

Online Convex Optimization of Programmable Quantum Computers to Simulate Time-Varying Quantum Channels

Hari Hara Suthan Chittoor, *Member, IEEE*, Osvaldo Simeone, *Fellow, IEEE*,
Leonardo Banchi and Stefano Pirandola

Abstract—Simulating quantum channels is a fundamental primitive in quantum computing, since quantum channels define general (trace-preserving) quantum operations. An arbitrary quantum channel cannot be exactly simulated using a finite-dimensional programmable quantum processor, making it important to develop optimal approximate simulation techniques. In this paper, we study the challenging setting in which the channel to be simulated varies adversarially with time. We propose the use of matrix exponentiated gradient descent (MEGD), an online convex optimization method, and analytically show that it achieves a sublinear regret in time. Through experiments, we validate the main results for time-varying dephasing channels using a programmable generalized teleportation processor.

Index Terms—Programmable quantum computing, convex optimization, online learning, quantum channel simulation

I. INTRODUCTION

A quantum computer can be programmed to carry out a given functionality in different ways, including the direct engineering of pulse sequences [1], the design of parametric quantum circuits via quantum machine learning [2, 3], the use of adaptive measurements on cluster states [4], and the optimization of a program state operating on a fixed quantum processor. A fundamental result derived in [5] states there is no universal programmable quantum processor that operates with finite-dimensional program states. Since a quantum processor is universal if it can implement any quantum operation, this conclusion implies that the exact simulation of an arbitrary quantum channel on a single programmable quantum processor is impossible. This, in turn, highlights the importance of developing tools for the optimization of quantum programs.

Hari Hara Suthan Chittoor and Osvaldo Simeone are with King’s Communications, Learning, and Information Processing (KCLIP) lab at the Department of Engineering of Kings College London, UK (emails: hari.hara@kcl.ac.uk, osvaldo.simeone@kcl.ac.uk). Their work has been supported by the European Research Council (ERC) under the European Union’s Horizon 2020 Research and Innovation Programme (Grant Agreement No. 725731), and Osvaldo Simeone has also been supported by an Open Fellowship of the EPSRC (EP/W024101/1). For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising. The authors acknowledge use of the research computing facility at King’s College London, Rosalind (<https://rosalind.kcl.ac.uk>).

Leonardo Banchi is with the Department of Physics and Astronomy, University of Florence & INFN sezione di Firenze, via G. Sansone 1, I-50019 Sesto Fiorentino (FI), Italy (email: leonardo.banchi@unifi.it). His work is supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Superconducting Quantum Materials and Systems Center (SQMS) under the contract No. DE-AC02-07CH11359.

Stefano Pirandola is with the Department of Computer Science, University of York, York YO10 5GH, UK (email: stefano.pirandola@york.ac.uk).




Fig. 1. Time-varying quantum channel \mathcal{E}^t (top) and its simulation \mathcal{E}_{π^t} via a programmable quantum processor \mathcal{Q} controlled by the time-varying program state π^t (bottom).

Reference [6] addressed the problem of approximately simulating a quantum channel using a finite-dimensional program state. The authors proved that the error between the target channel and simulated channel, as measured by the diamond distance, as well as other related metrics, is convex in the space of program states. Specifically, the optimal program state can be calculated using semidefinite programming. In this paper, we study the more challenging setting illustrated in Fig. 1, in which the channel to be simulated varies over time. We adopt a worst-case formulation in which channel variation are arbitrary, and chosen by “nature” in a possibly adversarial way.

To study this setting, we propose to adopt the framework of online convex optimization [7], which provides tools to track the optimal solution of time-varying convex problems. We specifically develop and analyze an online mirror descent algorithm over the space of positive definite matrices, yielding a matrix exponentiated gradient descent (MEGD) [8]. MEGD was previously used in the related, but distinct, problem of quantum state tracking [9, 10].

We prove that the regret of MEGD with respect to an optimized fixed program state is sublinear in time. We conduct experiments by adopting the generalized teleportation processor (GTP) as the programmable quantum processor. GTP can simulate exactly the class of teleportation-covariant channels, modeling Pauli and erasure channels [11], and is operated here in an adversarial setting with time varying dephasing channels. Numerical results validate the analysis.

Notations and Definitions: For any non-negative integer K , $[K]$ represents the set $\{1, 2, \dots, K\}$. The number of elements in a set \mathcal{A} is written as $|\mathcal{A}|$. Given two sets \mathcal{A} and \mathcal{B} , we write $\mathcal{A} \setminus \mathcal{B} = \{a : a \in \mathcal{A} \text{ and } a \notin \mathcal{B}\}$. The symbol \forall represents for all. The Kronecker product is denoted as \otimes ; I represents the identity matrix, with dimensions clear from

the context; M^\dagger represents the complex conjugate transpose of the matrix M ; and $\text{tr}(M)$ represents trace of the matrix M . We adopt standard notations for quantum states, computational basis, and quantum gates [11].

II. PROBLEM FORMULATION

In this section, we first review some background material and then describe the setting and problem of interest.

A. Preliminaries

Throughout this paper, we use the standard Dirac notation (see, e.g., [11]). Given n qubits, we let $\mathcal{D}(\mathcal{H})$ denote the space of all density matrices, i.e., positive semidefinite (PSD) matrices with unit trace, defined on the Hilbert space \mathcal{H} of dimension 2^n . Any $2^n \times 2^n$ Hermitian matrix A can be written in terms of its eigendecomposition $A = \sum_i \lambda_i |v_i\rangle\langle v_i|$, where eigenvalues $\{\lambda_i\}$ are real and the set $\{|v_i\rangle\}$ consists of a basis of orthonormal vectors for the 2^n -dimensional Hilbert space. Furthermore, the square root of a PSD matrix A is defined as $\sqrt{A} = \sum_i \sqrt{\lambda_i} |v_i\rangle\langle v_i|$. More generally, a function $f(A)$ of PSD matrix A is defined as $f(A) = \sum_i f(\lambda_i) |v_i\rangle\langle v_i|$.

A quantum channel \mathcal{E} is a completely positive trace preserving (CPTP) linear map that takes a density matrix $\rho \in \mathcal{D}(\mathcal{H})$ as input to produce a density matrix $\mathcal{E}(\rho) \in \mathcal{D}(\mathcal{H}')$ in a possibly distinct Hilbert space \mathcal{H}' of dimension $2^{n'}$ for an integer n' . Furthermore, given a system of $2n$ qubits, we denote as $I \otimes \mathcal{E}$, where I is the $2^n \times 2^n$ identity matrix, the channel that acts trivially on the first n qubits and applies channel \mathcal{E} to the last n qubits.

A quantum channel \mathcal{E} can be equivalently described by the $2^{2n} \times 2^{2n}$ PSD matrix obtained as the output of channel $I \otimes \mathcal{E}$ applied to a system of $2n$ qubits in the Bell state $|\Phi^+\rangle = 1/\sqrt{2^n} \sum_{i=0}^{2^n-1} |i\rangle \otimes |i\rangle$. This matrix, known as Choi matrix of the quantum channel \mathcal{E} , is hence defined as

$$C_{\mathcal{E}} = (I \otimes \mathcal{E})|\Phi^+\rangle\langle\Phi^+| = \frac{1}{2^n} \sum_{i=0}^{2^n-1} \sum_{j=0}^{2^n-1} |i\rangle\langle j| \otimes \mathcal{E}(|i\rangle\langle j|). \quad (1)$$

B. Setting

As shown in Fig. 1, we study the problem of simulating a time-varying quantum channel \mathcal{E}^t operating on n qubits using a programmable quantum processor \mathcal{Q} , where t is a discrete time index $t = 1, 2, \dots$. Specifically, the top part of Fig. 2 illustrates the Choi matrix $C_{\mathcal{E}^t}$ of the quantum channel \mathcal{E}^t . As depicted in the bottom part of Fig. 2, the programmable quantum processor is a fixed CPTP map \mathcal{Q} operating on a register of $n + n_\pi$ qubits. Examples of quantum processors \mathcal{Q} include generalized teleportation processor [12]–[14] and the port-based teleportation processor [15, 16]. The quantum processor \mathcal{Q} is “programmable” via a time-varying program state $\pi^t \in \mathcal{D}(\mathcal{H}_\pi)$, where \mathcal{H}_π is a 2^{n_π} -dimensional Hilbert space.

The register of $n + n_\pi$ qubits on which processor \mathcal{Q} operates is initially in state $\rho \otimes \pi^t$, where $\rho \in \mathcal{D}(\mathcal{H})$ is the input density and π^t the program state. After the application of processor




Fig. 2. Illustration of the Choi matrix $C_{\mathcal{E}^t}$ of a quantum channel \mathcal{E}^t (top) and of the Choi matrix $C_{\mathcal{E}_{\pi^t}} = C_{\pi^t}$ of its simulation \mathcal{E}_{π^t} via a programmable quantum processor \mathcal{Q} controlled by the program state π^t (bottom).

\mathcal{Q} , we wish for the reduced state of a subset of n qubits within the register of $n + n_\pi$ qubits to approximate the output state

$$\rho_{\mathcal{E}^t} = \mathcal{E}^t(\rho) \quad (2)$$

of channel \mathcal{E}^t for any possible input state ρ . We refer to the mentioned subset of n qubits as forming the output subregister. This may differ from the subregister consisting of the first n qubits initially in the input state ρ (see Sec. IV for an example).

To simulate the channel \mathcal{E}^t , we optimize the sequence of program states π^t sequentially over time $t \in [T]$. As we will detail in Section II-C, at each time t , the optimizer has access to information about the quality of the approximation of the channels \mathcal{E}^τ obtained with program states π^τ at previous times $\tau \in [t - 1]$.

Accordingly, given a program state $\pi^t \in \mathcal{D}(\mathcal{H}_\pi)$, the programmable quantum processor \mathcal{Q} implements the channel \mathcal{E}_{π^t} defined by the mapping

$$\mathcal{E}_{\pi^t}(\rho) = \text{tr}_\pi(\mathcal{Q}(\rho \otimes \pi^t)), \quad (3)$$

where $\text{tr}_\pi(\cdot)$ is the partial trace over the n_π qubits outside the output register. By (3), as illustrated in Fig. 1, these n_π qubits are discarded after the application of the operation \mathcal{Q} .

The simulation error at time t is a measure of the difference between the channel \mathcal{E}^t and the simulated channel \mathcal{E}_{π^t} . As illustrated in Fig. 2, this error can be measured by comparing the corresponding Choi matrices $C_{\mathcal{E}^t}$ and $C_{\mathcal{E}_{\pi^t}} = C_{\pi^t}$. Following [6], the loss is specifically quantified by either the trace distance

$$\ell_1(\mathcal{E}^t, \pi^t) = \frac{1}{2} \|C_{\mathcal{E}^t} - C_{\pi^t}\|_{\text{tr}}, \quad (4)$$

where $\|O\|_{\text{tr}} = \text{tr}(\sqrt{O^\dagger O})$ is the trace of positive square root of matrix $O^\dagger O$; or alternatively, by the complement of the squared fidelity

$$\ell_F(\mathcal{E}^t, \pi^t) = 1 - \left(\text{tr} \left(\sqrt{\sqrt{C_{\mathcal{E}^t}} C_{\pi^t} \sqrt{C_{\mathcal{E}^t}}} \right) \right)^2. \quad (5)$$

We write $\ell(\mathcal{E}^t, \pi^t)$ to denote either loss (4) or (5).

In [6], the authors studied the problem

$$\min_{\pi \in \mathcal{D}(\mathcal{H}_\pi)} \ell(\mathcal{E}, \pi) \quad (6)$$

of optimizing the program state π given a fixed quantum channel \mathcal{E} . Reference [6] proved that the optimization problem (6) is convex over the program states π . Note that the work [6] considered also the diamond distance, which can be upper bounded via the loss functions (4) and (5).

C. Problem Definition

Unlike [6], we consider the problem of online optimization of the program state π^t for time-varying channels \mathcal{E}^t over time $t \in [T]$. We formulate the problem in an adversarial setting to obtain worst-case performance guarantees in terms of the possible sequence of channels $\mathcal{E}^{[T]} = \mathcal{E}^1, \mathcal{E}^2, \dots, \mathcal{E}^T$. Accordingly, at every time $t \in [T]$, the optimizer produces a program state π^t . Then, a quantum channel \mathcal{E}^t is selected in an arbitrary way by “nature”, and the optimizer pays the loss $\ell(\mathcal{E}^t, \pi^t)$, which measures how poorly the simulated channel (3) obtained with program state π^t approximates channel \mathcal{E}^t . For every time $t \in [T]$, the optimizer produces an updated program state π^{t+1} assuming access to a subgradient of the current loss, i.e., $g(\mathcal{E}^t, \pi^t) = \nabla_{\pi} \ell(\mathcal{E}^t, \pi^t)$.

Let us define the policy followed by the optimizer as the function

$$\pi^{t+1} = f^t(\pi^t, g(\mathcal{E}^t, \pi^t)), \quad (7)$$

which maps the current program state π^t and subgradient $g(\mathcal{E}^t, \pi^t)$ to the next program state π^{t+1} . The goal is to design a sequence of functions $f^{[T]}(\cdot) = f^1(\cdot), f^2(\cdot), \dots, f^T(\cdot)$ that performs well when compared to a fixed reference program state π^* optimized based on knowledge of the sequence $\mathcal{E}^{[T]}$ of channels, i.e.,

$$\pi^* = \arg \min_{\pi \in \mathcal{D}(\mathcal{H}_{\pi})} \sum_{t=1}^T \ell(\mathcal{E}^t, \pi). \quad (8)$$

For a sequence of channels $\mathcal{E}^{[T]}$, the performance of policy $f^{[T]}(\cdot)$ relative to the program state in π^* (8) is hence defined by the regret

$$\text{Regret}(f^{[T]}, \mathcal{E}^{[T]}) = \sum_{t=1}^T \ell(\mathcal{E}^t, \pi^t) - \sum_{t=1}^T \ell(\mathcal{E}^t, \pi^*), \quad (9)$$

with programs π^t and π^* given in (7) and (8) respectively.

To elaborate on the regret as the performance criterion of interest, observe first that, if the channel sequence $\mathcal{E}^{[T]}$ to be simulated were constant, i.e., if we had $\mathcal{E}^t = \mathcal{E}$ for some channel \mathcal{E} , obtaining a per-step regret $\ell(\mathcal{E}, \pi^t) - \ell(\mathcal{E}, \pi^*)$ that decreases with t would indicate that the optimizer (7) converges to the program that best approximates the channel in terms of the loss function $\ell(\mathcal{E}, \pi)$. In the online setting under study, the channels \mathcal{E}^t are allowed to vary arbitrarily, and the goal of the optimizer is to track such changes as they occur, i.e., as t increases. Obtaining a small regret, irrespective of the channel sequence, provides evidence that the optimizer is extracting useful information about the single program π^* that would have been optimal in hindsight. Specifically, following the standard online optimization framework [7], we are interested in designing a policy $f^{[T]}$ that achieves a regret that grows sublinearly in T , implying that the per-step regret $\ell(\mathcal{E}^t, \pi^t) - \ell(\mathcal{E}^t, \pi^*)$ decreases over time t .

III. MATRIX EXPONENTIATED GRADIENT DESCENT

In this section, we propose a policy for the problem of on-line channel simulation introduced in the previous section that is based on matrix exponentiated gradient descent (MEGD) [8]. We also analyze its regret, showing that it is sublinear in T .

A. Matrix Exponentiated Gradient Descent (MEGD) for On-line Channel Simulation

The proposed MEGD algorithm is an online mirror descent algorithm [7, 8] over the space of PSD matrices with unit trace. MEGD initializes the program state as the maximally mixed state $\pi^1 = I/2^{n_{\pi}}$. For every time $t \in [T]$ an arbitrary channel \mathcal{E}^t is selected by nature, and the optimizer obtains the loss $\ell(\mathcal{E}^t, \pi^t)$. Based on the corresponding subgradient $g(\mathcal{E}^t, \pi^t)$, the optimizer updates the program state as

$$\pi^{t+1} = \frac{\exp(Z^t)}{\text{tr}(\exp(Z^t))}, \quad (10)$$

with matrix

$$Z^t = \log(\pi^t) - \eta \tilde{g}(\mathcal{E}^t, \pi^t), \quad (11)$$

where $\tilde{g}(\mathcal{E}^t, \pi^t) = (g(\mathcal{E}^t, \pi^t) + g(\mathcal{E}^t, \pi^t)^{\dagger})/2$ represents the Hermitian part of the subgradient $g(\mathcal{E}^t, \pi^t)$ and $\eta > 0$ is the learning rate.

To evaluate the subgradients $g(\mathcal{E}^t, \pi^t)$ for the losses (4) and (5), we first define the quantum channel Λ^t that maps program state π^t to the corresponding Choi matrix C_{π^t} . This channel can be specified by its Kraus decomposition $\Lambda^t(\pi) = \sum_i A_i \pi A_i^{\dagger}$, where the $2^{2n} \times 2^n$ Kraus operators $\{A_i\}$ satisfy the condition $\sum_i A_i^{\dagger} A_i = I$, with I being the $2^n \times 2^n$ identity matrix. The dual channel is defined as $\Lambda_*^t(\rho) = \sum_i A_i^{\dagger} \rho A_i$. Furthermore, we write the eigendecomposition of the Hermitian error operator $C_{\pi^t} - C_{\mathcal{E}^t}$ as $C_{\pi^t} - C_{\mathcal{E}^t} = \sum_i \lambda_i^t E_i^t$, with real eigenvalues $\{\lambda_i^t\}$ and eigenprojectors $\{E_i^t\}$. Following [6, Theorem 2], the subgradient with respect to PSD matrix π of the loss functions $\ell_1(\mathcal{E}^t, \pi)$ and $\ell_F(\mathcal{E}^t, \pi)$ evaluated at $\pi = \pi^t$ are given as

$$g_1(\mathcal{E}^t, \pi^t) = \sum_i \text{sign}(\lambda_i^t) \Lambda_*^t(E_i^t), \quad (12)$$

$$\text{and } g_F(\mathcal{E}^t, \pi^t) = -\sqrt{1 - \ell_F(\mathcal{E}^t, \pi^t)} \nabla L(\pi^t), \quad (13)$$

respectively, where $\text{sign}(x) = 1$ if $x \geq 0$ and $\text{sign}(x) = -1$ if $x < 0$, and

$$\nabla L(\pi^t) = \Lambda_*^t \left(\sqrt{C_{\mathcal{E}^t}} \left(\sqrt{C_{\mathcal{E}^t}} \Lambda^t(\pi^t) \sqrt{C_{\mathcal{E}^t}} \right)^{-\frac{1}{2}} \sqrt{C_{\mathcal{E}^t}} \right).$$

Finally, for numerical stability, the MEGD update (10) is implemented by replacing matrix Z^t with the time-unrolled update [8]

$$Z^t = d^t I + \log(\pi^1) - \eta \sum_{\tau=1}^t \tilde{g}(\mathcal{E}^{\tau}, \pi^{\tau}), \quad (14)$$

where $\{d^t\}$ is a sequence of fixed constants. These constants do not affect an infinite-precision implementation of update (10) with (14), but they can be useful to avoid numerical problems [8]. MEGD is summarized in Algorithm 1.

Algorithm 1 Matrix exponentiated gradient descent (MEGD) for online channel simulation

- 1: **Require:** Learning rate $\eta > 0$ and sequence of constants $d^t > 0$ for $t \in [T]$.
 - 2: Initialize the program state as $\pi^1 = I/2^{n_\pi}$
 - 3: **for** $t \in [T]$ **do**
 - 4: Adversary selects a new channel \mathcal{E}^t
 - 5: Optimizer obtains the loss $\ell(\mathcal{E}^t, \pi^t)$ using (4) or (5)
 - 6: Optimizer computes subgradient $g(\mathcal{E}^t, \pi^t)$ of the loss $\ell(\mathcal{E}^t, \pi^t)$ at π^t using (12) or (13)
 - 7: Optimizer updates the program state to π^{t+1} using (10) with (14)
 - 8: **end for**
-

B. Regret Analysis

To show that MEGD achieves sublinear regret in T , we start by interpreting the update rule (10) in terms of a regularized optimization problem that follows the mirror descent framework [7, 8, 17]. To this end, we introduce the negative von Neumann entropy, defined as

$$F(\pi) = \text{tr}(\pi \ln(\pi)) \quad (15)$$

for any state π . The Bregman divergence generated by function $F(\cdot)$ is given by

$$B_F(\pi_1; \pi_2) = F(\pi_1) - F(\pi_2) - \text{tr}[\nabla F(\pi_2)(\pi_1 - \pi_2)] = \text{tr}(\pi_1 \ln(\pi_1) - \pi_1 \ln(\pi_2)), \quad (16)$$

and it corresponds to the quantum relative entropy [11] between two density states π_1 and π_2 , defined on the same Hilbert space. Following [7, 8], the update rule (10) arises as the solution of the optimization problem

$$\pi^{t+1} = \arg \min_{\pi \in \mathcal{D}(\mathcal{H}_\pi)} (\eta \text{tr}[\pi \tilde{g}(\mathcal{E}^t, \pi^t)] + B_F(\pi; \pi^t)), \quad (17)$$

where the first term is a linearization of the per-step loss and the second is a regularizer penalizing deviations from the current program π^t . To formulate the main result in Proposition 1 we also introduce the spectral norm $\|\cdot\|_*$ of a matrix O as the square root of the largest eigenvalue of the matrix $O^\dagger O$.

Proposition 1. *The regret (9) of MEGD is upper bounded as*

$$\text{Regret}(f^{[T]}, \mathcal{E}^{[T]}) \leq \frac{B_F(\pi^*; \pi^1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \|g(\mathcal{E}^t, \pi^t)\|_*^2, \quad (18)$$

where subgradient $g(\mathcal{E}^t, \pi^t)$ and reference program state π^* are defined in (12)-(13) and (8), respectively. Furthermore, if the spectral norm of the subgradient is bounded as $\|g(\mathcal{E}^t, \pi^t)\|_* \leq L_*$ for every $t \in [T]$, choosing the learning rate as $\eta = \sqrt{2 \ln(2) n_\pi / T L_*^2}$, the following regret bound holds

$$\text{Regret}(f^{[T]}, \mathcal{E}^{[T]}) \leq L_* \sqrt{2 \ln(2) n_\pi T}. \quad (19)$$

Proof. Since the negative von Neumann entropy $F(\pi)$ is 1-strongly convex over the space of density states $\pi \in \mathcal{D}(\mathcal{H}_\pi)$

with respect to the trace norm [18], inequality (18) is a consequence of [7, Theorem 6.8]. To prove (19), we first upper bound $B_F(\pi^*; \pi^1)$ as

$$B_F(\pi^*; \pi^1) = F(\pi^*) + \ln(2^{n_\pi}) \leq \ln(2^{n_\pi}), \quad (20)$$

where the inequality holds from the non-negativity of the von Neumann entropy $-F(\cdot)$ [11]. Using (20) in (18), the regret can be upper bounded as

$$\text{Regret}(f^{[T]}, \mathcal{E}^{[T]}) \leq \frac{\ln(2) n_\pi}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \|g(\mathcal{E}^t, \pi^t)\|_*^2. \quad (21)$$

Finally, inequality (19) follows by using the assumed inequality $\|g(\mathcal{E}^t, \pi^t)\|_* \leq L_*$ and by selecting the learning rate as indicated in the proposition. \square

IV. EXPERIMENTS

In this Section, we provide experimental results to validate the proposed MEGD scheme in Algorithm 1. We start by describing the generalized teleportation processor (GTP) which will be adopted as the programmable quantum processor \mathcal{Q} .

A. Generalized Teleportation Processor (GTP)




Fig. 3. Generalized teleportation processor as a programmable processor \mathcal{Q} operating on one input qubit ($n = 1$) and on a two-qubit program state π ($n_\pi = 2$).

As illustrated in Fig. 3, the GTP operates on a register of three qubits with $n = 1$ and $n_\pi = 2$. A Bell measurement, defined by the projectors $\{P_0, P_1, P_2, P_3\}$, where

$$P_0 = |\Phi^+\rangle\langle\Phi^+|, P_1 = |\Psi^+\rangle\langle\Psi^+|, P_2 = |\Psi^-\rangle\langle\Psi^-|, P_3 = |\Phi^-\rangle\langle\Phi^-|,$$

with $\{|\Phi^+\rangle, |\Psi^+\rangle, |\Phi^-\rangle, |\Psi^-\rangle\}$ being the standard Bell states (see [11, Section 1.3.6]), is applied to the input qubit and to the first control qubit. Then, depending on the output $k \in \{0, 1, 2, 3\}$ of the measurement, with k corresponding to projector P_k , a unitary correction V_k is applied to the second control qubit, where

$$V_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, V_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, V_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, V_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (22)$$

are the Pauli operators. The last qubit encodes the output state $\mathcal{E}_\pi(\rho)$ of the GTP as shown in Fig. 3.

It is known that the GTP can simulate exactly all teleportation-covariant channels, which include Pauli and erasure channels [19]. This is done by choosing the program state π as the scaled Choi matrix $C_{\mathcal{E}}$ of the channel to be simulated. As a trivial special case, note that setting $\pi = |\Phi^+\rangle\langle\Phi^+|$ simulates the identity channel, since the corresponding Choi matrix is given by $|\Phi^+\rangle\langle\Phi^+|$. Not all channels are teleportation-covariant. For example, the amplitude damping channel is not teleportation-covariant [19], and hence it cannot be exactly simulated using the GTP.

In order to implement Algorithm 1 using the GTP, we need the quantum channel $\Lambda^t(\pi)$ mapping program state π to the corresponding Choi matrix C_π and its dual channel $\Lambda_*^t(\pi)$, for the evaluation of subgradients (12)-(13). For the GTP, these are given by [6]

$$\Lambda^t(\pi) = \Lambda_*^t(\pi) = \frac{1}{4} \sum_{k=0}^3 (V_k^\dagger \otimes V_k) \pi^t (V_k^\dagger \otimes V_k). \quad (23)$$

B. Results

We now study the simulation of the dephasing channel, which is a special case of Pauli channels. Specifically, at each time t , the single-qubit channel to be simulated is given as

$$\mathcal{E}^t(\rho) = (1 - p^t)\rho + p^t Z \rho Z, \quad (24)$$

where $Z = V_3$ represents the Pauli Z operator, and the sequence of probabilities p^1, p^2, \dots defines the sequence of channels (24).

In Fig. 4 we plot the normalized regret $T^{-1}\text{Regret}(f^{[T]}, \mathcal{E}^{[T]})$, where $\text{Regret}(f^{[T]}, \mathcal{E}^{[T]})$ is defined in (9), as a function of time T , by considering the trace distance (4) as the loss function. We consider $T \in [150]$ and use $\eta = 0.01$ as the learning rate and constants $d^t = 2$ for every $t \in [T]$ in (14). For each time window duration T , we obtain the optimal constant reference program π^* in (8) by optimizing the sum of loss functions via the MEGD update rule in (10) over 120 iterations with a learning rate $0.01/T$. We observed numerically that these choices yield convergent iterates.

We consider the setting in which the channel to be simulated changes independently at each time $t \in [T]$, with probability p^t drawn uniformly in the interval $[0.2, p_{max})$, with $p_{max} \in \{0.2, 0.4, 0.6, 0.8\}$. In all cases, we observe that, as stated in Proposition 1, MEGD is able to obtain a normalized regret that decreases sublinearly with T , hence approaching the performance of the reference optimal constant program π^* . Furthermore, as p_{max} increases, the performance of the fixed optimum program π^* decreases and the regret of MEGD is reduced accordingly.

REFERENCES

- [1] G. Wolfowicz and J. J. Morton, "Pulse techniques for quantum information processing," *eMagRes*, vol. 5, no. 4, pp. 1515–1528, 2016.
- [2] M. Schuld and F. Petruccione, *Machine learning with quantum computers*. Springer, 2021.

The PyTorch code for regenerating the results of this paper is available at <https://github.com/kclip/OCOPQC>.




Fig. 4. Normalized regret as a function of time T for MEGD (Algorithm 1) when simulating a time-varying dephasing channel with dephasing probabilities drawn independently and uniformly at each time in the interval $[0.2, p_{max})$ (setting $p_{max} = 0.2$ models a constant channel).

- [3] O. Simeone, "An introduction to quantum machine learning for engineers," *Foundations and Trends® in Signal Processing*, vol. 16, no. 1-2, pp. 1–223, 2022.
- [4] M. A. Nielsen and C. M. Dawson, "Fault-tolerant quantum computation with cluster states," *Phys. Rev. A*, vol. 71, p. 042323, Apr 2005.
- [5] M. A. Nielsen and I. L. Chuang, "Programmable quantum gate arrays," *Phys. Rev. Lett.*, vol. 79, pp. 321–324, Jul 1997.
- [6] L. Banchi, J. Pereira, S. Lloyd, and S. Pirandola, "Convex optimization of programmable quantum computers," *npj Quantum Information*, vol. 6, no. 1, pp. 1–10, 2020.
- [7] F. Orabona, "A modern introduction to online learning," *CoRR*, vol. abs/1912.13213, 2019. [Online]. Available: <http://arxiv.org/abs/1912.13213>
- [8] K. Tsuda, G. Rätsch, and M. K. Warmuth, "Matrix exponentiated gradient updates for on-line learning and Bregman projection," *Journal of Machine Learning Research*, vol. 6, no. 34, pp. 995–1018, 2005.
- [9] S. Aaronson, X. Chen, E. Hazan, S. Kale, and A. Nayak, "Online learning of quantum states," in *Advances in Neural Information Processing Systems*, vol. 31. Curran Associates, Inc., 2018.
- [10] M. Rambach, A. Youssry, M. Tomamichel, and J. Romero, "Efficient quantum state tracking in noisy environments," *arXiv preprint arXiv:2205.06389*, 2022.
- [11] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge University Press, 2010.
- [12] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, "Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels," *Phys. Rev. Lett.*, vol. 70, pp. 1895–1899, Mar 1993.
- [13] S. Pirandola, J. Eisert, C. Weedbrook, A. Furusawa, and S. L. Braunstein, "Advances in quantum teleportation," *Nature photonics*, vol. 9, no. 10, pp. 641–652, 2015.
- [14] R. Laurenza, C. Lupo, G. Spedalieri, S. L. Braunstein, and S. Pirandola, "Channel simulation in quantum metrology," *Quantum Measurements and Quantum Metrology*, vol. 5, no. 1, pp. 1–12, 2018.
- [15] J. Pereira, L. Banchi, and S. Pirandola, "Characterising port-based teleportation as universal simulator of qubit channels," *Journal of Physics A: Mathematical and Theoretical*, vol. 54, no. 20, p. 205301, Apr 2021.
- [16] S. Ishizaka and T. Hiroshima, "Quantum teleportation scheme by selecting one of multiple output ports," *Phys. Rev. A*, vol. 79, p. 042306, Apr 2009.
- [17] O. Simeone, *Machine learning for engineers*. Cambridge University Press, 2022.
- [18] Y.-L. Yu, "The strong convexity of von Neumann's entropy," *Unpublished note*, June, 2013.
- [19] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, "Fundamental limits of repeaterless quantum communications," *Nature communications*, vol. 8, no. 1, pp. 1–15, 2017.