PROPOSAL

4.1 Navigating the Dynamic Noise Landscape

In every iteration of VQA, a classical tuner evaluates the gradient along a particular direction. Subsequently, it chooses directions which are some function of the prior evaluated gradients. Although different optimizers might implement this in various ways, most gradient-based approaches are some interpretation of the above. For instance, in a simple steepest-path approach, the direction with the steepest gradient (towards the optimum) is chosen. This tuning approach is built on the fundamental assumption that the direction gradients evaluated are affected by the same noise environment.

Fig. 6 shows the VQA landscape for a minimization task under different scenarios. The blue dot represents the current configuration. The other dots represent different VQA configurations evaluated by the tuner towards picking a subsequent configuration. The dot positions along the Y-axis show their objective function estimates obtained on the machine in the presence of transients. Dots above the dashed line are *perceived* by the VQA tuner to worsen the objective while those below are *perceived* to improve it ('perceived' because the *true* transient-free estimates can be different from the

machine obtained estimates, which could be skewed by transients). The color and number of each dot is indicative of its *true* objective function estimate in the absence of transients - note that these true estimates are not simply known in reality. Green dots improve the objective while red dots worsen it and the numbers indicate the magnitude. In a steepest-gradient approach, the VQA tuner would ideally want to select the *true* minimum - a green dot with highest magnitude number. But in reality it picks the *perceived* minimum - i.e., the dot which is lowest in the vertical direction (indicated by arrow). This setting is key to understanding the impact of transients on VQA and the QISMET solution.

Ⓐ The ideal scenario with no transients is illustrated in Fig. 6.a. Here, the configuration selected by the tuner, which is the lowest vertically, is, in fact, the green dot with highest magnitude (green-6). This is because the true and perceived estimates are equal in an ideal transient-free setting. Thus, it selects the true minimum among the evaluated samples. The effective VQA objective progress is shown in the inset at the bottom (moving lower is better) - in this case, it is optimal.

Ⓑ Unfortunately, this is often not the case in reality as shown in Fig. 6.b. Transient noise induces discrepancies in the gradient evaluations, and this can cause the tuner to perceive configurations differently from their transient-free true estimates. This is indicated in the figure by the colors/magnitudes not being correlated with vertical positions. Some true good configurations are perceived as bad (green dots above the dashed line) and some true bad configurations are perceived as good (red dots below the dashed line). Furthermore, in this scenario, the perceived minimum which the tuner selects, i.e., the lowest dot in the vertical direction, is in fact a true bad configuration (red-3). Choosing this configuration throws the VQA optimization significantly farther from the optimum, creating upward spikes as seen in Fig. 5 and thereby derailing accurate convergence. VQA progress is shown in the inset - the VQA objective worsens.

Ⓒ Fig. 6.c shows how QISMET would navigate this realistic landscape with transient errors. By estimating the transient noise (discussed in Section 5.1), QISMET tries to identify which perceived gradients are opposite in direction to their true gradients. In other words, from Fig. 6.b QISMET tries to predict which green dots are above the dashed line and which red dots are below. It then re-executes these configurations on the quantum device until they are aligned with their true gradient directions (i.e., greens go below the dashed line and reds go above) or the retry budget is exhausted. In Fig. 6.c the retried configurations are circled in black. Over the retrials, most (but not all) configurations realigned with their true gradient directions. The realignment would happen for a configuration if it is able to execute in an instance of low transient noise - this is highly probable since transients of significant detrimental effect are fairly uncommon (as seen in Figures 3 and 4). Eventually when selecting the new configuration, the tuner again selects the perceived minimum, i.e., the lowest vertical dot, but now this is a good true minimum (green-5), even if it is not the best (which was green-6 in the ideal case). Since the selected configuration is still a good true minimum it does not derail the VQA optimization even in the presence transients. VQA progress is shown in the inset - the VQA objective improves, even if not optimal (compared to the ideal




Figure 6: QISMET navigating the VQA landscape. The pale blue dot represents the current configuration. Other dots represent different VQA configurations evaluated by the tuner. Y-axis positions indicate objective estimates (skewed by transients) as perceived by the VQA tuner - below the dashed line is perceived to be good. Numbers indicate true (transient-free) magnitude of objectives, while red and green indicate bad and good respectively. (a) In the ideal scenario, true and perceived estimates match and the tuner picks the true minimum (green-6). (b) In reality, the perceived minimum chosen by the tuner is a true bad configuration (red-3). (c) By estimating the transient effects, QISMET identifies some transient-affected configurations to be rerun. These are circled in black. Rerunning these configurations enables most of them to escape significant transients, thus most true and perceived estimates get reasonably aligned. The perceived minimum is then chosen, which is green-5. The inset squares show effective VQA objective progress (moving lower is better).

case). Thus, QISMET is more robust to the dynamic noise landscape of VQA.

4.2 Per-iteration QISMET Functionality

Fig. 7 shows an instance of how QISMET functions over multiple iterations of VQA. A sequence of VQA ‘jobs’ are run on the quantum machine - each job is a collection of independent circuits. If VQA is run with some error mitigation (measurement error mitigation, for example), job α would consist of circuits corresponding to the $(i)^{th}$ VQA iteration (orange inner box) and some related error mitigation circuits (dark gray inner box). QISMET adds some reference circuits to job α . These are the repeated circuits from the $(i - 1)^{th}$ VQA iteration (yellow box). Circuits in the other jobs shown can be similarly inferred from the figure. Progress in the VQA algorithm is determined by the QISMET controller which is shown in the triangle labeled ‘C’ between jobs. The following steps can occur:

- ① Upon executing job α , the QISMET controller takes as input: a) iteration $(i - 1)$ circuits from the previous job ω (yellow), b) its repetitions in the current job, i.e., iteration $(i - 1)$ circuits in job α (yellow), and c) the primary circuits of job α , which correspond to iteration (i) (orange). In this case, the controller is shown to make the decision that the impact of transient noise on the $(i)^{th}$ iteration is reasonable.
- ② Thus, VQA proceeds to job β with its primary circuit as VQA iteration $(i + 1)$ (blue), along with other circuits, as shown in figure.
- ③ After completion of job β , the QISMET controller now takes as input: a) the iteration (i) circuits from the previous job α (orange), b) its repetitions in the current job, i.e., iteration (i) circuits in job β (orange), and c) the primary circuit of job β , which correspond to iteration $(i + 1)$ (blue). In this instance, the controller is shown to

make the decision that the transient noise impact is significant and can be detrimental to VQA.

- ④ Therefore, VQA does not proceed. Instead the VQA iteration $(i + 1)$ circuits along with other circuits are repeated via job γ (hence same coloring as job β). Note: since some transients effects can stay for extended time periods, the number of such rejections and repetitions is limited to some small max-out.

- ⑤ When the controller checks again (note that it obtains the iteration (i) circuits in orange from job α), it is deemed that significant fluctuation has passed and the current impact of transient noise is within reasonable limits.

- ⑥ Thus, VQA proceeds to job δ with the primary circuit as VQA iteration $(i + 2)$, along with other circuits.

By taking steps as described above, the QISMET controller is able to navigate the dynamic noise VQA landscape, making decisions to the effect of Fig. 9 as discussed in Section 4.1. Transient error estimation is discussed in Section 5.1 and controller logic in Section 5.2.

5 QISMET DESIGN

The QISMET proposal has two major components. The first component estimates transient error in each VQA iteration and uses this to make transient-free predictions. The second component uses the transient-free predictions to skip/retry particular iterations which are deemed harmful to VQA convergence.

5.1 Transient Estimation, Transient-Free Prediction

Fig. 8 is a detailed expansion of any single iteration of Fig. 7. Note that no error mitigation circuits from Fig. 7 are shown. The left of Fig. 8 shows the VQA iteration (i) circuit that was run in the ‘previous’




Figure 7: Multiple VQA ‘jobs’ (in light gray) run on the quantum machine. Job α consists of circuits corresponding to the $(i)^{th}$ VQA iteration (orange box), some error mitigation circuits (dark gray box), and QISMET adds circuits from the previous VQA iteration $(i - 1)$ (yellow box). Other jobs/iterations can be inferred similarly from the figure. Progress in the VQA algorithm is determined by the QISMET controller (triangle labeled ‘C’). The following steps can occur: (1) After job α , the QISMET controller takes as input the $(i - 1)^{th}$ iteration circuits from previous job ω (yellow), the repeated $(i - 1)^{th}$ iteration circuits in job α (yellow), and the i^{th} VQA iteration from job α (orange). The controller then makes a decision that the transient noise impact on the $(i)^{th}$ iteration is reasonable. (2) VQA thus proceeds to the job β with the primary circuit as the $(i + 1)^{th}$ VQA iteration (blue), along with other circuits. (3) In this case, the controller deems that the transient noise impact is significant and can be detrimental to VQA. (4) Therefore, VQA does not proceed. Instead, the circuits are repeated via job γ . (5) When the controller checks again, it is deemed that the transient impact is within reasonable limits. (6) Thus, VQA proceeds to the next iteration and beyond.




Figure 8: Detailed expansion of Fig. 7. VQA iteration (i) in job α produces energy estimate $E_m(i)$. Iteration $(i + 1)$ in job β produces $E_m(i + 1)$. The rerun of iteration (i) in job β produces $E_{mR}(i)$. The gradient with transients, as observed by a traditional VQA tuner, is $G_m(i + 1)$. QISMET estimates the machine transient noise on iteration $(i + 1)$ as $T_m(i + 1)$ and then predicts transient-free $E_p(i + 1)$ and $G_p(i + 1)$. A VQA iteration is accepted by QISMET only if G_p and G_m point in the same direction as shown in Fig. 9.

job α . Its objective function (or energy) estimation on the machine is denoted by $E_m(i)$. To its right is ‘current’ job β in which two circuits are shown. The left circuit is the rerun of iteration (i) and its energy estimation on the machine is denoted by $E_{mR}(i)$. The right circuit is the iteration $(i + 1)$ VQA circuit and its energy estimation on the machine is denoted by $E_m(i + 1)$. The machine gradient for iteration $(i + 1)$ is $G_m(i + 1)$ and is the difference between the machine energy estimates of iterations $(i + 1)$ and (i) . This is the gradient estimate that a traditional VQA tuner would use towards parameter selection for the next iteration. QISMET estimates the transient error in iteration $(i + 1)$ i.e., $T_m(i + 1)$, as the energy difference between the rerun of iteration (i) in job β and its original run in job α . QISMET then predicts the transient-free energy for iteration $(i + 1)$ i.e., $E_p(i + 1)$, relative to (i) by removing the transient error component from $E_m(i + 1)$. The predicted transient-free gradient

is then $G_m(i + 1)$. Equations are shown to the right of Fig. 8. This data is passed onto the QISMET controller which is discussed next in Section 5.2.

Note that this assumes that the impact of the transient error on iteration $(i + 1)$ is the same as that on iteration (i) . Although this is not necessarily exactly true, circuit (i) is the closest possible reference circuit for transient noise estimation on circuit $(i + 1)$, so its choice is justified.

5.2 Gradient Faithful QISMET Controller

Fig. 9 is an illustration of different scenarios presented to the QISMET controller and the corresponding decisions made so that it achieves the navigation illustrated in Fig. 6.

The intuitive idea is that the controller accepts VQA iterations only if the directions of the true gradients (i.e., gradients obtained




Figure 9: QISMET's gradient faithful controller. Green dots: VQA machine estimates for iteration (i) . Blue dots: machine estimate for iteration $(i + 1)$. Yellow lines: negation of transient noise between (i) and $(i + 1)$. Orange dots: predictions of transient-free estimates for iteration $(i + 1)$. Difference scenarios and the reaction of the controller are shown. (a) and (b): Both the machine and the transient-free estimates have a positive gradient. Thus, the estimates are acceptable. (d) and (e): Both machine and predicted transient-free estimates have a negative gradient, and are thus acceptable. (c) Machine estimate has a positive gradient but the predicted estimate has a negative gradient. This instance is unacceptable, since it makes a potentially bad VQA configuration be interpreted as a good one. (f) This is the opposite of (c) and is again unacceptable. Note that gradient swings within the shaded threshold region are always accepted.

in an ideal transient-free scenario) match the direction of the perceived gradients (i.e., gradients observed by the VQA tuner based on machine energy estimations). This prevents scenarios such as truly bad VQA configurations being perceived as good configurations, which could throw the VQA optimization away from its path towards the optimum target.

In the figure, the green dots are VQA energy estimates made on the quantum machine for iteration (i) , i.e., $E_m(i)$. The blue dots are machine energy estimates for iteration $(i + 1)$, i.e., $E_m(i + 1)$. The yellow lines are the negation of the estimated transient error between (i) and $(i + 1)$, i.e., $-T_m(i + 1)$. They are used to produce the orange dots which are the predictions of the transient-free energy for iteration $(i + 1)$ relative to (i) , i.e., $E_p(i + 1)$. Different scenarios and controller decisions are discussed below.

① In Fig. 9.a and b, both the machine estimates and the predicted transient-free estimates for iteration $(i + 1)$ have a positive gradient wrt (i) , i.e., $G_m(i + 1)$ and $G_p(i + 1)$ are both positive. Thus, the VQA iteration is acceptable, as the direction of the gradient is maintained.

② In (d) and (e), the machine gradients and predicted transient-free gradients of $(i + 1)$ wrt (i) are both negative, thus this iteration is acceptable.

③ In (c), the machine estimate for $(i + 1)$ has a positive gradient wrt (i) , but since the transient noise is significantly large, the predicted estimate (after negation of the transient noise) has a negative gradient. This is unacceptable since it makes a potentially bad VQA configuration be interpreted as a good one.

④ (f) is the inverse of (c) - again unacceptable because it makes a good VQA configuration be interpreted as bad.

⑤ Finally, the gradient swings within a specified error threshold region (pink shaded region) are always accepted. This is to avoid frequent skipping on less impacting transients. This is discussed and evaluated in Section 7.7.

5.3 Other Skipping and Filtering Techniques

It could be argued that an alternative intuitive solution is to avoid all instances in which transient noise is greater than some set threshold. The controller could simply be designed to skip a VQA iteration (i) if $\text{abs}(T_m(i)) > \tau$. We find that such a solution is unsuitable because many transients (even of reasonable magnitude) that are constructive to VQA progress, are skipped, thus, affecting accuracy and substantially delaying convergence. A quantitative comparison is discussed in Section 7.3.

The above is a fundamental difference between the QISMET gradient faithful controller and other filtering techniques that are common in classical computing (such as in signal processing). For example, Kalman filtering [45] is a signal processing approach that uses a system's environment model, some known inputs to that system, and multiple noisy measurements (say, from sensors), to build an understanding of the system's noise/variation, and then calculate noise-free estimates. It is generally suitable to application settings with dynamic variations similar to our focus in this work. A key feature of such filtering techniques is that they treat all forms of measurement variance to be the same and eliminate the variance in a principled way by inferring some information from multiple measurements.

Unfortunately, this is less suited to the VQA tuning landscape. As discussed in Section 5.2, not all transients are harmful to VQA progress and only those that flip the gradient directions are detrimental and need to be avoided. From Fig. 9, let's consider scenarios (a) and (c). While both have significant transient error magnitudes, (a) is not detrimental to VQA progress and is accepted by QISMET while (c) is deemed detrimental and is rejected. In contrast, the Kalman Filter (and related techniques) will incorporate some information from both (a) and (c) towards its elimination of noise and estimation of the true value. When carefully tuned, the differences in these scenarios can be incorporated into the filter to some extent, but its efficiency will be limited. With less careful tuning, the filtering might, in fact, prove to be detrimental. We quantitatively showcase different filtering scenarios in Sec. 7.4.

It should also be noted that Kalman and other classical filtering techniques work best when the signal and noise follow some specific fairly stable characteristics. Unfortunately, as per our current understanding of quantum devices, the transient noise dynamism

Table 1: TFIM VQA applications for simulation

Application	Qubits	Ansatz	Reps	Machine + trial
App1	6	SU2	2	Toronto (v1)
App2	6	RA	4	Guadalupe (v1)
App3	6	RA	4	Guadalupe (v2)
App4	6	SU2	4	Toronto (v2)
App5	6	RA	8	Cairo (v1)
App6	6	RA	8	Casablanca (v1)

is largely non-deterministic and constantly varying over fine granularities of time.

6 METHODOLOGY

6.1 Applications and General Infrastructure

QISMET is a software optimization that can be integrated with any VQA classical optimizer. We implement it in Python and use it with the Qiskit VQA framework [5], through which it interacts with quantum execution.


We limit ourselves to one VQA domain, the VQE, which was introduced in Section 2. Due to noisy machine limitations on circuit width and depth, applications are restricted to 6 qubits. Our evaluations encompass 2 Hamiltonians, 6 different ansatz, and 7 IBMQ machines (as well as simulation). We primarily use the SPSA classical tuner across all our evaluations.

Our primary focus is the one-dimensional TFIM (Transverse Field Ising Model) Hamiltonian, which is an ubiquitous model that has applications in understanding phase transitions in magnetic materials [42]. The TFIM is a desirable system since it is exactly solvable via classical means. Further, the Hamiltonian is easily constructed to create small problem sizes. We also evaluate the potential energy of the Hydrogen molecule over bond lengths of 0.4 - 2 Å.

We use the hardware efficient SU2 [2] and RA [1] ansatz which are of low depth and therefore suited to NISQ devices. The number of block repetitions in the ansatz is varied between 2, 4 and 8 repetitions. IBM Q machines targeted for machine runs are Guadalupe (27q), Toronto (27q), Sydney (27q), Casablanca (7q), Jakarta (7q), and Mumbai (27q). IBM Q machines from which traces are generated for simulation are Guadalupe (27q), Toronto (27q), Cairo (27q), and Casablanca (7q). Machine details can be found on the IBM Quantum Systems page [3]. TFIM simulation application details are summarized in Table 1.

6.2 Building Transient Noise Models for Simulation

While execution on real quantum machines is a must for the holistic evaluation of application fidelity, sufficient access to machines is still limited today, especially for long-running applications such as VQAs. Further, fair evaluation requires a deterministic environment across different points of comparison, as well as reproducible results to validate and improve upon the proposal. To enable longer VQA runs in a deterministic environment and for reproducible comparisons, we are required to explore the use of IBM's quantum


Figure 10: VQA simulation on the Qiskit simulator with our added transient noise model.

simulator (via Qiskit). Unfortunately, the noise models traditionally available to Qiskit are static over the period of a machine's calibration cycle - they are refreshed roughly once per day.

To enable fine-granularity transient noise analysis in simulation, we build transient error traces for our target application-machine combinations. Per-iteration transient effects on VQA are captured and normalized to the magnitude of the VQA estimations. These transient effects are composed into a data structure and integrated into Qiskit's VQA framework. In each simulated VQA iteration, an instance of transient noise is accessed from the data structure corresponding to the particular application-machine pair. This transient noise is factored in with existing static components from Qiskit's noise models.

Furthermore, the magnitude of transient noise can also be varied to perform a wider range of analyzes, as shown in Fig.10. The figure shows a variety of transient noise magnitudes applied to a VQA application, varying from 0-50% of the ideal VQA objective estimations. Intuitively, it is evident that as the impact of transient noise increases, the accuracy and convergence of the VQA estimates worsen. Similar experiments can be performed to test the capability of new proposals - we use this simulator to extensively evaluate QISMET.

6.3 Evaluation Comparisons

We evaluate QISMET based on VQA energy estimates (using the standard Hartree Energy metric when applicable). In one or more experiments, we compare:

Baseline: Traditional VQA employing measurement error mitigation but no mitigation for transient errors.

QISMET: Baseline + control to estimate transients and skip/rerun iterations when appropriate. The QISMET error threshold is set so as to skip at most 10% of the iterations (found to be the most suitable - Section 7.7).

Blocking: A Qiskit SPSA optimization that only accepts VQA updates that move towards the objective.

Resampling: SPSA optimization that increases the number of times the gradient is sampled (we use 2x).

2nd-order: A Qiskit SPSA optimization that estimates second-order derivatives to condition the gradient.

Noise-free: An ideal scenario without any form of error (run on the Qiskit simulator).




Figure 11: QISMET benefits for a 6-qubit TFIM VQA on IBMQ Guadalupe run over a 48-hour period. Two instances of moderate transient error are circled and are avoided by QISMET.




Figure 12: QISMET benefits for a 6-qubit TFIM VQA on IBMQ Sydney run over a 48-hour period with one sharp instance of transient error (circled) and are avoided by QISMET.

QISMET-conservative: QISMET with threshold set to skip at most 1% of the VQA iterations.

QISMET-aggressive: QISMET with threshold set to skip at most 25% of the VQA iterations.

Only-transients: An alternative proposed technique that skips VQA iterations if the estimated transient error is larger than some specified threshold.

Kalman: A classical filtering approach discussed in Section 5.3. It requires tuning filter hyper-parameters such as the Transition Coefficient (T), which is a linear estimation of the slope of the noise-free curve, and the Measurement Variance (MV), which is related to the variance in the noisy measurements.

7 EVALUATION RESULTS




Figure 13: QISMET benefits for a 6-qubit TFIM VQA on six IBMQ machines. Iterations varies across applications, depending on machine availability.




Figure 14: Simulating VQA with transient errors for App2, using the SPSA tuner, over 2000 iterations. QISMET is compared against SPSA optimizations.

7.1 Real Machine Experiments

First, we analyze VQA experiments run on the real IBMQ quantum machines. Note that the number of iterations that we can run on these machines is limited due to access constraints.

Fig.11 shows a comparison of QISMET against the baseline on a 6-qubit TFIM VQA on IBMQ Guadalupe. Both versions of the application are run for around 270 iterations, over a 48 hour period. The versions are run in a synchronous manner, so each iteration of the baseline is run temporally adjacent to the corresponding iteration of the QISMET version. Many instances of moderate transient error impact the baseline - two phases with transient errors are circled. While the baseline quickly recovers from the first transient phase, the second phase is more harmful, causing the VQA accuracy to somewhat stagnate over the next 50-100 iterations. These transient phases are predominantly avoided with QISMET, leading to a 40% improvement in VQA estimation over 270 iterations.

Fig.12 shows a comparison of QISMET against the baseline for the same 6-qubit TFIM VQA on IBMQ Sydney. Both versions of the application are run for around 350 iterations, over a 48 hour period. While the machine behavior is smooth for most of the tuning period, there is a single phase of significant turbulence (circled) which heavily impacts the baseline. QISMET is able to avoid this turbulent phase - QISMET is especially effective here since the transients only occur for a few iterations - once QISMET skips those errors, it is able to continue its steady progress, achieving a 50% improvement over the 350 iterations.

Fig.13 shows the QISMET benefits over the baseline on the primary vertical axis and the number of iterations run on the secondary vertical axis. Results are shown across 6 different IBM Q machines. QISMET improves the measured VQA expectation by by 29-51% across the machines, with a mean improvement of 39%. These improvements are achieved over 200-450 VQA iterations run on the devices. We expect that the relative QISMET benefits would be greater if more iterations were run since there would be increased potential for more transients to impact the baseline.




Figure 15: Only-transients skipping approach for App1.

7.2 Comparing against SPSA Optimization Schemes

Fig.14 shows simulated evaluations of App2 (see Table 1) for five schemes - Baseline, QISMET, Resampling, Blocking and 2nd-order. These are introduced in Section 6. Simulations are run for 2000 iterations using the SPSA tuner, with convergence generally beginning at around 1250 iterations. QISMET performs best, achieving a VQA expectation of -1.5, which is a 65% improvement over the baseline.

The Blocking and Resampling schemes also show some improvement over the baseline. While the Blocking scheme can avoid transient errors that move VQA away from the objective, it also hurts the ability to escape from local minima. Similarly, while the Resampling scheme's increase in samples offers some transient robustness, it has twice the computational cost per iteration and therefore can be 2x slower to converge. Therefore, the VQA improvements obtained from these are around 30% lower than QISMET. The 2nd-order scheme uses 2nd-order derivatives to influence the gradients calculated in each iteration. We observe this scheme to in fact be detrimental in the presence of transients - imperfect 2nd-order derivatives can potentially skew the gradients even further away from optimal. This scheme performs 35% worse than the baseline and roughly 2.5x worse than QISMET.

7.3 Alternative Only-Transients Skipping Approach

Fig.15 shows the the alternative skipping approach discussed in Section 5.3, evaluated for App1. Thresholds (based on which the controller decides to skip or not) vary from a very high threshold, which keeps the % of skips below 1% (99p), to a low threshold, which allows up to 50% of the iterations to be skipped (50p). In all scenarios, the resulting VQA estimates are worse than the baseline and the higher skipping thresholds always perform better. This is clearly indicative that the skipping technique is more harmful than useful, motivating the need for a more intelligent approach such as QISMET.

7.4 Comparing against Kalman filtering

Furthering the discussion in Section 5.3, we compare different Kalman instances (with different values of T and MV parameters -




Figure 16: Simulated comparison of Kalman filter against Qismet and baseline, for App6. Two Kalman hyper-parameters (described in Section 6.3) are tuned over a limited search space. While good choices of Kalman hyper-parameters can provide some improvements over the baseline, the benefits are substantially lower than those from Qismet.

details in Section 6.3) in Fig.16. We show its evaluation for App6, against the Baseline and Qismet, for 500 iterations, in simulation. Here, the Kalman filtering is applied on top of the noisy VQA tuning performed with SPSA. Note here that the T and MV values are tuned via an oracle approach, with noise known apriori, purely to maximize benefits.

We show 6 different Kalman instances with MV=0.01/0.1 and T=0.9/0.99/1 in Fig.16. A low MV value indicates that the filter assumes low influence of variation in the machine measurements - thus the filter does not actively combat it. Thus, any instance of high transient noise significantly impacts the filter (seen in the pink line, for example). A high MV value indicates the possibility of high variation influence in the measurements - so the filter actively avoids this. Unfortunately, the filter is unable to differentiate between machine noise and actual algorithmic variance (navigating local minima etc) - thus, this filtering technique saturates quickly and poorly (seen in the dark blue line, for example). A T value farther away from 1 imposes a forced downward descent on the estimated values - while this is helpful in true downward slopes, it is bad for escaping minima. In this application example, the most stable low estimate is produced by (MV=0.1, T=1) whereas the lowest but an erratic estimate is produced by (MV=0.01, T=1). While the estimates are as much as 1.4x better than the baseline, QISMET estimates are 3x better than the best Kalman variant, while also being stable.

7.5 Simulated Comparisons for Multiple Benchmarks

Next, in Fig.17 we perform evaluation comparisons for five schemes - QISMET, Blocking, Resampling, 2nd-order, and Kalman filter, across the six applications listed in Table 1. All evaluations are simulated for 2000 iterations and optimized with the SPSA tuner. QISMET provides consistent benefits, always outperforming the baseline and competing schemes, achieving mean improvements




Figure 17: QISMET benefits from the six simulated applications using the SPSA tuner run for 2000 iterations. QISMET provides consistent benefits, outperforming the baseline and competing schemes (Blocking, Resampling, 2nd-order and Kalman filtering).

of 2x (up to 3x), 1.7x, 1.6x, 2.4x, and 1.85x respectively. Note that Blocking and Resampling schemes are inconsistent, offering a mean improvement over the baseline (1.2-1.25x), but perform worse on some applications (highlighted in red). As noted earlier, while these schemes have some robustness to transients, they are slow to accurate convergence due to increased per-iteration evaluation. The 2nd-order scheme consistently performs worse than the baseline and is not observed to be a suitable solution for transient errors. The Kalman filter hyper-parameters are tuned as discussed in Section 7.4 and only the best-case results are shown. While there is non-negligible improvement over the baseline, the benefits are fairly low due to the reasons described in Section 5.3. Clearly, Qismet achieves a considerably higher improvement. Also, the best Kalman instance varies across different applications, further highlighting the challenges. Finally, it should be noted that QISMET benefits are generally greater for deeper circuits and/or on traces from noisier machines. This is intuitive and in line with the discussion in Section 3.2.




Figure 18: Effect of transient noise on multi-VQA experiments, shown by energy estimation of H_2 .

7.6 Multi-VQA Experiments for Molecules

In molecular chemistry, it is important to estimate the difference in VQE estimates across multiple Hamiltonians [43]. Here, each Hamiltonian models a particular geometry of the molecule (usually different bond lengths), and the VQE estimate is a measure of the

molecular potential energy. The difference between the energy estimates at different bond lengths is indicative of the chemical reaction rates for the molecule of interest. If transient noise affects some of the VQA experiments more than others, then the difference in energy estimates can be very skewed.

Fig.18 shows potential energy estimation for the H_2 molecule through simulation. This simulation setup only uses transient noise and no static noise component. Estimations are performed for 10 different $H-H$ bond lengths, each corresponding to a unique Hamiltonian and VQE experiment. The ideal noise-free scenario is shown, along with estimates obtained with QISMET and the baseline. Clearly, QISMET provides high-accuracy estimates, and the shape of the bell curve closely models the noise-free scenario. On the other hand, the baseline steadily deviates away from noise-free. Greater deviation is expected at higher bond lengths since the quantum component of the potential energy is more significant in this region (and therefore, so is the quantum error).




Figure 19: Different QISMET error thresholds are analyzed on two simulated use cases which experience low and high transient noise respectively.

7.7 Sweeping the QISMET Error Threshold

In Section 5.2, we discussed the use of a threshold to control the fraction of transients that are skipped by the QISMET controller. In Fig.19 three different thresholds are analyzed along with the baseline - these were introduced in Section 6. The evaluation is shown for two simulated use cases which experience low and high transient noise respectively. For the chosen use cases, the conservative threshold skips too few instances to have a significant impact and, therefore, performs similarly to the baseline. The aggressive threshold pushes QISMET to be worse than the baseline in the scenario

with low transient noise - too many iterations with low transient noise are skipped unnecessarily which delay convergence. In the scenario with high transient noise, the aggressive threshold still achieves considerable benefit since there are enough such impacting transients to skip. The best-case threshold is a good trade-off and is able to achieve the highest benefits in both scenarios, achieving a 1.2x and 3x improvement over the baseline, respectively. Note that intelligent dynamic thresholding can potentially be used to improve these benefits further, but is beyond our current scope.

8 DISCUSSION

8.1 Tuning QISMET Parameters

There are two tunable parameters in the QISMET framework. The more impactful one is the error threshold, which is discussed in Section 5.2 and evaluated in Section 7.7. If transient errors are very few and of very large magnitude, then conservative thresholds are suitable—QISMET will predominantly follow the baseline and only avoid the transient errors which clearly stand out as egregious occurrences. On the other hand, transient errors which are more frequent and of magnitudes that are not too large but still detrimental, will benefit from more aggressive thresholds.

The other tunable parameter is the retry budget or the maximum number of repetitions that QISMET will perform on any iteration that does not meet the error threshold. In this work we fix this to 5 repetitions. The optimal number for this budget could be derived from an understanding of the typical time duration for which transient effects last. As long as the execution time taken by the retry-budget number of repetitions is slightly greater than the transient error duration, then an instance without the transient error's effects can be achieved for the given iteration. At the same time, if some device characteristic change occurs for a long time duration, i.e., it is not transient (eg., machine re-calibration), then we would want the change to be quickly accepted by QISMET, and adapted to by the VQA tuner in future iterations. From our real device experiments, we observe that transient errors disappear in one or two repetitions.

8.2 QISMET Suitability

QISMET is most effective and has most significant benefits when transient errors are of high magnitude and of reasonably short duration. Such instances can be identified and skipped in just a single trial. Our empirical real machine observations suggest that this is often the case. Transients of lower magnitude are not harmful to either QISMET or the baseline. There are some particular scenarios in which QISMET could perform poorly. One, if transient errors are gradually accumulating, then the errors would always be acceptable to the threshold but they might push VQA far from optimality. In this scenario though, QISMET would perform no worse than the baseline. Two, if high magnitude transient errors are of very long duration, then the effect of the transient is accepted after a retry budget worth of skips—thus, the negative impact on VQA estimation is similar to the baseline, but QISMET also suffers from lost iterations in this scenario—thus QISMET is worse off. While there could be other adversarial scenarios that negatively impact

QISMET, quantum devices are too noisy today for contrived adversarial patterns to have any realistic likelihood of occurrence over a prolonged period.

8.3 QISMET Overheads

QISMET requires that each execution instance runs the current VQA iteration as well as a rerun of the previous iteration. In the absence of other error mitigation and/or supporting circuits, this means that the circuit execution overhead of QISMET is at least 2x (compared to a baseline with no transient errors). For the above scenario, in the absence of no skips, the overhead is exactly 2x, whereas the overhead increases with the number of skips. On most fairly reliable machines, the number of skips is extremely small—only the egregious transients are being avoided. However, it should be noted that error mitigation circuits (such as those to combat measurement errors) are often run alongside the primary circuit. These supporting circuits are often many in number, scaling in proportion to the number of measurement bits. In the presence of such circuits, QISMET overheads are reduced. Further, any QISMET overhead will only be negatively reflected if transients are entirely absent—in practical settings which do have some transients, avoiding them with QISMET substantially lowers the overall iteration and execution time cost.

Kalman comparison: While, at first glance, it may seem that there are no overheads to the Kalman filtering (since circuits are run only once), it should be noted that the oracle (or any) approach used to tune the hyper-parameters requires some profiling, which would likely be much more expensive than the QISMET scenario, as described above.

8.4 Other Classical Techniques

In Sec.7.4, we showed that QISMET can substantially outperform classical filtering techniques such as the Kalman filter. Other filtering techniques are also available, such as Constant False Alarm Rate (CFAR) [35] detection, which is used in radar systems to detect target returns against a background of noise and interference. Similar to Kalman, they are limited in their capability to avoid specific detrimental transients. Finally, while the above filtering techniques are fairly unsupervised, supervised techniques are impractical since there is no useful labeled training data available, since the data would have to be machine-application pair dependent and stable over long periods of time.

9 CONCLUSION

Transient errors can be very detrimental to quantum computing in the NISQ era, especially for long-running applications such as VQAs. This paper proposes QISMET to navigate the dynamic noise landscape of VQA. It actively avoids instances of high fluctuating noise which can have a significant transient error impact on specific VQA iterations. To achieve this, QISMET estimates transient noise in VQA iterations and designs a controller to keep the machine-obtained iteration gradients (specifically, their directions) faithful to their corresponding transient-free predicted estimates.

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; in part by NSF award 2110860; 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) and in part based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers. This work was completed in part with resources provided by the University of Chicago's Research Computing Center. GSR is supported as a Computing Innovation Fellow at the University of Chicago. This material is based upon work supported by the National Science Foundation under Grant # 2030859 to the Computing Research Association for the CIFellows Project. KNS is supported by IBM as a Postdoctoral Scholar at the University of Chicago and the Chicago Quantum Exchange. HH is supported by NSF (grants CCF-2119184, CNS-1956180, CNS-1952050, CCF-1823032, CNS-1764039), ARO (grant W911NF1920321), and a DOE Early Career Award (grant DESC0014195 0003). FTC is Chief Scientist for Quantum Software at ColdQuanta and an advisor to Quantum Circuits, Inc.

REFERENCES

- [1] 2021. IBM Quantum RA ansatz. <https://qiskit.org/documentation/stubs/qiskit.circuit.library.RealAmplitudes.html>.
- [2] 2021. IBM Quantum SU2 ansatz. <https://qiskit.org/documentation/stubs/qiskit.circuit.library.EfficientSU2.html>.
- [3] 2021. IBM Quantum Systems. <https://quantum-computing.ibm.com/services?systems=all>.
- [4] 2021. SPISA: Simultaneous Perturbation Stochastic Approximation method. <https://www.jhuapl.edu/spisa/>.
- [5] Héctor Abraham, AduOfiei, Rochisha Agarwal, Ismail Yunus Akhalwaya, Gadi Aleksandrowicz, Thomas Alexander, Matthew Amy, Eli Arbel, Arijit02, Abraham Asfaw, Artur Avkhadiiev, Carlos Azaustre, AzizNgoueya, Abhik Banerjee, Aman Bansal, Panagiotis Barkoutsos, George Barron, George S. Barron, Luciano Bello, Yael Ben-Haim, Daniel Bevenius, Arjun Bhobe, Lev S. Bishop, Carsten Blank, Sorin Bolos, Samuel Bosch, Brandon, Sergey Bravyi, Bryce-Fuller, David Bucher, Artemiy Burov, Fran Cabrera, Padraic Calpin, Lauren Capelluto, Jorge Carballo, Ginés Carrascal, Adrian Chen, Chun-Fu Chen, Edward Chen, Jielun (Chris) Chen, Richard Chen, Jerry M. Chow, Spencer Churchill, Christian Claus, Christian Clauss, Romilly Cocking, Filipe Correa, Abigail J. Cross, Andrew W. Cross, Simon Cross, Juan Cruz-Benito, Chris Culver, Antonio D. Córcoles-Gonzales, Sean Dague, Tareq El Dandachi, Marcus Daniels, Matthieu Dartiaill, DavideFrr, Abdón Rodríguez Davila, Anton Dekusar, Delton Ding, Jun Doi, Eric Drechsler, Drew, Eugene Dumitrescu, Karel Dumon, Ivan Duran, Kareem EL-Safty, Eric Eastman, Grant Eberle, Pieter Eendebak, Daniel Egger, Mark Everitt, Paco Martín Fernández, Axel Hernández Ferrera, Romain Fouilland, FranckChevallier, Albert Frisch, Andreas Fuhrer, Bryce Fuller, MELVIN GEORGE, Julien Gacon, Borja Godoy Gago, Claudio Gambella, Jay M. Gambetta, Adhisha Gammanpila, Luis Garcia, Tanya Garg, Shelly Garion, Austin Gilliam, Aditya Giridharan, Juan Gomez-Mosquera, Salvador de la Puente González, Jesse Gorzinski, Ian Gould, Donny Greenberg, Dmitry Grinko, Wen Guan, John A. Gunnels, Mikael Haglund, Isabel Haide, Ikko Hamamura, Omar Costa Hamido, Frank Harkins, Vojtech Havlicek, Joe Hellmers, Łukasz Herok, Stefan Hillmich, Hiroshi Horii, Connor Howington, Shaohan Hu, Wei Hu, Junye Huang, Rolf Huisman, Haruki Imai, Takashi Imamichi, Kazuaki Ishizaki, Raban Iten, Toshinari Itoko, JamesSeaward, Ali Javadi-Abhari, Jessica, Madhav Jivrajani, Kiran Johns, Scott Johnston, Jonathan-Shoemaker, Vismai K, Tal Kachmann, Naoki Kanazawa, Kang-Bae, Anton Karazeev, Paul Kassebaum, Josh Kelson, Spencer King, Knabberjoe, Yuri Kobayashi, Arseny Kovyrshin, Rajiv Krishnakumar, Vivek Krishnan, Kevin Krsulich, Prasad Kumkar, Gawel Kus, Ryan LaRose, Enrique Lacal, Raphaël Lambert, John Lapeyre, Joe Latone, Scott Lawrence, Christina Lee, Gushu Li, Dennis Liu, Peng Liu, Yunho Maeng, Kahan Majmudar, Aleksei Malyshev, Joshua Manela, Jakub Marecek, Manoel Marques, Dmitri Maslov, Dolph Mathews, Atsushi Matsuo, Douglas T. McClure, Cameron McGarry, David McKay, Dan McPherson, Srujan Meesala, Thomas Metcalfe, Martin Mevissen, Andrew Meyer, Antonio Mezzacapo, Rohit Midha, Zlatko Mineev, Abby Mitchell, Nikolaj Moll, Jhon Montanez, Michael Duane Mooring, Renier Morales, Niall Moran, Mario Motta, MrF, Prakash Murali, Jan Müggenburg, David Nadlinger, Ken Nakanishi, Giacomo Nannicini, Paul Nation, Edwin Navarro, Yehuda Naveh, Scott Wyman Neagle, Patrick Neuweiler, Johan Nicander, Pradeep Niroula, Hassi Norlen, NuoWenLei, Lee James O'Riordan, Oluwatobi Ogunbayo, Pauline Ollitrault, Raul Otaolea, Steven Oud, Dan Padilha, Hanhee Paik, Soham Pal, Yuchen Pang, Simone Perriello, Anna Phan, Francesco Piro, Marco Pistoia, Christophe Piveteau, Pierre Pocreau, Alejandro Pozas-Kerstjens, Viktor Prutyayov, Daniel Puzzuoli, Jesús Pérez, Quintiii, Rafael Iqbal Rahman, Arun Raja, Nipun Ramagiri, Anirudh Rao, Rudy Raymond, Rafael Martín-Cuevas Redondo, Max Reuter, Julia Rice, Marcello La Rocca, Diego M. Rodriguez, RohithKarur, Max Rossmannek, Mingi Ryu, Tharmashastha SAPV, SamFerracin, Martin Sandberg, Hirmay Sandesara, Ritvik Sapra, Hayk Sargsyan, Aniruddha Sarkar, Ninad Sathaye, Bruno Schmitt, Chris Schnabel, Zachary Schoenfeld, Travis L. Scholten, Eddie Schoute, Joachim Schwarm, Ismael Faro Sertage, Kanav Setia, Nathan Shammah, Yunong Shi, Adenilton Silva, Andrea Simonetto, Nick Singstock, Yukio Siraichi, Iskandar Sitdikov, Seyon Sivarajah, Magnus Berg Sletfjerding, John A. Smolin, Mathias Soeken, Igor Olegovich Sokolov, Igor Sokolov, SooluThomas, Starfish, Dominik Steenzen, Matt Stypulkoski, Shaojun Sun, Kevin J. Sung, Hitomi Takahashi, Tanvish Takawale, Ivano Tavernelli, Charles Taylor, Pete T aylour, Soolu Thomas, Mathieu Tillet, Maddy Tod, Miroslav Tomasik, Enrique de la Torre, Kenso Trabling, Matthew Treinish, TrishaPe, Davindra Tulsi, Wes Turner, Yotam Vaknin, Carmen Recio Valcarce, Francois Varchon, Alमुdena Carrera Vazquez, Victor Villar, Desiree Vogt-Lee, Christophe Vuillot, James Weaver, Johannes Weidenfeller, Rafal Wieczorek, Jonathan A. Wildstrom, Erick Winston, Jack J. Woehr, Stefan Woerner, Ryan Woo, Christopher J. Wood, Ryan Wood, Stephen Wood, Steve Wood, James Wootton, Daniyar Yeralin, David Yonge-Mallo, Richard Young, Jessie Yu, Christopher Zachow, Laura Zdanski, Helena Zhang, Christa Zoufal, Zoufal, a kapila, a matsuo, bcamorrison, brandhsn, nick bronn, chlorophyll zz, dekel.meirom, dekelmeirom, dekoool, dime10, drholmie, dtrenev, ehchen, elfr-campeador, faisaldebouni, fanizzamarco, gabrieleagl, gadial, galeinston, georgios ts, gruu, hhorii, hykavitha, jagunther, jliu45, jscott2, kanejess, klinvill, krutik2966, kurarr, lerongil, ma5x, merav aharoni, michelle4654, ordmoj, sagar pahwa, rmo-yard, saswati qiskit, scottkelso, sethmerkel, strickroman, sumitpuri, tigerjack, toural, tsura crisaldo, vvilpas, welien, willhbang, yangluh, yotamvakninibm, and Mantas Čepulkovskis. 2019. Qiskit: An Open-source Framework for Quantum Computing. <https://doi.org/10.5281/zenodo.2562110>
- [6] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. 2017. Quantum machine learning. *Nature* 549, 7671 (2017), 195–202.
- [7] M J Biercuk, A C Doherty, and H Uys. 2011. Dynamical decoupling sequence construction as a filter-design problem. *Journal of Physics B: Atomic, Molecular and Optical Physics* 44, 15 (Jul 2011), 154002. <https://doi.org/10.1088/0953-4075/44/15/154002>
- [8] Sergey Bravyi, Sarah Sheldon, Abhinav Kandala, David C McKay, and Jay M Gambetta. 2021. Mitigating measurement errors in multiqubit experiments. *Physical Review A* 103, 4 (2021), 042605.
- [9] JJ Burnett, A Bengtsson, M Scigliuzzo, D Niepce, M Kudra, P Delsing, and J Bylander. 2019. Decoherence benchmarking of superconducting qubits. *npj Quantum Inf.* 5, (2019).
- [10] Yongshan Ding, Pranav Gokhale, Sophia Fuhui Lin, Richard Rines, Thomas Propson, and Frederic T Chong. 2020. Systematic Crosstalk Mitigation for Superconducting Qubits via Frequency-Aware Compilation. *arXiv preprint arXiv:2008.09503* (2020).
- [11] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm. *arXiv:1411.4028 [quant-ph]*
- [12] Jay M Gambetta, Jerry M Chow, and Matthias Steffen. 2017. Building logical qubits in a superconducting quantum computing system. *npj Quantum Information* 3, 1 (2017), 1–7.
- [13] Tudor Giurgica-Tiron, Yousef Hindy, Ryan LaRose, Andrea Mari, and William J Zeng. 2020. Digital zero noise extrapolation for quantum error mitigation. In *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 306–316.
- [14] Pranav Gokhale, Yongshan Ding, Thomas Propson, Christopher Winkler, Nelson Leung, Yunong Shi, David I. Schuster, Henry Hoffmann, and Frederic T. Chong. 2019. Partial Compilation of Variational Algorithms for Noisy Intermediate-Scale Quantum Machines. *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture* (Oct 2019). <https://doi.org/10.1145/3352460.3358313>
- [15] E. L. Hahn. 1950. Spin Echoes. *Phys. Rev.* 80 (Nov 1950), 580–594. Issue 4. <https://doi.org/10.1103/PhysRev.80.580>
- [16] Petar Jurcevic, Ali Javadi-Abhari, Lev S Bishop, Isaac Lauer, Daniela Borgorin, Markus Brink, Lauren Capelluto, Oktay Gunluk, Toshinari Itoko, Naoki Kanazawa, et al. 2021. Demonstration of quantum volume 64 on a superconducting quantum computing system. *Quantum Science and Technology* (2021).
- [17] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. 2017. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* 549, 7671 (2017), 242–246.

- [18] Kaveh Khodjasteh and Daniel A. Lidar. 2007. Performance of deterministic dynamical decoupling schemes: Concatenated and periodic pulse sequences. *Physical Review A* 75, 6 (Jun 2007). <https://doi.org/10.1103/physreva.75.062310>
- [19] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Yu Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, Erik Lucero, J. Y. Mutus, O. Naaman, C. Neill, C. Quintana, P. Roushan, Daniel Sank, A. Vainsencher, J. Wenner, T. C. White, S. Boixo, R. Babbush, V. N. Smelyanskiy, H. Neven, and John M. Martinis. 2018. Fluctuations of Energy-Relaxation Times in Superconducting Qubits. *Physical Review Letters* 121, 9 (aug 2018). <https://doi.org/10.1103/physrevlett.121.090502>
- [20] Ryan LaRose, Andrea Mari, Sarah Kaiser, Peter J. Karalekas, Andre A. Alves, Piotr Czarnik, Mohamed El Mandouh, Max H. Gordon, Yousef Hindy, Aaron Robertson, Purva Thakre, Nathan Shammah, and William J. Zeng. 2021. Mitiq: A software package for error mitigation on noisy quantum computers. [arXiv:2009.04417](https://arxiv.org/abs/2009.04417) [quant-ph]
- [21] Wim Lavrijsen, Ana Tudor, Juliane Müller, Costin Iancu, and Wibe de Jong. 2020. Classical Optimizers for Noisy Intermediate-Scale Quantum Devices. In *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 267–277. <https://doi.org/10.1109/QCE49297.2020.00041>
- [22] Ying Li and Simon C. Benjamin. 2017. Efficient Variational Quantum Simulator Incorporating Active Error Minimization. *Phys. Rev. X* 7 (Jun 2017), 021050. Issue 2. <https://doi.org/10.1103/PhysRevX.7.021050>
- [23] Angus Lowe, Max Hunter Gordon, Piotr Czarnik, Andrew Arrasmith, Patrick J Coles, and Lukasz Cincio. 2020. Unified approach to data-driven quantum error mitigation. *arXiv preprint arXiv:2011.01157* (2020). <https://arxiv.org/abs/2011.01157>
- [24] John M Martinis, Ken B Cooper, Robert McDermott, Matthias Steffen, Markus Ansmann, KD Osborn, Katarina Cicak, Seongshik Oh, David P Pappas, Raymond W Simmonds, et al. 2005. Decoherence in Josephson qubits from dielectric loss. *Physical review letters* 95, 21 (2005), 210503.
- [25] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. 2016. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics* 18, 2 (2016), 023023.
- [26] Nikolaj Moll, Panagiotis Barkoutsos, Lev S Bishop, Jerry M Chow, Andrew Cross, Daniel J Egger, Stefan Filipp, Andreas Fuhrer, Jay M Gambetta, Marc Ganzhorn, et al. 2018. Quantum optimization using variational algorithms on near-term quantum devices. *Quantum Science and Technology* 3, 3 (2018), 030503.
- [27] Clemens Müller, Jared H Cole, and Jürgen Lisenfeld. 2019. Towards understanding two-level-systems in amorphous solids: insights from quantum circuits. *Reports on Progress in Physics* 82, 12 (2019), 124501.
- [28] Prakash Murali, Jonathan M Baker, Ali Javadi-Abhari, Frederic T Chong, and Margaret Martonosi. 2019. Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. 1015–1029.
- [29] Prakash Murali, David C McKay, Margaret Martonosi, and Ali Javadi-Abhari. 2020. Software mitigation of crosstalk on noisy intermediate-scale quantum computers. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 1001–1016.
- [30] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'brien. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature communications* 5 (2014), 4213.
- [31] Bibek Pokharel, Namit Anand, Benjamin Fortman, and Daniel A Lidar. 2018. Demonstration of fidelity improvement using dynamical decoupling with superconducting qubits. *Physical review letters* 121, 22 (2018), 220502.
- [32] Bibek Pokharel, Namit Anand, Benjamin Fortman, and Daniel A. Lidar. 2018. Demonstration of Fidelity Improvement Using Dynamical Decoupling with Superconducting Qubits. *Physical Review Letters* 121, 22 (Nov 2018). <https://doi.org/10.1103/physrevlett.121.220502>
- [33] John Preskill. 2018. Quantum Computing in the NISQ era and beyond. *Quantum* 2 (2018), 79.
- [34] Gokul Subramanian Ravi, Kaitlin N. Smith, Pranav Gokhale, Andrea Mari, Nathan Earnest, Ali Javadi-Abhari, and Frederic T. Chong. 2021. VAQEM: A Variational Approach to Quantum Error Mitigation. [arXiv:2112.05821](https://arxiv.org/abs/2112.05821) [quant-ph]
- [35] F.C. Robey, D.R. Fuhrmann, E.J. Kelly, and R. Nitzberg. 1992. A CFAR adaptive matched filter detector. *IEEE Trans. Aerospace Electron. Systems* 28, 1 (1992), 208–216. <https://doi.org/10.1109/7.135446>
- [36] Steffen Schlör, Jürgen Lisenfeld, Clemens Müller, Alexander Bilmes, Andre Schneider, David P Pappas, Alexey V Ustinov, and Martin Weides. 2019. Correlating decoherence in transmon qubits: Low frequency noise by single fluctuators. *Physical review letters* 123, 19 (2019), 190502.
- [37] Kaitlin N Smith, Gokul Subramanian Ravi, Prakash Murali, Jonathan M Baker, Nathan Earnest, Ali Javadi-Abhari, and Frederic T Chong. 2021. Error Mitigation in Quantum Computers through Instruction Scheduling. *arXiv preprint arXiv:2105.01760* (2021).
- [38] Alexandre M Souza, Gonzalo A Álvarez, and Dieter Suter. 2012. Robust dynamical decoupling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 370, 1976 (2012), 4748–4769.
- [39] Swamit S Tannu and Moinuddin K Qureshi. 2019. Mitigating measurement errors in quantum computers by exploiting state-dependent bias. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*. 279–290.
- [40] Swamit S Tannu and Moinuddin K Qureshi. 2019. Not all qubits are created equal: a case for variability-aware policies for NISQ-era quantum computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. 987–999.
- [41] Kristan Temme, Sergey Bravyi, and Jay M Gambetta. 2017. Error mitigation for short-depth quantum circuits. *Physical review letters* 119, 18 (2017), 180509.
- [42] AV Uvarov, AS Kardashin, and Jacob D Biamonte. 2020. Machine learning phase transitions with a quantum processor. *Physical Review A* 102, 1 (2020), 012415.
- [43] Steven Vancoillie, Per Åke Malmqvist, and Valera Veryazov. 2016. Potential energy surface of the chromium dimer re-revisited with multiconfigurational perturbation theory. *Journal of chemical theory and computation* 12, 4 (2016), 1647–1655.
- [44] Lorenza Viola, Emanuel Knill, and Seth Lloyd. 1999. Dynamical decoupling of open quantum systems. *Physical Review Letters* 82, 12 (1999), 2417.
- [45] Greg Welch, Gary Bishop, et al. 1995. An introduction to the Kalman filter. (1995).

Received 2022-07-07; accepted 2022-09-22