Pre-Distillation of Magic States via Composite Schemes

Muhammad Erew,[1] Moshe Goldstein,[1] Yaron Oz,[1] and Haim Suchowski[1]

[1] *Raymond and Beverly Sackler School of Physics and Astronomy, Tel-Aviv University, Tel-Aviv 6997801, Israel.*

(Dated: October 2, 2025)

Magic state distillation (MSD) is a cornerstone of fault-tolerant quantum computing, enabling non-Clifford gates via state injection into stabilizer circuits. However, the substantial overhead of current MSD protocols remains a major obstacle to scalable implementations. We propose a general framework for *pre-distillation*, based on composite pulse sequences that suppress systematic errors in the generation of magic states. Unlike typical composite designs that target simple gates such as X, Z, or Hadamard, our schemes directly implement the non-Clifford $\mathcal{T}$ gate with enhanced robustness. We develop composite sequences tailored to the dominant control imperfections in superconducting, trapped-ion, neutral-atom, and integrated photonic platforms. To quantify improvement in the implementation, we introduce an operationally motivated fidelity measure specifically tailored to the $\mathcal{T}$ gate: the T-magic error, which captures the gate's effectiveness in preparing high-fidelity magic states. We further show that the error in the channel arising from the injection of faulty magic states scales linearly with the leading-order error of the states. Across all platforms, our approach yields high-fidelity $\mathcal{T}$ gates with reduced noise, lowering the number of distillation levels by up to three. This translates to exponential savings in qubit overhead and offers a practical path toward more resource-efficient universal quantum computation.

I. Introduction

Quantum computing promises computational advantages over classical methods, with celebrated examples such as Shor's factoring and Grover's search algorithms [1, 2]. Yet building scalable machines remains challenging because quantum information is fragile, and error correction is restricted by fundamental no-go theorems. In particular, the Eastin-Knill theorem shows that no quantum error-correcting code admits a fully transversal universal gate set [3, 4], forcing us to seek alternative paths to universality.

A central strategy is to augment the efficiently and fault-tolerantly implementable Clifford operations, realized, for instance, through stabilizer codes, with special non-stabilizer ancilla states known as *magic states*. Injecting these states into Clifford circuits enables the implementation of otherwise inaccessible non-Clifford gates, thereby promoting the Clifford group to universality [4, 5]. Since such states are typically noisy, they must be purified. This is achieved through *magic state distillation* (MSD) [6], a systematic protocol that consumes multiple noisy copies of a state and, via stabilizer code measurements, produces fewer but higher-fidelity copies.

Yet, the overhead of MSD is severe: each distillation level increases the number of physical qubits exponentially [6–8]. For example, preparing a single logical $|T\rangle$ state with error below 10^{-15} may require five to ten levels of recursive distillation, consuming thousands to millions of physical qubits depending on the initial error.

To mitigate this cost, we propose a platform-aware strategy of *pre-distillation* based on robust composite gate constructions that reduce the error in magic states prior to entering a distillation protocol. Unlike most composite schemes, which are designed for elementary Clifford gates like X, Z, or Hadamard [9–16], our sequences are tailored to directly implement the $\mathcal{T}$ gate- the gate that produces $|T\rangle$ when acting on $|0\rangle$. By targeting the gate responsible for magic state generation itself, we reduce the physical noise at the source, thereby improving their fidelity and reducing the overhead of distillation.

We develop and analyze composite pulse schemes for a range of platforms to improve the fidelity of magic state preparation prior to distillation. For systems with X-Y control Hamiltonians, such as superconducting, trapped-ion, and neutral-atom qubits, we construct symmetric three- and five-segment sequences that implement the $\mathcal{T}$ gate while canceling first- and second-order global Rabi frequency errors, respectively. In systems governed by X-Z dynamics, we show how such gates can be synthesized when they are not directly implementable by single-segment, and design composite sequences that achieve robustness against global detuning-like errors. Finally, in integrated photonics, we apply our framework to directional couplers subject to correlated fabrication errors and design four-segment geometries that reduce the magic state error below the distillation threshold. Across all platforms, we quantify the improvement in terms of the number of required distillation iterations and demonstrate reductions of one to three levels, corresponding to orders-of-magnitude savings in physical qubit overhead. To support this analysis, we also introduce a natural, operationally motivated fidelity measure for the $\mathcal{T}$ gate, *magic T-gate fidelity*, which directly reflects the success of magic state preparation and guides the estimation of distillation depth. In this way, our work bridges two previously separate directions- robust gate synthesis and magic state distillation- by showing how pulse-level error suppression translates directly into fewer distillation rounds and reduced resource overhead.

This work is organized as follows. In Section II, we review the theoretical background of MSD and the role of the $\mathcal{T}$ gate in enabling universal fault-tolerant quantum computation. Section III introduces the physical real-

ization of the $\mathcal{T}$ and $\mathcal{H}$ gates across multiple quantum platforms, along with their error models. In Section IV, we construct composite pulse sequences that implement the $\mathcal{T}$ and $\mathcal{H}$ gates in X-Y systems while achieving first- and second-order robustness to global Rabi frequency errors. In Section V, we extend this approach to X-Z systems, providing composite schemes that reduce semi-detuning-induced errors. Section VI applies these ideas to integrated photonic platforms, presenting directional coupler designs that implement the $\mathcal{T}$ gate with enhanced resilience to global fabrication deviations. In Section VII, we show that the channel error induced by faulty magic states increases linearly with their leading-order error, confirming that the error analysis presented earlier is sufficient. Finally, in Section VIII, we conclude with a discussion of implications for scalable fault-tolerant architectures and potential extensions of our pre-distillation framework.

II. Background: Distillation and Pre-Distillation

MSD lies at the heart of most approaches to universal fault-tolerant quantum computation, providing a practical means of suppressing noise in non-Clifford resource states. While the general framework of MSD has been developed in many variants, we illustrate the main principles here through the canonical five-qubit protocol for the T-state, which serves as a representative example. A detailed technical discussion of these preliminaries, along with circuit-level constructions, is deferred to Appendix A.

The canonical magic T-state is defined by $|T\rangle = \cos\beta\,|0\rangle + e^{i\pi/4}\sin\beta\,|1\rangle$, where $\cos^2(2\beta) = \frac{1}{3}$. This state is an eigenvector of the T-gate, $T = e^{i\pi/4}SH$, which its two orthogonal eigenstates are denoted as $|T_0\rangle = |T\rangle$, $|T_1\rangle = -e^{-i\pi/4}\sin\beta\,|0\rangle + \cos\beta\,|1\rangle$.

Figure 1 provides an overview. In Figure 1(a), five noisy input states of the form $\rho_T(\epsilon) = (1-\epsilon)\,|T_0\rangle\langle T_0| + \epsilon\,|T_1\rangle\langle T_1|$ are consumed in a single round of distillation. With probability $p(\epsilon) = \frac{1}{6} + O(\epsilon)$, the protocol outputs a logical state $\rho_T^{(L)}(1-\epsilon')$ whose error is quadratically suppressed, $\epsilon'(\epsilon) = \frac{\epsilon^5+5\epsilon^2(1-\epsilon)^3}{\epsilon^5+5\epsilon^2(1-\epsilon)^3+5\epsilon^3(1-\epsilon)^2+(1-\epsilon)^5} = 5\epsilon^2 + O(\epsilon^3)$. Thus, whenever the input error lies below a critical threshold, $\epsilon_c \approx 0.173$, the protocol successfully suppresses errors, albeit with only probabilistic success and at the cost of additional resource overhead. The technical details of the distillation circuit are in Figure 6 in Appendix A.

To further suppress the error, the protocol must be recursively applied by nesting and concatenating the code: each physical qubit is elevated to a logical qubit, which in turn is encoded into five new physical qubits. Achieving fault-tolerant thresholds therefore demands multiple rounds of recursion, resulting in an exponential growth of physical-qubit overhead per logical qubit. Figures 1(c-d) characterize the performance of the protocol. Figure 1(b)

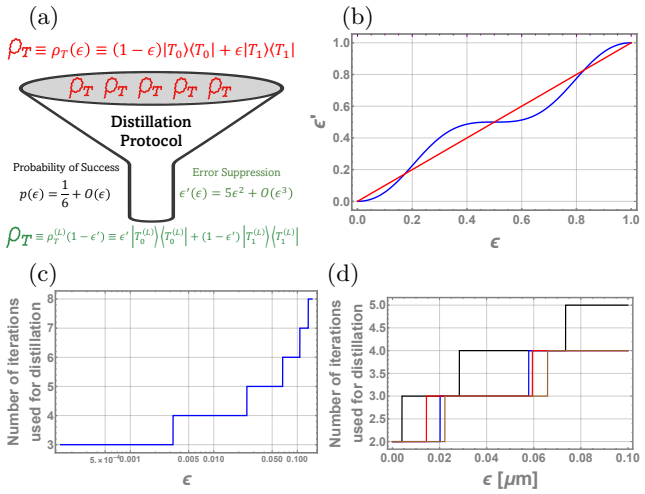

FIG. 1. Five-qubit T-state distillation and its improvement through pre-distillation. (a) Conceptual schematic of a single round of the T-magic-state distillation protocol. One round consumes five noisy copies of $\rho_T(\epsilon) \equiv (1-\epsilon)\,|T_0\rangle\langle T_0| + \epsilon\,|T_1\rangle\langle T_1|$ (denoted by ρ_T with heavily wavy red symbols to emphasize noise). With probability $p(\epsilon) = \frac{1}{6} + O(\epsilon)$, the protocol outputs the logical state $\rho_T^{(L)}(1-\epsilon')$, whose error is quadratically suppressed, $\epsilon'(\epsilon) = 5\epsilon^2 + O(\epsilon^3)$. We represent this improved state using green symbols with reduced waviness, reflecting its lower noise. (b) Distillation curve $\epsilon'(\epsilon)$ (blue): the output error rate as a function of the input error rate, showing quadratic suppression after one round. For reference, the red line shows the identity function $f(\epsilon) = \epsilon$, highlighting the improvement gained over the uncorrected case. (c) Required number of recursive distillation rounds to achieve an error below the 10^{-15} fault-tolerance threshold, illustrating the exponential scaling of qubit overhead with the initial error ϵ. (d) Comparison of $\mathcal{T}$-gate implementations under systematic width errors in directional couplers. The plot shows the number of distillation levels needed to reach 10^{-15} error. Results are shown for a two-segment design (black) and three composite four-pulse schemes: (A) blue, (B) red, (C) brown.

shows the distillation curve $\epsilon'(\epsilon)$, making explicit how the output error depends on the input. Figure 1(c) highlights the recursive structure of MSD: by concatenating multiple rounds, one can reduce the error below the fault-tolerance threshold (here taken to be 10^{-15}). However, the required number of rounds grows rapidly with the initial error, reflecting the exponential qubit overhead inherent to this process.

Therefore, it is more efficient to reduce the error *prior* to encoding, as this can save multiple levels of concatenation and lead to an exponential reduction in the qubit overhead. In this work, we employ composite pulse schemes, traditionally used for implementing gates such as X and Hadamard, to enhance the fidelity of T-state preparation. Figure 1(d) illustrates how *pre-distillation techniques* -in this case, composite pulse sequences designed to suppress systematic width errors in directional couplers- can improve overall efficiency. The figure compares several two- and four-pulse composite schemes by

the number of distillation levels required to reach an error rate of 10^{-15}. While some schemes offer modest advantages only at small errors, others consistently outperform across the full range, thereby reducing the depth of concatenation required for fault-tolerant $\mathcal{T}$-gate implementation.

III. $\mathcal{T}$ and $\mathcal{H}$ Gates and Physical Realizations

In this section, we examine the explicit unitary operations required to generate the magic states $|T\rangle$ and $|H\rangle$ (the eigenstates of the Hadamard gate H) from the stabilizer state $|0\rangle$. We derive the corresponding $\mathcal{T}$ and $\mathcal{H}$ gates and analyze how these unitaries can be physically realized across leading quantum hardware platforms. Emphasis is placed on the structure of the underlying control Hamiltonians and the types of systematic errors that arise in each setting.

The terminology surrounding the T gate and associated magic states can be confusing due to inconsistent naming conventions in the literature. Depending on the context, the T gate may refer to either $T = e^{i\pi/4}SH = \frac{e^{i\pi/4}}{\sqrt{2}}\begin{pmatrix}1 & 1\\ i & -i\end{pmatrix}$, or to the so-called $\pi/4$ gate (or the $\pi/8$ gate historically), $T = \begin{pmatrix}1 & 0\\ 0 & e^{-i\pi/4}\end{pmatrix}$. The $|T\rangle$ states are eigenstates of the former and, when injected into a circuit, enable the fault-tolerant implementation of a $\pi/12$ rotation. This differs from the latter T gate (i.e., the $\pi/4$ gate or the $\pi/8$ gate historically), which is used to realize the Z-axis rotation $\exp\left(-i\frac{\pi}{8}Z\right)$ and is instead implemented via magic state injection of the $|H\rangle$ states-eigenstates of the Hadamard gate H.

Here we define new gates, $\mathcal{T}$ and $\mathcal{H}$, associated with the $|T\rangle$ and $|H\rangle$ states, respectively. These are the unitary operators that map the stabilizer state $|0\rangle$ to $|T\rangle$ and $|H\rangle$, and they diagonalize the corresponding non-Clifford unitaries. To avoid introducing new ambiguities, we use the notation $\mathcal{T}$ and $\mathcal{H}$ rather than overloading existing symbols.

A. $\mathcal{T}$ and $\mathcal{H}$ Gates

To prepare the $|T\rangle$ and $|H\rangle$ magic states, one typically starts with a stabilizer state and applies an appropriate unitary transformation. In this work, we take the initial state to be $|0\rangle$, which we assume can be prepared either fault-tolerantly or with negligible error. Since qubits reside in a two-dimensional Hilbert space, the unitary operator that maps $|0\rangle$ to a target pure state is uniquely determined up to a global phase. The unitaries required to produce $|T\rangle$ and $|H\rangle$ from $|0\rangle$ are given by:

$$\begin{aligned}\mathcal{T} &= \begin{pmatrix}\cos\beta & -e^{-i\pi/4}\sin\beta\\ e^{i\pi/4}\sin\beta & \cos\beta\end{pmatrix}\\ &= \begin{pmatrix}\cos\beta & -ie^{-i\ 3\pi/4}\sin\beta\\ -ie^{i\ 3\pi/4}\sin\beta & \cos\beta\end{pmatrix},\end{aligned} \tag{1a}$$

$$\begin{aligned}\mathcal{H} &= \begin{pmatrix}\cos\left(\frac{\pi}{8}\right) & -\sin\left(\frac{\pi}{8}\right)\\ \sin\left(\frac{\pi}{8}\right) & \cos\left(\frac{\pi}{8}\right)\end{pmatrix}\\ &= \begin{pmatrix}\cos\left(\frac{\pi}{8}\right) & -ie^{-i\pi/2}\sin\left(\frac{\pi}{8}\right)\\ -ie^{i\pi/2}\sin\left(\frac{\pi}{8}\right) & \cos\left(\frac{\pi}{8}\right)\end{pmatrix},\end{aligned} \tag{1b}$$

where $\cos^2(2\beta) = \frac{1}{3}$, respectively.

These gates correspond to specific single-qubit rotations on the Bloch sphere, characterized by nontrivial rotation axes and complex phase factors. Their physical implementation depends on the underlying quantum hardware platform, which dictates the available control mechanisms and error models. In the following subsection (III B), we examine how various quantum computing platforms realize arbitrary single-qubit unitaries, how these controls can be tailored to implement the specific gates $\mathcal{T}$ and $\mathcal{H}$, and how composite pulse techniques can be employed to mitigate common systematic errors. In the next sections (IV,V,VI), we find composite schemes for them.

B. Physical Realizations and Errors

Physical realizations of qubits span several leading experimental platforms, each with distinct control mechanisms, error sources, and robustness strategies. Among the most developed are superconducting circuits, where microwave drives manipulate anharmonic oscillators; trapped ions, where internal states of ions are controlled by laser-driven Raman transitions; neutral atoms in optical tweezers or lattices, where hyperfine ground states are coupled by microwave or Raman fields; and integrated photonic systems, where single photons are guided in waveguide circuits and manipulated through directional couplers. While these platforms differ substantially in hardware implementation, they share common challenges such as systematic amplitude and detuning errors. In many cases, composite pulse or composite device strategies are employed to mitigate these errors in the realization of gates such as $\mathcal{T}$ and $\mathcal{H}$. The following overview highlights the key features and error mechanisms of these platforms, forming the basis for the composite schemes developed in subsequent sections. A more detailed technical discussion is provided in Appendix B.

- **Superconducting qubits:** [17, 18] Transmons implement quantum gates via resonant microwave drives. In the rotating frame, the drive produces rotations around equatorial Bloch-sphere axes, with pulse amplitude and phase determining the rotation angle and direction. Gates like $\mathcal{T}$ and $\mathcal{H}$ are

synthesized through calibrated pulse sequences, often combining R_x and R_z rotations. Dominant errors include amplitude miscalibration, phase errors from control electronics, and leakage into higher excited states. Composite pulse techniques such as BB1 and SK1 mitigate systematic amplitude errors and enhance robustness.

- **Trapped ions:** [19–22] Qubits are encoded in long-lived internal states of ions (e.g., $^{171}Yb^{+}$), with laser-driven Raman transitions enabling high-fidelity gates. The laser amplitude, detuning, and phase directly set the unitary, allowing flexible implementation of $\mathcal{T}$ and $\mathcal{H}$. Common errors stem from laser intensity noise, detuning from resonance, and AC Stark shifts. Composite control methods such as CORPSE or SCROFULOUS cancel systematic errors by distributing rotations into carefully chosen segments.

- **Neutral atoms:** [23, 24] Atoms confined in optical tweezers or lattices encode qubits in hyperfine ground states. Microwave or Raman transitions drive single-qubit rotations, with arbitrary unitaries (including $\mathcal{T}$ and $\mathcal{H}$) constructed from combinations of R_y and R_z operations. Errors arise from intensity fluctuations, frequency drifts, and spatial crosstalk between neighboring atoms. Composite schemes such as BB1, CORPSE, and Walsh-modulated pulses are widely applied to suppress coherent amplitude and detuning errors.

- **Integrated photonics:** [16, 25–28] Qubits are encoded in photon paths, with directional couplers implementing coherent mode mixing. The effective Hamiltonian naturally generates X and Z rotations, while gates involving Y (e.g., $\mathcal{T}$) require multi-segment designs with varied parameters. Photonic devices are limited by static fabrication imperfections in waveguide widths, gaps, and etch depths. Mitigation strategies include post-fabrication tuning, symmetric design (e.g., Mach-Zehnder interferometers), and programmable interferometers. Composite directional coupler designs provide a passive, static route to error resilience, important for scalable and fault-tolerant architectures.

Each platform effectively restricts the set of accessible *generators* of SU(2) in its Hamiltonian. Consequently, one must adopt an appropriate model for both the Hamiltonian and its associated errors in order to accurately describe the system. Consider a general single-qubit gate parameterized by

$$U(\theta,\delta,\phi)=\begin{pmatrix} e^{i\delta}\cos(\theta/2) & -ie^{-i\phi}\sin(\theta/2)\\ -ie^{i\phi}\sin(\theta/2) & e^{-i\delta}\cos(\theta/2)\end{pmatrix}, \quad (2)$$

and examine the effect of dephasing or twirling twirling (see footnote 3 in Appendix A) on the state produced by applying this gate to $|0\rangle$. Define

$$\rho(\theta,\delta,\phi)=U(\theta,\delta,\phi)\,|0\rangle\,\langle 0|\,U^{\dagger}(\theta,\delta,\phi). \quad (3)$$

The dephased (twirled) state is

$$\begin{aligned} D(\rho(\theta,\delta,\phi)) &= \langle T_0|\,\rho\,|T_0\rangle\,|T_0\rangle\,\langle T_0| + \langle T_1|\,\rho\,|T_1\rangle\,|T_1\rangle\,\langle T_1| \\ &= \tfrac{1}{6}\big[3+\sqrt{3}\cos\theta-\sqrt{3}\sin\theta(\cos(\delta-\phi)+\sin(\delta-\phi))\big]\,|T_0\rangle\,\langle T_0| \\ &\quad + \tfrac{1}{6}\big[3-\sqrt{3}\cos\theta+\sqrt{3}\sin\theta(\cos(\delta-\phi)+\sin(\delta-\phi))\big]\,|T_1\rangle\,\langle T_1|\,. \end{aligned} \quad (4)$$

In the special case of zero detuning ($\delta = 0$) and zero phase ($\phi = 0$), the error rate (the population weight on $|T_1\rangle$) is minimized by optimizing

$$\frac{3+\sqrt{3}(\sin\theta-\cos\theta)}{6}, \quad (5)$$

which achieves a minimum value of approximately 9.175×10^{-2}. According to Figure 1(c) (and Eq. (A11) in Appendix A), this error level requires six levels of concatenated distillation using the five-qubit code to reach a fault-tolerant threshold below 10^{-15}, underscoring the steep resource overhead.

This analysis assumes an idealized implementation. By contrast, if one allows a general phase choice, e.g. $\phi = -\pi/4$, the desired rotation can be exactly realized in the absence of imperfections. Thus, wherever physically feasible, it is advantageous to exploit the full generality of the gate, even at some implementation cost, as this can significantly reduce the number of distillation rounds by improving the fidelity of the magic state before purification.

With this foundation in place, we next examine how different quantum computing platforms implement arbitrary single-qubit unitaries, how their control parameters can be tuned to realize specific gates such as $\mathcal{T}$ and $\mathcal{H}$, and how composite pulse techniques can be employed to mitigate systematic errors.

IV. Composite Pulses Scheme for X-Y Systems with Global Rabi Frequency Errors

This section presents composite pulse schemes for implementing the $\mathcal{T}$ gate in systems governed by X-Y control Hamiltonians, which are common in superconducting qubits, trapped ions, and neutral atoms. For more details, see Appendix B. We focus on mitigating global systematic errors in the Rabi frequency- errors that uniformly scale the rotation angle across all segments. By constructing symmetric composite sequences, we analytically derive gate implementations that are robust to first- and second-order errors while exactly realizing the target unitary.

A. Notation

Our objective is to determine the set of parameters $\{\theta_i,\phi_i\}_{i=1}^{N}$ that define a composite pulse sequence imple-

menting a target single-qubit gate $G = U(\theta^*, \phi^*) \equiv \theta^*_{\phi^*}$, where a general rotation of θ_ϕ is given by

$$\theta_\phi \equiv U(\theta, \phi) = \begin{pmatrix} \cos(\theta/2) & -ie^{-i\phi}\sin(\theta/2) \\ -ie^{i\phi}\sin(\theta/2) & \cos(\theta/2) \end{pmatrix}. \quad (6)$$

Here, θ denotes the rotation angle and ϕ the phase of the driving field, corresponding to a rotation axis in the equatorial plane of the Bloch sphere. The overall composite gate constructed from N such segments is given by:

$$U_N \equiv U(\{\theta_i, \phi_i\}_{i=1}^N) = \prod_{i=1}^{N} U(\theta_i, \phi_i). \quad (7)$$

To improve robustness against systematic errors, we optimize the parameters such that the derivatives of the composite gate $U(\{\theta_i, \phi_i\}; \theta^*, \phi^*)$ with respect to a global error parameter vanish. The performance of the composite gate can be evaluated for example using the trace fidelity,

$$\mathcal{F}_{\text{trace}} = \frac{1}{2}\left|\text{Tr}\left[U(\{\theta_i, \phi_i\}_i; \theta^*, \phi^*)G^\dagger(\theta^*, \phi^*)\right]\right|, \quad (8)$$

which serves as a figure of merit, although it does not fully capture the detailed nature of gate deviations.

B. The Composite Pulses Scheme

To suppress control errors in such systems, we employ a symmetric sequence of composite pulses (or segmented gates), motivated by numerical evidence supporting its robustness. The unitary transformation is given by:

$$U_{2n-1} = \theta_{\phi_1}\pi_{\phi_2}\pi_{\phi_3}\cdots\pi_{\phi_{n-1}}\pi_{\phi_n}\pi_{\phi_{n-1}}\cdots\pi_{\phi_3}\pi_{\phi_2}\theta_{\phi_1}. \quad (9)$$

The use of this symmetric construction enhances robustness against systematic errors in the Rabi frequency or in the overall scaling of the Hamiltonian.

We denote the k^{th} derivative of the unitary rotation $U(\theta, \phi)$ with respect to θ by $U^{(k,0)}(\theta, \phi)$. These derivatives take the following forms:

$$U^{(2k,0)}(\theta, \phi) = (-1)^k \frac{1}{2^k}\begin{pmatrix} \cos(\theta/2) & -ie^{-i\phi}\sin(\theta/2) \\ -ie^{i\phi}\sin(\theta/2) & \cos(\theta/2) \end{pmatrix}, \quad (10a)$$

$$U^{(2k+1,0)}(\theta, \phi) = (-1)^k \frac{1}{2^{k+1}}\begin{pmatrix} -\sin(\theta/2) & -ie^{-i\phi}\cos(\theta/2) \\ -ie^{i\phi}\cos(\theta/2) & -\sin(\theta/2) \end{pmatrix}. \quad (10b)$$

Evaluating these expressions at $\theta = \pi$, we find:

$$\begin{aligned} U^{(2k,0)}(\pi, \phi) &= (-1)^k \frac{1}{2^k}\begin{pmatrix} 0 & -ie^{-i\phi} \\ -ie^{i\phi} & 0 \end{pmatrix} \\ &= -(-1)^{k+1}\frac{i}{2^{k+1}}\left(\cos\phi\, X + \sin\phi\, Y\right), \end{aligned} \quad (11a)$$

$$\begin{aligned} U^{(2k+1,0)}(\pi, \phi) &= (-1)^{k+1}\frac{1}{2^{k+1}}\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \\ &= (-1)^{k+1}\frac{1}{2^{k+1}} I. \end{aligned} \quad (11b)$$

Here, X, Y, and I denote the standard Pauli matrices and the identity operator, respectively. These expressions reveal the periodic structure and symmetry in the higher-order derivatives, which will be useful in analyzing the sensitivity of composite gates to systematic errors.

The error model of the systems we are talking about is $\theta_i \to \theta_i(1+\epsilon)$ for each pulse or segment. Therefore, the k^{th} variation of $U(\theta, \phi)$, denoted by $\delta^k U(\theta, \phi)$ satisfies

$$k!\,\delta^k U(\theta, \phi) = \theta^k U^{(k,0)}(\theta, \phi). \quad (12)$$

C. Three-Pulse Sequence ($n = 2$)

For the minimal case $n = 2$, the symmetric composite sequence consists of three segments. To achieve robustness to first-order global errors while implementing a desired gate $U(\theta^*, \phi^*)$, we impose the following conditions:

$$U^{(0,0)}(\theta, \phi) = U(\theta^*, \phi^*), \quad (13a)$$

$$U^{(1,0)}(\theta, \phi) = 0. \quad (13b)$$

These constraints yield the system:

$$-\cos(\phi_1 - \phi_2)\sin(\theta) = \cos\left(\frac{\theta^*}{2}\right), \quad (14a)$$

$$e^{i\phi^*}\sin\left(\frac{\theta^*}{2}\right) = e^{i\phi_1}\left(e^{i(\phi_2-\phi_1)}\cos^2\left(\frac{\theta}{2}\right) - e^{-i(\phi_2-\phi_1)}\sin^2\left(\frac{\theta}{2}\right)\right), \quad (14b)$$

$$(2\theta\cos(\phi_1 - \phi_2) + \pi)\cos(\theta) = 0, \quad (14c)$$

$$(2\theta\cos(\phi_1 - \phi_2) + \pi)\sin(\theta) = 0. \quad (14d)$$

This system admits a unique solution given by:

$$\frac{\sin\theta}{\theta} = \frac{2}{\pi}\cos\left(\frac{\theta^*}{2}\right), \quad (15a)$$

$$\cos(\phi_2 - \phi_1) = -\frac{\pi}{2\theta}, \quad (15b)$$

$$e^{i\phi_{1\pm}} = \frac{\sin\left(\frac{\theta^*}{2}\right)e^{i\phi^*}}{\cos\theta\cos x \pm i\sin x}, \quad (15c)$$

where $x = \phi_2 - \phi_1$. This solution ensures first-order insensitivity to global Rabi frequency errors while exactly implementing the target gate.

In Table I, we present representative solutions for selected values of the target angle θ^*. Note that for $\theta^* = 1$ and $\phi^* = 0$, the resulting gate is $-iX$, and for $\theta^* = \frac{1}{2}$ and $\phi^* = \frac{\pi}{4}$, the gate corresponds to ZH

θ^*	$(\theta\ ;\ \phi_1-\phi^*\ ,\ \phi_2-\phi^*)$
$\frac{1}{10}$	(0.50614 ; −0.46114 , 0.48923)
$\frac{1}{5}$	(0.52421 ; −0.42464 , 0.47824)
$\frac{3}{10}$	(0.55331 ; −0.39234 , 0.46678)
$\frac{2}{5}$	(0.59226 ; −0.36536 , 0.45458)
$\frac{1}{2}$	(0.63990 ; −0.34419 , 0.44129)
$\frac{3}{5}$	(0.69538 ; −0.32894 , 0.42648)
$\frac{7}{10}$	(0.75827 ; −0.31966 , 0.40952)
$\frac{4}{5}$	(0.82882 ; −0.31661 , 0.38952)
$\frac{9}{10}$	(0.90828 ; −0.32055 , 0.36501)
1	(1.00000 ; −0.33333 , 0.33333)

TABLE I. Composite pulse parameters for three-segment sequences: each row lists the values of $(\theta;\phi_1-\phi^*,\phi_2-\phi^*)$ corresponding to a target angle θ^*, all expressed in units of π. The values were obtained by Eqs. (15).

D. Five-Pulse Sequence ($n = 3$)

For the case $n = 3$, the symmetric composite sequence consists of five segments. To achieve robustness against both first- and second-order global Rabi frequency errors while implementing a target gate $U(\theta^*,\phi^*)$, we impose the following derivative cancellation conditions:

$$U^{(0,0)}(\theta,\phi) = U(\theta^*,\phi^*), \tag{16a}$$
$$U^{(1,0)}(\theta,\phi) = 0, \tag{16b}$$
$$U^{(2,0)}(\theta,\phi) = 0. \tag{16c}$$

We define the phase parameters of the sequence as $\phi_3 = \phi_2+\alpha$ and $\phi_1 = \phi_2+\beta$. This results in the following system of constraints:

$$\cos(\alpha+\beta)\sin\theta = \cos\left(\frac{\theta^*}{2}\right), \tag{17a}$$
$$-e^{i\phi_1}\left(\cos\theta\cos(\alpha+\beta)+i\sin(\alpha+\beta)\right) = e^{i\phi^*}\sin\left(\frac{\theta^*}{2}\right), \tag{17b}$$
$$\pi + 2\pi\cos\alpha + 2\theta\cos(\alpha+\beta) = 0, \tag{17c}$$
$$4\pi\theta + 4\pi^2\cos\beta + 2\pi^2\cos(\beta-\alpha) + 8\pi\theta\cos\alpha + (3\pi^2+4\theta^2)\cos(\alpha+\beta) = 0, \tag{17d}$$
$$4\sin\beta + 2\sin(\beta-\alpha) + 3\sin(\alpha+\beta) = 0. \tag{17e}$$

This set comprises five nonlinear equations in four variables: α, β, θ, and ϕ_2. While the system is not solvable analytically for arbitrary target parameters θ^* and ϕ^*, particular solutions can be constructed. For example, the choice $\alpha = -\arccos(-4/5)$, $\beta = 3\pi/2$, and $\theta = 3\pi/2$ satisfies the equations and yields the implemented gate angle $\theta^* = \pm 2\arccos(3/5)$ with a corresponding relation for ϕ_2 for every ϕ^*.

A general parametric solution can be obtained by solving the last three equations first. In particular, Eqs. (17c-e) can be solved in closed form as:

$$\beta = \frac{\pi}{2} \pm \frac{\pi}{2} - \arctan\left(\frac{\sin\alpha}{4+5\cos\alpha}\right), \tag{18a}$$
$$\theta = -\frac{\pi}{2}\cdot\frac{1+2\cos\alpha}{\cos(\alpha+\beta)}. \tag{18b}$$

Substituting these into the first two conditions, we obtain:

$$\cos(\alpha+\beta)\sin\theta = \cos\left(\frac{\theta^*}{2}\right), \tag{19a}$$
$$\phi_2 = \phi^* - \beta + \arctan\left(\frac{\tan(\alpha+\beta)}{\cos\theta}\right) + \frac{\pi}{2} \pm \frac{\pi}{2}. \tag{19b}$$

Using Eqs. (18b) and (19a), we find that $\cos(\alpha+\beta)$ must satisfy the condition:

$$\cos(\alpha+\beta) = -\frac{\cos\left(\frac{\theta^*}{2}\right)}{\sin\left(\frac{\pi}{2}\cdot\frac{1+2\cos\alpha}{\cos(\alpha+\beta)}\right)}. \tag{20}$$

Inserting Eq. (18a) into this, we derive a self-consistency condition for α:

$$0 = \cos\left(\frac{\theta^*}{2}\right) + \cos\left(\alpha - \arctan\left(\frac{\sin\alpha}{4+5\cos\alpha}\right)\right) \cdot \sin\left(\frac{\pi}{2}\cdot\frac{1+2\cos\alpha}{\cos\left(\alpha-\arctan\left(\frac{\sin\alpha}{4+5\cos\alpha}\right)\right)}\right) \tag{21}$$

This is a transcendental equation in α that must be solved numerically. Once a solution for α is found, the remaining parameters β, θ, and ϕ_2 are determined via Eqs. (18) and (19). In Table II, we present representative solutions for selected values of the target angle θ^*. Note that for $\theta^* = 1$ and $\phi^* = 0$, the resulting gate is $-iX$, and for $\theta^* = \frac{1}{2}$ and $\phi^* = \frac{\pi}{4}$, the gate corresponds to ZH.

θ^*	$(\theta\ ;\ \phi_1-\phi^*\ ,\ \phi_2-\phi^*\ ,\ \phi_3-\phi^*)$
$\frac{1}{10}$	(1.50291 ; 0.48163 , −0.51210 , −0.45592)
$\frac{1}{5}$	(1.51182 ; 0.46354 , −0.52386 , −0.41193)
$\frac{3}{10}$	(1.52730 ; 0.44603 , −0.53487 , −0.36810)
$\frac{2}{5}$	(1.55034 ; 0.42955 , −0.54464 , −0.32444)
$\frac{1}{2}$	(1.58238 ; 0.41472 , −0.55249 , −0.28086)
$\frac{3}{5}$	(1.62536 ; 0.40238 , −0.55752 , −0.23704)
$\frac{7}{10}$	(1.68194 ; 0.39372 , −0.55852 , −0.19226)
$\frac{4}{5}$	(1.75611 ; 0.39037 , −0.55380 , −0.14500)
$\frac{9}{10}$	(1.85530 ; 0.39495 , −0.54056 , −0.09202)
1	(2.00000 ; 0.41312 , −0.51245 , −0.02491)

TABLE II. Composite pulse parameters for five-segment sequences: each row lists the values of $(\theta;\phi_1-\phi^*,\phi_2-\phi^*,\phi_3-\phi^*)$ corresponding to a target angle θ^*, all expressed in units of π, shown to five decimal digits. The values were obtained by choosing the positive sign in Eqs. (18a) and the negative sign in Eq. (19b), ensuring physically valid solutions.

E. Pre-Distillation: The Improvement in the $\mathcal{T}$ Gate

In quantum computing research, several fidelity and distance measures are used to evaluate how closely an implemented gate approximates its ideal version, each suited for different purposes. The average gate fidelity is the most commonly reported measure, especially in experiments, as it captures how well a gate performs on average across all possible input states and can be efficiently estimated using randomized benchmarking. For more detailed characterization, process fidelity compares the full quantum processes and is often used in quantum process tomography, although it's more resource-intensive to compute. Trace fidelity and trace distance are state-based measures used to assess how distinguishable the output states are, useful in situations like error diagnostics or state preparation. For fault-tolerant quantum computing, the diamond norm distance is the gold standard- it measures the worst-case deviation of the actual gate from the ideal one, accounting for all possible inputs and entanglements with ancillas. While very informative, it's also the most computationally demanding. In contrast, the Frobenius distance is a simple matrix norm used occasionally in simulations for quick comparisons but lacks the physical rigor of other measures. Each metric balances between ease of computation and how well it captures meaningful quantum behavior, with average fidelity favored in benchmarking, diamond norm in theoretical error thresholds, and trace-based metrics in state analysis.

To assess the improvement in implementing the $\mathcal{T}$ gate, we do not rely on standard fidelity or distance measures. Instead, we adopt a natural and operationally motivated metric, denoted by $\varepsilon(G(\epsilon))$, which quantifies the deviation of the output state from the ideal $|T_0\rangle$ when a noisy gate $G(\epsilon)$, affected by a global control error ϵ, is applied to the input state $|0\rangle$. Specifically, we define

$$\varepsilon(G(\epsilon)) = \langle T_1| \, \rho(G(\epsilon)) \, |T_1\rangle \,, \tag{22}$$

where $\rho(G(\epsilon))$ is the pure state resulting from the gate action:

$$\rho(G(\epsilon)) = \left(G_{1,1}(\epsilon)\,|0\rangle + G_{2,1}(\epsilon)\,|1\rangle\right)\left(G^*_{1,1}(\epsilon)\,\langle 0| + G^*_{2,1}(\epsilon)\,\langle 1|\right). \tag{23}$$

We Call $\varepsilon(G(\epsilon))$ *the T-magic error* of the gate. This measure is particularly relevant for applications in magic state injection into stabilizer codes, where the output of $G(\epsilon)\,|0\rangle$ is intended to serve as a noisy magic state. The metric $\varepsilon(G(\epsilon))$ reflects the overlap with the orthogonal state $|T_1\rangle$, thereby capturing the component of the error that undermines distillability via, for example, the five-qubit code. Our choice of this metric is guided by the operational considerations discussed in Section III A.

We introduce a natural and operationally motivated definition of the *magic T-gate fidelity*:

$$\mathcal{F}_{\text{magic } T\text{-gate}} = 1 - \sqrt{\varepsilon(G(\epsilon))}. \tag{24}$$

This measure quantifies the success of preparing the ideal $|T\rangle$ state using the noisy gate $G(\epsilon)$, thereby providing a direct estimate of the number of distillation levels or concatenation rounds required for fault-tolerant encoding. Indeed, this is the central figure of merit we will plot throughout our analysis, whereas the T-magic fidelity itself is plotted only in Appendix D, to illustrate its nontrivial behavior and its lack of monotonicity with respect to other standard gate fidelity metrics. It is important to note that $\mathcal{F}_{\text{magic } T\text{-gate}}$ does not capture the full fidelity of the gate $G(\epsilon)$ when applied to arbitrary input states or within general qudit or multi-qubit contexts. Rather, its definition is tailored to assess the quality of the output state obtained when the noisy gate acts on a standard stabilizer input, specifically in the context of MSD. For completeness, we also recall two commonly used fidelity fidelity measures:

$$\mathcal{F}_{\text{frob}}(G(\epsilon),\mathcal{T}) = 1 - \sqrt{\frac{1}{4}\text{Tr}\left[\left(G(\epsilon)-\mathcal{T}\right)\left(G^\dagger(\epsilon)-\mathcal{T}^\dagger\right)\right]}\,, \tag{25a}$$

$$\mathcal{F}_{\text{trace}}(G(\epsilon),\mathcal{T}) = 1 - \frac{1}{2}\text{Tr}\left[G(\epsilon)\mathcal{T}^\dagger\right]\,. \tag{25b}$$

The parameters defining the target $\mathcal{T}$ gate are given by (see Eqs. (1) and (6)):

$$\theta^* = \arccos\left(\frac{1}{\sqrt{3}}\right) \quad , \quad \phi^* = \frac{3\pi}{4}\,. \tag{26}$$

The parameters θ, ϕ_1, ϕ_2 for the three-pulse composite implementation of the $\mathcal{T}$ gate are (see Eqs. (15)):

$$\theta \approx 1.74270 \quad , \quad \phi_1 \approx 1.12743 \quad , \quad \phi_2 \approx 3.82112\,. \tag{27}$$

For the five-pulse composite implementation, the corresponding parameters $\theta, \phi_1, \phi_2, \phi_3$ are:

$$\theta \approx 4.80062\,,\ \phi_1 \approx 3.75525\,,\ \phi_2 \approx 0.67449\,,\ \phi_3 \approx 1.20539\,. \tag{28}$$

A numerically obtained solution for a seven-pulse composite sequence that achieves third-order robustness to control errors in the gate implementation is given by:

$$\begin{aligned}\theta \approx 1.78928\,,\ \phi_1 \approx 3.4837\,,\ \phi_2 \approx 4.23899\,,\\ \phi_3 \approx 1.15951\,,\ \phi_4 \approx 0.894556\,.\end{aligned} \tag{29}$$

Figure 2 illustrates the effectiveness of composite pulse sequences in improving the fidelity of the $\mathcal{T}$ gate and in reducing the resource overhead required for MSD. The black curves correspond to the single-pulse implementation, while the blue, red, and brown curves correspond to the three-, five-, and seven-pulse composite gates. Panel (a) compares two fidelity measures, Frobenius distance fidelity (solid lines) and trace distance fidelity (dashed lines), for single-, three-, five-, and seven-pulse implementations under global systematic errors. The multi-pulse sequences clearly enhance gate robustness, reducing sensitivity to over- and under-rotations.

Panel (b) highlights the practical significance of this improvement by showing the number of distillation iterations required to reach a target gate error of 10^{-15}. In several physical error regimes, the use of composite sequences reduces the number of required distillation levels by one, two, or even three, substantially lowering the overhead for fault-tolerant quantum computation.

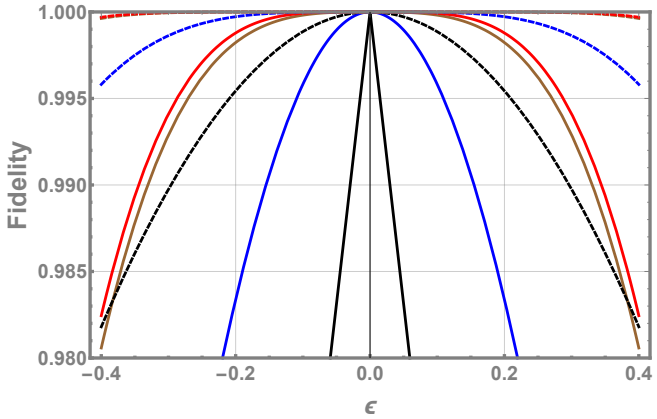

(a) Frobenius fidelity $\mathcal{F}_{\text{frob}}$ for different composite pulse schemes.

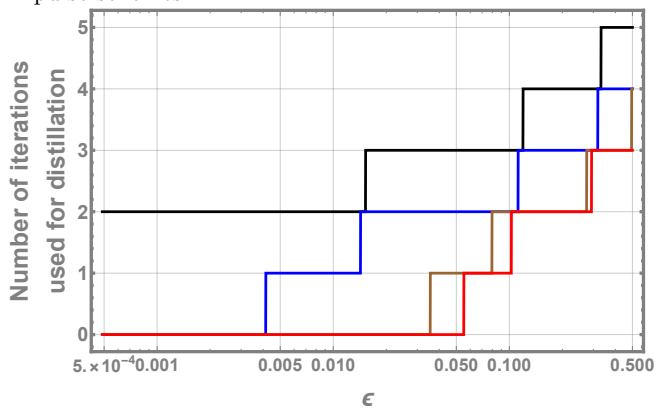

(b) Number of distillation levels required to reach an error of 10^{-15}.

FIG. 2. Improvements enabled by composite pulse sequences for implementing the $\mathcal{T}$ gate in X-Y Systems with Global Rabi Frequency Errors. The black curves correspond to the single-pulse implementation, while the blue, red, and brown curves correspond to the three-, five-, and seven-pulse composite gates. (a) Comparison of Frobenius distance fidelity $\mathcal{F}_{\text{frob}}$ (solid lines) and trace distance fidelity $\mathcal{F}_{\text{trace}}$ (dashed lines) for single-, three-, five-, and seven-pulse implementations under global systematic errors. (b) Reduction in the number of MSD iterations required to achieve a target gate error of 10^{-15}, where composite sequences reduce the required levels by one, two, or even three in some error regimes.

F. Pre-distillation of H-states

We now apply the pre-distillation scheme presented in Tables I, II to the case of H-states. Although T-states are considered to possess more magic than H-states, exhibiting higher symmetry and lower stabilizer fidelity, the threshold fidelity required for any protocol that distills H-states is larger than that of T-states [6].

Moreover, in terms of error-correcting codes, the simplest code capable of protecting a logical T-state qubit against arbitrary single-qubit errors encodes one logical qubit into five physical qubits, whereas for H-states the smallest such code requires a 15-to-1 encoding. Nevertheless, the error-suppression performance of H-state distillation is significantly better: the output error scales as

$$\epsilon' = 35\epsilon^3 + O(\epsilon^4),$$

compared to

$$\epsilon' = 5\epsilon^2 + O(\epsilon^3)$$

for T-states [6]. This cubic suppression makes H-states a more efficient resource in certain regimes, highlighting the practical motivation for pursuing H-state pre-distillation.

The parameters defining the target $\mathcal{H}$ gate are given by (see Eqs. (1) and (6)):

$$\theta^* = \frac{\pi}{4} \quad , \quad \phi^* = \frac{\pi}{2}\,. \tag{30}$$

The parameters θ, ϕ_1, ϕ_2 for the three-pulse composite implementation of the $\mathcal{T}$ gate are (see Eqs. (15)):

$$\theta \approx 1.68846 \quad , \quad \phi_1 \approx 0.28941 \quad , \quad \phi_2 \approx 3.05547\,. \tag{31}$$

For the five-pulse composite implementation, the corresponding parameters $\theta, \phi_1, \phi_2, \phi_3$ are:

$$\begin{aligned} \theta \approx 4.77109\ ,\ \phi_1 \approx 2.99923\ , \\ \phi_2 \approx -0.09264\ ,\ \phi_3 \approx 0.34559\,. \end{aligned} \tag{32}$$

A numerically obtained solution for a seven-pulse composite sequence that achieves third-order robustness to control errors in the gate implementation is given by:

$$\begin{aligned} \theta \approx 1.72181\ ,\ \phi_1 \approx 2.76539\ ,\ \phi_2 \approx 3.39854\ , \\ \phi_3 \approx 0.30736\,,\ \phi_4 \approx 0.08854\,. \end{aligned} \tag{33}$$

The error suppression achieved by the 15-1 code is given in [6] as:

$$\epsilon'(\epsilon) = \frac{1 - 15(1-2\epsilon)^7 + 15(1-2\epsilon)^8 - (1-2\epsilon)^{15}}{2\big(1 + 15(1-2\epsilon)^8\big)}. \tag{34}$$

Figure 3 illustrates the effectiveness of composite pulse sequences in improving the fidelity of the $\mathcal{H}$ gate and in reducing the resource overhead required for MSD.

V. Example: Pre-Distillation in X-Z Systems with Global Z Errors

As discussed in Section III B (and in Appendix B), errors in integrated photonic platforms typically arise in

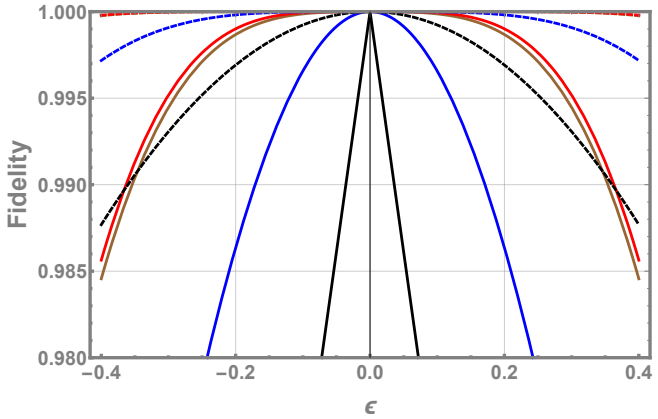

(a) Frobenius fidelity $\mathcal{F}_{\text{frob}}$ for different composite pulse schemes.

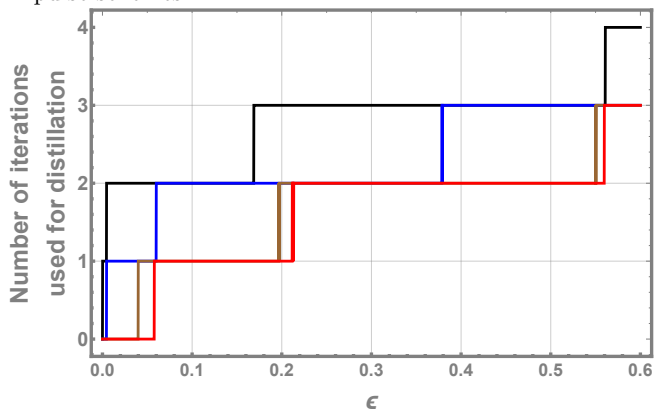

(b) Number of distillation levels required to reach an error of 10^{-15}.

FIG. 3. Improvements enabled by composite pulse sequences for implementing the $\mathcal{H}$ gate in X-Y Systems with global Rabi frequency errors. The black curves correspond to the single-pulse implementation, while the blue, red, and brown curves correspond to the three-, five-, and seven-pulse composite gates. (a) Comparison of Frobenius distance fidelity $\mathcal{F}_{\text{frob}}$ (solid lines) and trace distance fidelity $\mathcal{F}_{\text{trace}}$ (dashed lines) for single-, three-, five-, and seven-pulse implementations under global systematic errors. (b) Reduction in the number of MSD iterations required to achieve a target gate error of 10^{-15}. In certain error regimes, composite sequences lower the required distillation levels by one or two.

both the Rabi frequency and the detuning parameters. For more details, see Appendix B. To demonstrate the effectiveness of composite pulse schemes in constructing the $\mathcal{T}$ gate, we consider a simplified error model previously studied in Ref. [16]. In this model, the system is restricted to real-valued Rabi frequencies Ω, and we focus on a single error parameter: a systematic Z error denoted by $\epsilon_k = \epsilon$ for all segments $k = 1, \ldots, N$. In this model, we neglect errors in the Rabi frequency Ω and assume that the detuning errors are fully correlated across all segments.

This model captures the essential features of error behavior in a variety of physical platforms governed by SU(2) dynamics, including trapped-ion qubits [20, 21], nonlinear optical processes such as sum-frequency generation [29], and atomic systems [30–32].

The total unitary evolution under an N-segment control sequence is given by

$$U_N = \prod_{k=1}^{N} U\left(\theta_k, \phi_k, \epsilon\right), \tag{35}$$

where each segment corresponds to the unitary

$$U\left(\theta, \phi, \epsilon\right) = \exp\left[i\theta\left(\cos\phi X + \sin\phi Z\right) + i\epsilon Z\right], \tag{36}$$

and where θ_k and ϕ_k denote the parameters of the k^{th} segment, respectively.

For the purpose of facilitating numerical optimization in the next section, which also deals with an $X - Z$ system, we redefine the $\mathcal{T}$ gate and the associated T-states as follows.

$$|T_0\rangle = \cos\beta\ |0\rangle - e^{-i\pi/4}\sin\beta\ |1\rangle, \tag{37a}$$

$$|T_1\rangle = e^{i\pi/4}\sin\beta\ |0\rangle + \cos\beta\ |1\rangle, \tag{37b}$$

$$\mathcal{T} = \begin{pmatrix} \cos\beta & e^{i\pi/4}\sin\beta \\ -e^{-i\pi/4}\sin\beta & \cos\beta \end{pmatrix}. \tag{37c}$$

These states are Clifford-equivalent to the previously defined T-states, and thus this alternative representation is fully consistent for our purposes.

Since the $\mathcal{T}$ gate involves a non-zero component along the Y axis, it cannot be implemented using a single segment with real-valued control fields. At least two segments are required to synthesize such a unitary. Solving the constraint equation $U_2 = \mathcal{T}$ yields the following relations:

$$\begin{aligned} \cot\theta_1 &= -\sin\phi_1 \quad , \quad \cot\theta_2 = \sin\phi_2, \\ \cot\phi_2 &= \frac{(1+\sqrt{3})\cot\phi_1 \pm 2}{(1+\sqrt{3}) \mp 1\cot\phi_1}, \end{aligned} \tag{38}$$

where ϕ_1 can be chosen freely as a parameter.

To achieve first-order robustness against detuning errors, we design a three-segment composite pulse sequence by solving the coupled conditions $U_3 = \mathcal{T}$ and $\frac{dU_3}{d\epsilon}\big|_{\epsilon=0} = 0$. This leads to a continuous family of solutions that cannot be expressed analytically. From this family, we select the following three representative solutions given in Table III to demonstrate the feasibility such robust constructions.

Sequence	ϕ_1	ϕ_2	ϕ_3	θ_1	θ_2	θ_3
(a)	3.31558	−3.14159	−0.17399	1.26121	−2.86242	−1.26121
(b)	0.49741	−0.06006	−0.41585	3.25204	−1.56495	−1.35397
(c)	0.55004	−0.08606	−3.58545	−3.02982	−1.56520	1.35359

TABLE III. Optimized pulse parameters (θ_k, ϕ_k) for three-segment composite sequences achieving first-order robustness to detuning errors. All values are in radians, shown to five decimal digits.

Figure 4 illustrates the effectiveness of these composite pulse schemes in enhancing the fidelity of the $\mathcal{T}$

gate under correlated systematic Z errors, as arise from global imperfections in coupler widths in X-Z control systems. The black curve corresponds to the baseline two-segment implementation, while the red, blue, and brown curves represent three distinct three-segment composite sequences labeled (a), (b), and (c), respectively.

Panel 4(a) displays the Frobenius fidelity $\mathcal{F}_{\text{frob}}$ as a function of the global error parameter, highlighting the superior robustness of the composite designs. Panel 4(b) quantifies the resulting resource savings by plotting the number of distillation levels required to reach a target gate error of 10^{-15}, based on the T-magic error of the gate- the operational error metric defined in Eq. (22). Notably, in wide regions of the error domain, the composite pulse schemes reduce the required number of distillation rounds by up to two levels- achieving this improvement solely through a more robust gate synthesis using three pulses instead of two.

VI. Pre-Distillation in Integrated Photonics

As discussed in Section III B (and in Appendix B), we model the dominant source of error in integrated photonic implementations as a fully correlated global perturbation in the waveguide widths across all segments. Specifically, each segment experiences a common deviation δw added to both waveguide widths. For more details, see Appendix B.

The total unitary evolution of an N-segment control sequence is given by

$$U_N = \prod_{k=1}^{N} U(z_k, w_{1,k}, w_{2,k}; \delta w), \tag{39}$$

where z_k is the length of the k^{th} segment, $w_{1,k}$ and $w_{2,k}$ are the nominal widths of the two waveguides, and δw is the global width error. The unitary for each segment takes the form

$$U(z_k, w_{1,k}, w_{2,k}; \delta w) = \exp\Big[-i\frac{z_k}{2}\Big(\Omega(w_{1,k}+\delta w, w_{2,k}+\delta w)X - \Delta(w_{1,k}+\delta w, w_{2,k}+\delta w)Z\Big)\Big], \tag{40}$$

where Ω_k and Δ_k denote the effective coupling strength and detuning for the k^{th} segment, determined by the geometry.

The interpolation functions $\Omega(w_1, w_2)$ and $\Delta(w_1, w_2)$ map the physical widths to the effective Hamiltonian parameters. These were obtained via numerical simulations in Ref. [16] for width values between 350 nm and 450 nm. The fitted polynomial expressions (with widths expressed in micrometers and resulting values in units of $1/\mu$m) are:

$$\begin{aligned}\Delta(w_1, w_2) = {} & 3.94502808\, w_2 - 18.0203544\, w_2^2 \\ & + 27.94843595\, w_2^3 - 15.42295066\, w_2^4 \\ & - 3.94502818\, w_1 + 18.02035521\, w_1^2 \\ & - 27.94843797\, w_1^3 + 15.42295233\, w_1^4,\end{aligned} \tag{41a}$$

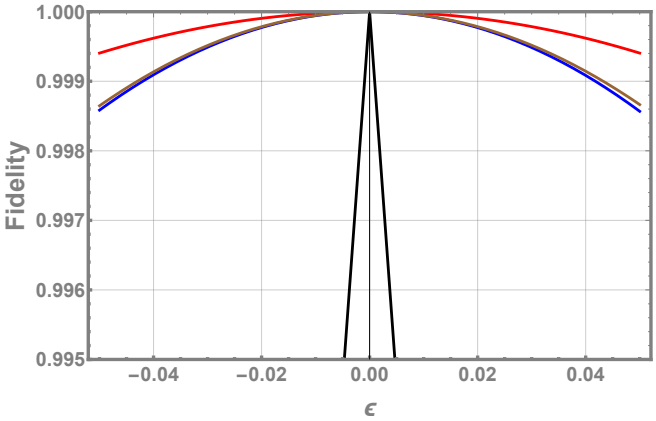

(a) Frobenius fidelity $\mathcal{F}_{\text{frob}}$ for different composite pulse schemes.

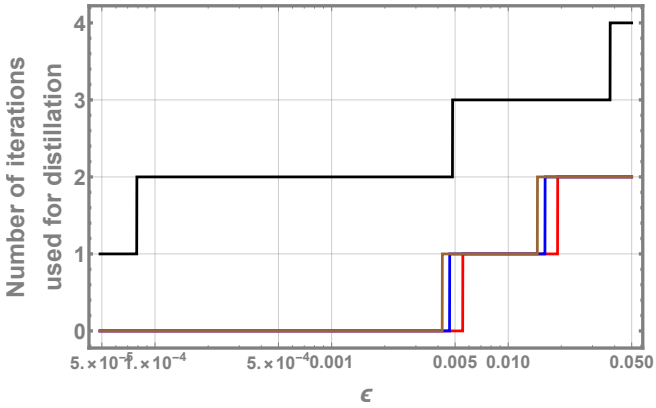

(b) Number of distillation levels required to reach an error of 10^{-15}.

FIG. 4. Performance of $\mathcal{T}$ gate implementations using composite pulse sequences under correlated systematic Z errors in X-Z Hamiltonians. The black curve shows the two-segment baseline, while the red, blue, and brown curves correspond to three different three-pulse designs labeled (a), (b), and (c). Panel (a) presents the Frobenius fidelity $\mathcal{F}_{\text{frob}}$ across these schemes. Panel (b) shows the corresponding reduction in the number of magic state distillation levels required to achieve a gate T-magic error below 10^{-15}, based on Eq. (22).

$$\begin{aligned}\Omega(w_1, w_2) = {} & 0.38044405 - 1.48138422(w_1 + w_2) \\ & + 2.51783632(w_1 + w_2)^2 - 1.9993113(w_1 + w_2)^3 \\ & + 0.60771393(w_1 + w_2)^4.\end{aligned} \tag{41b}$$

To ensure compatibility with fabrication constraints, we restrict all geometric parameters, specifically, the waveguide widths and segment lengths, to integer multiples of 1 nm. For each candidate solution to gate synthesis, we round these values up or down and select the combination that yields the highest fidelity to the target unitary. The fidelity is evaluated by comparing the synthesized gate to the desired one, and the optimal configuration is retained.

Table IV summarizes the implementations of the $\mathcal{T}$ gate using both two- and four-segment pulse sequences. The two-segment design directly implements the target

unitary without robustness considerations. In contrast, the four-segment case yields a continuous family of solutions that achieve first-order robustness against variations in the coupler width, δw. Although this family cannot be succinctly described in closed form, we identify three representative configurations that exhibit high fidelity and practical feasibility.

Design	Segment	w_1 (nm)	w_2 (nm)	z (μm)
Two-segment	1	449	356	23.813
	2	356	449	23.813
Four-segment: Design (a)	1	449	387	23.220
	2	424	448	17.190
	3	450	351	21.152
	4	353	450	36.094
Four-segment: Design (b)	1	448	387	23.822
	2	395	422	15.660
	3	449	350	21.215
	4	355	449	37.081
Four-segment: Design (c)	1	450	389	22.675
	2	395	407	18.369
	3	449	350	20.343
	4	359	450	39.073

TABLE IV. Segment parameters for two-, three-, and four-segment composite implementations of the $\mathcal{T}$ gate. Widths w_1, w_2 are in nm, lengths z in μm. The three four-segment designs are representative samples from a continuous robust family.

Figure 5 illustrates the performance of composite pulse schemes in enhancing the fidelity of the $\mathcal{T}$ gate under global systematic errors in the coupler widths. The black curve corresponds to the two-segment implementation, while the blue, red, and brown curves represent three distinct four-pulse schemes labeled (a), (b), and (c), respectively.

Panel 5(a) compares the Frobenius fidelity $\mathcal{F}_{\text{frob}}$ across these schemes. Due to limitations in the available generators and fabrication constraints, the fidelity achieved by the four-pulse composite schemes can be slightly lower than that of the two-segment design for very small error values. However, across a broader error range, the composite sequences demonstrate superior robustness.

Panel 5(b) shows the corresponding reduction in the number of distillation levels required to achieve a target gate T-magic error of 10^{-15}, as defined in Eq. (22). While the fidelity differences observed at small error rates might suggest an advantage for the two-segment scheme, all implementations ultimately require only two levels of MSD in this regime. Notably, although scheme (a) initially appears advantageous, scheme (c) consistently yields the lowest number of distillation levels across the entire error range, making it the most effective choice overall.

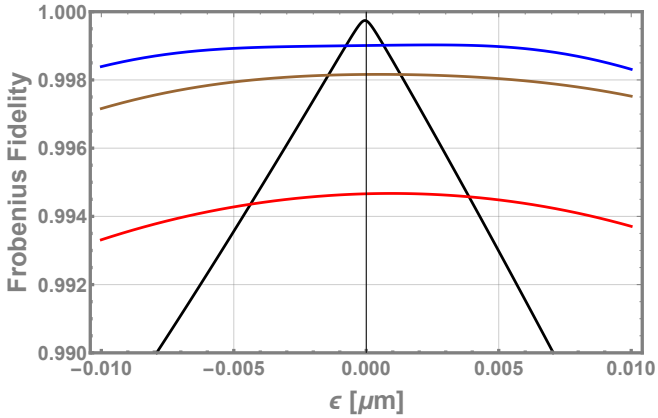

(a) Frobenius fidelity $\mathcal{F}_{\text{frob}}$ for different composite pulse schemes.

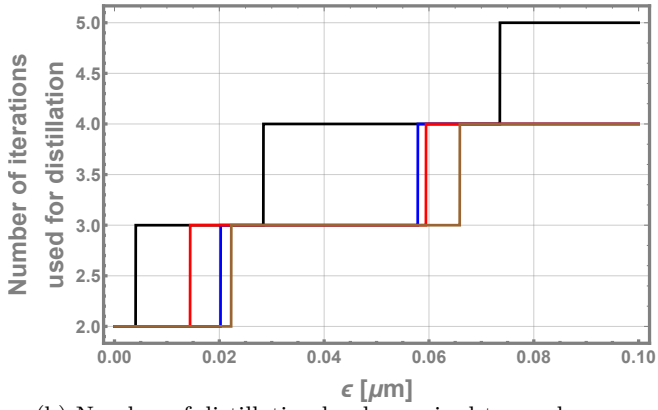

(b) Number of distillation levels required to reach an error of 10^{-15}.

FIG. 5. Performance of $\mathcal{T}$ gate implementations using composite pulse sequences under global systematic width errors in directional couplers. The two-segment design (black) is compared with three four-pulse composite schemes: (a) blue, (b) red, and (c) brown. While scheme (a) shows an initial advantage in fidelity, scheme (c) offers superior performance over the entire error range in terms of the number of concatenation levels.

VII. Errors in Gate-Channel Implementation

As presented in Section II and Appendix A, distilled magic states are injected into Clifford-based, fault-tolerant circuits to realize non-Clifford gates. These gates complete the universal gate set required for fault-tolerant quantum computation. So, in practice, one must also evaluate the error of the implemented gate, since imperfections in the injected magic state directly propagate into the non-Clifford operation and may compromise fault tolerance. Here we show that the error in the operation scales linearly with the leading-order error of the noisy prepared T-states.

A standard technique of the injection employs an ancilla-assisted circuit that applies the desired non-Clifford gate probabilistically via postselection. Consider

the ancilla state

$$|A_\theta\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + e^{i\theta}|1\rangle\right), \tag{42}$$

and assume the initial state is $|\psi\rangle \otimes |A_\theta\rangle$. A stabilizer measurement of $\sigma_z \otimes \sigma_z$ is performed, followed by a CNOT gate with the first qubit as control.

We define the single-qubit rotation

$$\Lambda(e^{i\theta}) \equiv \begin{pmatrix} 1 & 0 \\ 0 & e^{i\theta} \end{pmatrix}. \tag{43}$$

The measurement outcomes yield the following transformations:

- If the outcome is $+1$, the resulting state is $\Lambda(e^{i\theta})\,|\psi\rangle \otimes |0\rangle$.
- If the outcome is -1, the resulting state is $e^{i\theta}\Lambda(e^{-i\theta})\,|\psi\rangle \otimes |1\rangle$.

Hence, the rotation $\Lambda(e^{i\theta})$ is realized with probability $1/2$. At this point, the second qubit can be discarded, i.e., traced out.

In particular, given access to the magic state $|T\rangle$, one can prepare the two-qubit state $|T\rangle \otimes |T\rangle$, perform a $\sigma_z \otimes \sigma_z$ measurement, apply a CNOT, and then a Hadamard gate on the first qubit. For a $+1$ measurement outcome, the postselected state is

$$e^{i\gamma}\,|A_{-2\gamma}, 0\rangle, \quad \text{with } \gamma = \tfrac{\pi}{12},$$

which enables the implementation of the non-Clifford unitary $\Lambda(e^{-2i\gamma})$. Again, the second qubit can be discarded here, i.e., traced out.

Adopting the Bravyi-Kitaev model, we assume the ability to prepare $|0\rangle$ states ideally, to apply Clifford group operations ideally, and to perform ideal measurements of single-qubit Pauli operators. The only faulty resource is the preparation of magic states from stabilizer states, modeled by the mixed state $\rho_T(\epsilon)$. Under these assumptions, one can compute the effective operation corresponding to the intended $\Lambda(e^{-2i\gamma})$ rotation. In fact, the resulting operation is not a unitary, but rather a quantum channel. Explicitly, the channel including the two errors, ϵ_1 and ϵ_2, in the preparataion of the T-states is given by

$$|\psi\rangle\langle\psi| \mapsto \begin{pmatrix} \cos^2\theta & -ia\,e^{-i\phi}\,\cos\theta\sin\theta \\ ia^*\,e^{i\phi}\,\cos\theta\sin\theta & \sin^2\theta \end{pmatrix}, \tag{44}$$

where $|\psi\rangle \equiv \cos\theta\,|0\rangle + e^{i\phi}\sin\theta\,|1\rangle$,

$$a = \frac{1 - i\sqrt{3} + i\left(2i + \sqrt{3}\right)\epsilon_2 + \epsilon_1\left(-2 + i\sqrt{3} + 4\epsilon_2\right)}{-2 + \epsilon_1 + \epsilon_2 - 2\epsilon_1\epsilon_2}, \tag{45}$$

and we used the density matrix representation for the resulting state.

To quantify the performance of this channel, we can use Jozsa's fidelity for quantum states [33] [1] between the achieved state and the desired state. The resulting fidelity is

$$\begin{aligned} F &= 1 - \frac{3}{4}\frac{\epsilon_1 + \epsilon_2}{2 - \epsilon_1 - \epsilon_2 + 2\epsilon_1\epsilon_2}\sin^2(2\theta) \\ &= 1 - \frac{3\sin^2(2\theta)}{8}(\epsilon_1 + \epsilon_2) + O(\epsilon^2). \end{aligned} \tag{46}$$

Hence, the error in the channel scales linearly with the leading-order error of the noisy prepared T-states.

A more suitable fidelity measure is the one proposed in [34], which is based on Jozsa's fidelity for quantum states. Nevertheless, the calculation presented here is representative and sufficient to demonstrate that the gate error depends linearly on the error in the T-states.

More generally, when losses and non-Hermitian dynamics are neglected, any n-qubit gate G implemented in a physical realization with small imperfections can be modeled as $Ge^{i\epsilon_G H_G}$, where $\epsilon_G \ll 1$ quantifies the error strength and H_G is a traceless Hermitian operator, conventionally normalized such that $\mathrm{Tr}\left(H_G^2\right) = 1/2$, specifying the error direction. Both ϵ_G and H_G are determined by the particular physical realization and its underlying error mechanisms. Although one can consider this more general case, the calculation we presented already captures the essential point: the gate error is linear in the error of the T-states.

VIII. Conclusion

We have presented a general and platform-aware method for reducing the overhead of magic state distillation (MSD) through the use of composite pulse sequences that robustly implement the non-Clifford $\mathcal{T}$ gate. In contrast to conventional composite designs focused on Clifford gates such as X or Hadamard, our approach directly targets the synthesis of the $\mathcal{T}$ gate-improving the fidelity of the resulting magic states prior to any distillation protocol.

We derived both analytic and numerical constructions for symmetric three- and five-segment sequences that suppress first- and second-order global Rabi frequency errors in X-Y control systems. Similar techniques were extended to X-Z Hamiltonians, where we demonstrated robustness to global detuning-like imperfections. In integrated photonics, we adapted our framework to directional couplers, designing geometry-optimized segments

[1] The Jozsa's fidelity between two quantum states ρ and σ, expressed as density matrices, is:

$$F(\rho,\sigma) = \left(\mathrm{Tr}\sqrt{\sqrt{\rho}\,\sigma\,\sqrt{\rho}}\right)^2.$$

that yield high-fidelity $\mathcal{T}$ gates resilient to global correlated fabrication imperfections.

To quantify the improvement, we introduced a natural, operationally meaningful fidelity measure tailored to the $\mathcal{T}$ gate, providing a direct link between physical gate performance and distillation cost. Across all platforms, our schemes reduce the required number of distillation iterations by up to three levels, resulting in exponential savings in qubit overhead. This pre-distillation approach complements existing fault-tolerant architectures, and offers a scalable route to higher-fidelity magic states at reduced cost, and paves the way toward more scalable and resource-efficient universal quantum computation.

A natural direction for extending our work is the integration of pre-distillation into fault-tolerant architectures. Benchmarking our composite Clifford- and non-Clifford- gate designs within surface code frameworks would enable a direct assessment of reductions in logical qubit counts and space-time overhead for large-scale algorithms, while exploring their adaptation to LDPC-type quantum codes could highlight advantages for emerging high-threshold, low-overhead schemes. Beyond single-qubit resources, composite sequences may be generalized to prepare two-qubit non-Clifford resources, such as entangled magic states required for Toffoli or controlled-S gates, which are typically much more costly to distill. This line of development naturally leads to the concept of "pre-distilled resource factories" implemented at the hardware level, where composite gates act as the entry point for generating batches of high-fidelity non-Clifford states that can be directly consumed by error-corrected computation, thereby substantially reducing overhead at scale.

Acknowledgements

This work has been supported by the Israel Science Foundation (ISF), the Directorate for Defense Research and Development (DDR&D) under Grant No. 3427/21, and the Tel Aviv University Center for Quantum Science and Technology. M.G. acknowledges additional support from ISF Grant No. 1113/23 and from the US–Israel Binational Science Foundation (BSF) under Grant No. 2020072. Y.O. acknowledges further support from the ISF Excellence Center, the BSF, and the Israel Ministry of Science.

[1] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.

[2] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 212–219, New York, NY, USA, 1996. Association for Computing Machinery.

[3] Bryan Eastin and Emanuel Knill. Restrictions on transversal encoded quantum gate sets. *Phys. Rev. Lett.*, 102:110502, Mar 2009.

[4] Daniel Gottesman. *Stabilizer Codes and Quantum Error Correction*. PhD thesis, California Institute of Technology, 1997.

[5] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[6] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Phys. Rev. A*, 71:022316, Feb 2005.

[7] Sergey Bravyi and Jeongwan Haah. Magic-state distillation with low overhead. *Phys. Rev. A*, 86:052329, Nov 2012.

[8] Earl T. Campbell and Dan E. Browne. On the structure of protocols for magic state distillation. In Andrew Childs and Michele Mosca, editors, *Theory of Quantum Computation, Communication, and Cryptography*, pages 20–32, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.

[9] Malcolm H. Levitt and Ray Freeman. Nmr population inversion using a composite pulse. *Journal of Magnetic Resonance (1969)*, 33(2):473–476, 1979.

[10] A. J. Shaka. Composite pulses for ultra-broadband spin inversion. *Chemical Physics Letters*, 120(2):201–205, 1985.

[11] A.J Shaka and Alexander Pines. Symmetric phase-alternating composite pulses. *Journal of Magnetic Resonance (1969)*, 71(3):495–503, 1987.

[12] Malcolm H. Levitt. Composite pulses. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 18(2):61–122, 1986.

[13] Nuala Timoney, V. Elman, Steffen J. Glaser, C. Weiss, M. Johanning, W. Neuhauser, and Chr. Wunderlich. Error-resistant single-qubit gates with trapped ions. *Physical Review A*, 77(5), 2008.

[14] Elica Kyoseva, Hadar Greener, and Haim Suchowski. Detuning-modulated composite pulses for high-fidelity robust quantum control. *Physical Review A*, 100(3):032333, 2019.

[15] Moshe Katzman, Yonatan Piasetzky, Evyatar Rubin, Ben Barenboim, Maayan Priel, Muhammad Erew, Avi Zadok, and Haim Suchowski. Robust directional couplers for state manipulation in silicon photonic-integrated circuits. *Journal of Lightwave Technology*, pages 1–1, 2022.

[16] Ido Kaplan, Muhammad Erew, Yonatan Piasetzky, Moshe Goldstein, Yaron Oz, and Haim Suchowski. Segmented composite design of robust single-qubit quantum gates. *Phys. Rev. A*, 108:042401, Oct 2023.

[17] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. A quantum engineer's guide to superconducting qubits. *Applied Physics Reviews*, 6(2):021318, 06 2019.

[18] Morten Kjaergaard, Mollie E. Schwartz, Jochen Braumüller, Philip Krantz, Joel I.-J. Wang, Simon Gustavsson, and William D. Oliver. Superconducting qubits:

Current state of play. *Annual Review of Condensed Matter Physics*, 11(Volume 11, 2020):369–395, 2020.
[19] Rainer Blatt and David Wineland. Entangled states of trapped atomic ions. *Nature*, 453(7198):1008–1015, Jun 2008.
[20] R. Ozeri, W. M. Itano, R. B. Blakestad, J. Britton, J. Chiaverini, J. D. Jost, C. Langer, D. Leibfried, R. Reichle, S. Seidelin, J. H. Wesenberg, and D. J. Wineland. Errors in trapped-ion quantum gates due to spontaneous photon scattering. *Phys. Rev. A*, 75:042329, Apr 2007.
[21] A. E. Webb, S. C. Webster, S. Collingbourne, D. Bretaud, A. M. Lawrence, S. Weidt, F. Mintert, and W. K. Hensinger. Resilient entangling gates for trapped ions. *Phys. Rev. Lett.*, 121:180501, Nov 2018.
[22] Colin D. Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M. Sage. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2):021314, 05 2019.
[23] Mark Saffman. Quantum computing with neutral atoms. *National Science Review*, 6(1):24–25, 09 2018.
[24] Loïc Henriet, Lucas Beguin, Adrien Signoles, Thierry Lahaye, Antoine Browaeys, Georges-Olivier Reymond, and Christophe Jurczak. Quantum computing with neutral atoms. *Quantum*, 4:327, September 2020.
[25] Fulvio Flamini, Nicolò Spagnolo, and Fabio Sciarrino. Photonic quantum information processing: a review. *Reports on Progress in Physics*, 82(1):016001, nov 2018.
[26] Jianwei Wang, Fabio Sciarrino, Anthony Laing, and Mark G. Thompson. Integrated photonic quantum technologies. *Nature Photonics*, 14(5):273–284, May 2020.
[27] Ido Kaplan, Haim Suchowski, and Yaron Oz. Correlation thresholds for effective composite pulse quantum error mitigation. *Physical Review A*, 110(5):052614, 2024.
[28] Khen Cohen, Haim Suchowski, and Yaron Oz. Robust photonic quantum gates with a large number of waveguide segments. *Advanced Quantum Technologies*, page e2500304.
[29] R.W. Boyd. *Nonlinear Optics*. Elsevier Science, 2020.
[30] J. E. Lang, T. Madhavan, J.-P. Tetienne, D. A. Broadway, L. T. Hall, T. Teraji, T. S. Monteiro, A. Stacey, and L. C. L. Hollenberg. Nonvanishing effect of detuning errors in dynamical-decoupling-based quantum sensing experiments. *Phys. Rev. A*, 99:012110, Jan 2019.
[31] Kevin Cox, Matthew Norcia, Joshua Weiner, Justin Bohnet, and James Thompson. Reducing collective quantum state rotation errors with reversible dephasing. *Applied Physics Letters*, 105, 07 2014.
[32] J. Randall, A. M. Lawrence, S. C. Webster, S. Weidt, N. V. Vitanov, and W. K. Hensinger. Generation of high-fidelity quantum control methods for multilevel systems. *Phys. Rev. A*, 98:043414, Oct 2018.
[33] Richard Jozsa. Fidelity for mixed quantum states. *Journal of Modern Optics*, 41(12):2315–2323, 1994.
[34] Maxim Raginsky. A fidelity measure for quantum channels. *Physics Letters A*, 290(1):11–18, 2001.
[35] Daniel Gottesman. The heisenberg representation of quantum computers. *arXiv preprint*, 1998.
[36] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70:052328, Nov 2004.
[37] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, 52:3457–3467, Nov 1995.

A. Magic State Distillation

Quantum computing has emerged as a powerful paradigm that challenges classical notions of efficiency in computation. Yet, despite significant theoretical breakthroughs and experimental progress, realizing scalable quantum computers remains a daunting task. At the heart of this challenge lies the fragility of quantum information and the limitations imposed by fundamental theorems in quantum error correction. In particular, achieving both fault-tolerance and universality requires circumventing deep structural constraints on quantum codes. This section explains how one can overcome these obstacles through the technique of MSD, which supplements Clifford-based computation with carefully prepared non-stabilizer resources.

1. Quantum Computing and Theoretical Limitations

Quantum computing holds the promise of solving certain problems exponentially faster than classical algorithms. Landmark examples include Shor's algorithm for integer factorization [1] and Grover's algorithm for unstructured search [2]. Quantum systems harness the phenomena of superposition, entanglement, and interference, enabling entirely new paradigms of information processing. A single qubit can be described by the pure state

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle\,, \quad \text{with } \alpha,\beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1. \tag{A1}$$

The Clifford group, which plays a central role in quantum error correction and classical simulation of quantum circuits, is defined as the normalizer of the Pauli group within the unitary group. It is generated by the following gates:

- The Hadamard gate: $H = \frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}$,
- The Phase gate: $S = \begin{pmatrix}1 & 0\\ 0 & i\end{pmatrix}$,
- The CNOT gate: $\mathrm{CNOT} = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes X$.

Clifford gates map Pauli operators to Pauli operators under conjugation. The Gottesman-Knill theorem [35, 36] shows that quantum circuits composed solely of Clifford gates, Pauli measurements, and stabilizer state preparations can be efficiently simulated on a classical computer. Hence, such circuits offer no computational advantage.

Moreover, despite their theoretical power, physical quantum computers face a critical challenge: decoherence and noise. Quantum systems interact unavoidably with their environment, leading to loss of coherence and errors. Unlike classical bits, qubits cannot be copied due to the no-cloning theorem, making traditional error correction inapplicable.

Quantum error correction offers a path forward. Stabilizer codes [4], a powerful family of quantum error-correcting codes (QECCs), encode logical qubits into entangled states of multiple physical qubits. For instance, the five-qubit code encodes one logical qubit and can correct any arbitrary single-qubit error. These codes rely on stabilizer generators-commuting

Pauli operators that define the code space and protect the logical information.

A key requirement for fault-tolerant quantum computation is to prevent error propagation or to control the errors' spread. Transversal gates, which apply the same unitary to each qubit in a code block, naturally suppress correlated errors. However, the Eastin-Knill theorem [3] proves that no QECC can implement a universal set of gates transversally. While Clifford gates often admit transversal implementations, at least one non-Clifford gate must be realized via a different, fault-tolerant mechanism.

2. Magic State Distillation and Universal Computation

Bravyi and Kitaev [6] proposed *magic state distillation* (MSD) as a fault-tolerant method to overcome the Eastin-Knill constraint. The key idea is to prepare special non-stabilizer states, known as *magic states*, and use them to promote the Clifford group to universality via state injection. We begin by explaining how the injection of magic states enables universal computation.

a. Universality via Ancilla-Assisted State Injection

To achieve universality, we often need to implement non-Clifford gates such as phase rotations of the form

$$\Lambda(e^{i\theta}) \equiv \begin{pmatrix} 1 & 0 \\ 0 & e^{i\theta} \end{pmatrix}. \tag{A2}$$

That is because the subgroup of U(2) generated by the gate $\Lambda(e^{i\theta})$ and the Clifford Group for single qubits is dense in U(2). Therefore, this set provides a universal basis for quantum computation [2] .

A standard technique is to use an ancilla-based circuit that probabilistically applies the desired gate by postselection. Consider the ancilla state

$$|A_\theta\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta}|1\rangle), \tag{A3}$$

and assume the initial state $|\psi\rangle \otimes |A_\theta\rangle$. A stabilizer measurement of $\sigma_z \otimes \sigma_z$ is performed, followed by a CNOT gate with the first qubit as control.

- If the measurement outcome is $+1$, the resulting state is $\Lambda(e^{i\theta})|\psi\rangle \otimes |0\rangle$.
- If the outcome is -1, the state becomes $e^{i\theta}\Lambda(e^{-i\theta})|\psi\rangle \otimes |1\rangle$.

Discarding the ancilla leaves the data qubit in a state that has undergone a known unitary-either $\Lambda(e^{i\theta})$ or $\Lambda(e^{-i\theta})$. Repeating this process allows one to simulate the desired non-Clifford gate with high probability.

b. Magic States and Fault-Tolerance

The states that can be injected and allow fault tolerant quantum computation are the magic states. A canonical magic state is the T-state:

$$\begin{aligned} |T\rangle &= \cos\beta|0\rangle + e^{i\pi/4}\sin\beta|1\rangle \\ &= \sqrt{\frac{3+\sqrt{3}}{6}}|0\rangle + e^{i\pi/4}\sqrt{\frac{3-\sqrt{3}}{6}}|1\rangle, \end{aligned} \tag{A4}$$

where $\cos^2(2\beta) = \frac{1}{3}$. It is an eigenstate of the T-gate,

$$T = e^{i\pi/4}SH = \frac{e^{i\pi/4}}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ i & -i \end{pmatrix}. \tag{A5}$$

We denote its two orthogonal eigenstates as

$$|T_0\rangle = |T\rangle, \tag{A6a}$$

$$|T_1\rangle = -e^{-i\pi/4}\sin\beta|0\rangle + \cos\beta|1\rangle. \tag{A6b}$$

Given access to $|T\rangle$, one can prepare $|T\rangle \otimes |T\rangle$, measure $\sigma_z \otimes \sigma_z$, and apply a CNOT, and then a Hadamard gate on the first qubit. For a $+1$ measurement outcome, the postselected state is

$$e^{i\gamma}|A_{-2\gamma}\rangle, \quad \text{with } \gamma = \frac{\pi}{12},$$

which enables implementation of the non-Clifford unitary $\Lambda(e^{-2i\gamma})$.

What makes the state $|T\rangle$ *magic* is not only its role in enabling universality, but also the fact that it can be distilled using the five-qubit code. Starting from five copies of the noisy mixed state [3]

$$\rho_\epsilon = (1-\epsilon)|T_0\rangle\langle T_0| + \epsilon|T_1\rangle\langle T_1|, \tag{A8}$$

a syndrome measurement projects (with probability $p(\epsilon) = \frac{\epsilon^5+5\epsilon^2(1-\epsilon)^3+5\epsilon^3(1-\epsilon)^2+(1-\epsilon)^5}{6} = \frac{1}{6}+O(\epsilon^2)$) onto the same form $\rho_{1-\epsilon'(\epsilon)}$ with reduced noise:

$$\begin{aligned} \epsilon'(\epsilon) &= \frac{\epsilon^5 + 5\epsilon^2(1-\epsilon)^3}{\epsilon^5 + 5\epsilon^2(1-\epsilon)^3 + 5\epsilon^3(1-\epsilon)^2 + (1-\epsilon)^5} \\ &= 5\epsilon^2 + O(\epsilon^3). \end{aligned} \tag{A9}$$

---

[2] The Universal Gate Set Theorem states that any unitary operation on n qubits can be approximated to arbitrary accuracy using a finite sequence of gates drawn from:

- All single-qubit unitary operations (i.e., elements of U(2)), and
- The two-qubit controlled-NOT gate (CNOT).

Therefore, the set consisting of all single-qubit gates and the CNOT gate is universal for quantum computation. This result implies that any quantum algorithm can be implemented using only single-qubit operations and entangling CNOT gates, making this set of gates a *universal basis* for quantum computation [5, 37].

[3] Why do we consider mixed states of the form $\rho_\epsilon = (1-\epsilon)|T_0\rangle\langle T_0| + \epsilon|T_1\rangle\langle T_1|$ rather than pure states of the form $\sqrt{1-\epsilon}|T_0\rangle + e^{i\phi}|T_1\rangle$ for arbitrary ϕ? While distillation is still possible in the latter case, the threshold ϵ_c becomes significantly lower. To mitigate this, one can apply a dephasing transformation to each copy of the pure state, defined as

$$D(\eta) = \frac{1}{3}\left(\eta + T\eta T^\dagger + T^\dagger \eta T\right), \tag{A7}$$

thereby obtaining the mixed state ρ_ϵ described above. The transformation D can be realized by randomly applying one of the operators I, T, or T^{-1}, each with probability $\frac{1}{3}$. This preprocessing makes the distillation process significantly more efficient.

This function is illustrated in Figure 6a. The distillation works for $\epsilon < \epsilon_c \approx 0.173$, where

$$\epsilon_c = \frac{1}{2}\left(1 - \sqrt{\frac{3}{7}}\right). \tag{A10}$$

So the five-qubit code, like every distillation protocol, exhibits a distillation threshold beyond which distillation fails to converge. This marks the maximum tolerable noise level for the input states under this scheme. Above this threshold, the output fidelity deteriorates with successive iterations, rendering the protocol ineffective for magic state purification.

By recursively applying the distillation procedure, through concatenation or nesting of the five-qubit code, the error can be suppressed to arbitrarily low levels, enabling the preparation of magic states with arbitrarily high fidelity. Each iteration corresponds to a level of code nesting; that is, performing n iterations implies that each logical qubit requires 5^n physical qubits due to recursive encoding.

We denote by ϵ_{c_i} the critical value of ϵ at which the protocol transitions from requiring $i-1$ iterations to i iterations. In Figure 6b, we present the required number of distillation iterations, using the five-qubit code, to reduce the physical error rate ϵ below the fault-tolerance threshold of 10^{-15}, plotted on a logarithmic scale. The following are the critical error values ϵ_{c_i} for the first ten iteration thresholds:

$$\begin{aligned}
&\epsilon_{c_1} = 10^{-15} \quad , \ \epsilon_{c_2} = 1.414 \times 10^{-8}, \\
&\epsilon_{c_3} = 5.318 \times 10^{-5}, \ \epsilon_{c_4} = 3.251 \times 10^{-3}, \\
&\epsilon_{c_5} = 2.490 \times 10^{-2}, \ \epsilon_{c_6} = 6.676 \times 10^{-2}, \\
&\epsilon_{c_7} = 1.072 \times 10^{-1}, \ \epsilon_{c_8} = 1.353 \times 10^{-1}, \\
&\epsilon_{c_9} = 1.520 \times 10^{-1}, \ \epsilon_{c_{10}} = 1.615 \times 10^{-1}.
\end{aligned} \tag{A11}$$

These thresholds illustrate the sharp trade-off between input noise and resource overhead: as ϵ increases, the number of required iterations (and thus the number of physical qubits per logical qubit) increases rapidly. For instance, an initial error just below 10^{-2} requires four levels of distillation, resulting in $5^4 = 625$ physical qubits per logical qubit, while an error around 10^{-1} would require seven or more iterations, consuming over $78,000$ qubits. This exponential growth in resource overhead underscores the importance of reducing the initial magic state error as much as possible prior to distillation, as even small improvements can drastically reduce the number of required physical qubits.

Figure 1 provides an overview to illustrate these ideas. In panel (a), five noisy input states of the form $\rho_T(\epsilon) = (1-\epsilon)\left|T_0\right\rangle\left\langle T_0\right| + \epsilon\left|T_1\right\rangle\left\langle T_1\right|$ are consumed in a single round of distillation. With probability $p(\epsilon) = \frac{1}{6} + O(\epsilon)$, the protocol outputs a logical state $\rho_T^{(L)}(\epsilon')$ whose error is quadratically suppressed, $\epsilon'(\epsilon) = 5\epsilon^2 + O(\epsilon^3)$. Thus, the protocol achieves both error reduction and probabilistic success, at the cost of increased resource overhead. The technical details of the distillation circuit are depicted in panel (b). The input state can be written in the coherent form $|\psi\rangle = \sqrt{1-\epsilon}\left|T_0\right\rangle + e^{i\phi}\sqrt{\epsilon}\left|T_1\right\rangle$. After a random dephasing step, implemented by applying I, T, or T^2 with equal probability, it reduces to the mixed form $\rho_T(\epsilon)$. This mixed state serves as the standard input to the stabilizer-based distillation procedure. Panels (c) and (d) characterize the performance of the protocol. Panel (c) shows the distillation curve $\epsilon'(\epsilon)$, making explicit how the output error depends on the input. Panel (d) highlights the recursive structure of MSD: by concatenating multiple rounds, one can reduce the error below the fault-tolerance threshold (here taken to be 10^{-15}). However, the required number of rounds grows rapidly with the initial error, reflecting the exponential qubit overhead inherent to this process.

Another notable single-qubit magic state is the Hadamard eigenstate:

$$\left|H\right\rangle = \cos\left(\frac{\pi}{8}\right)\left|0\right\rangle + \sin\left(\frac{\pi}{8}\right)\left|1\right\rangle. \tag{A12}$$

Similar to $\left|T\right\rangle$, this state can be distilled and used to implement non-Clifford operations. In particular, noting that

$$\left|H\right\rangle = e^{i\pi/8} S^3 H \left|A_{-\pi/4}\right\rangle, \tag{A13}$$

reveals that $\left|H\right\rangle$ enables the injection of the non-Clifford unitary

$$\Lambda(e^{-i\pi/4}) = \begin{pmatrix} 1 & 0 \\ 0 & e^{-i\pi/4} \end{pmatrix}, \tag{A14}$$

which lies outside the Clifford group and is commonly referred to as the T gate in other contexts.

The set of magic states for single qubits consists of $\left|T\right\rangle$, $\left|H\right\rangle$, and their Clifford equivalents. These are visualized on the Bloch sphere in Figure 7.

3. Toward Fault-Tolerant Universality

To summarize, combining the following ingredients enables fault-tolerant, universal quantum computation:

- Stabilizer codes to protect quantum information,
- Clifford gates implemented transversally,
- MSD to generate high-fidelity non-stabilizer states,
- State injection to realize non-Clifford gates.

While the resource overhead of MSD remains substantial, ongoing research seeks more efficient protocols. MSD continues to play a central role in the quest for scalable, fault-tolerant quantum computing.

B. Physical Realizations of Qubits and Operations for Quantum Computing

In this appendix, we provide a brief overview of several leading physical platforms for realizing qubits and implementing single-qubit operations. These include superconducting circuits, trapped ions, neutral atoms, and integrated photonic devices. Each platform offers distinct control mechanisms - from microwave or laser-driven rotations in matter-based systems to coherent interference in photonic waveguides- and each suffers from characteristic error sources such as amplitude miscalibration, detuning, crosstalk, or fabrication imperfections. Understanding these platforms and their dominant error models establishes the foundation for the composite schemes introduced in the main text, which are designed to mitigate such systematic errors and enable more robust implementations of non-Clifford gates like $\mathcal{T}$ and $\mathcal{H}$.

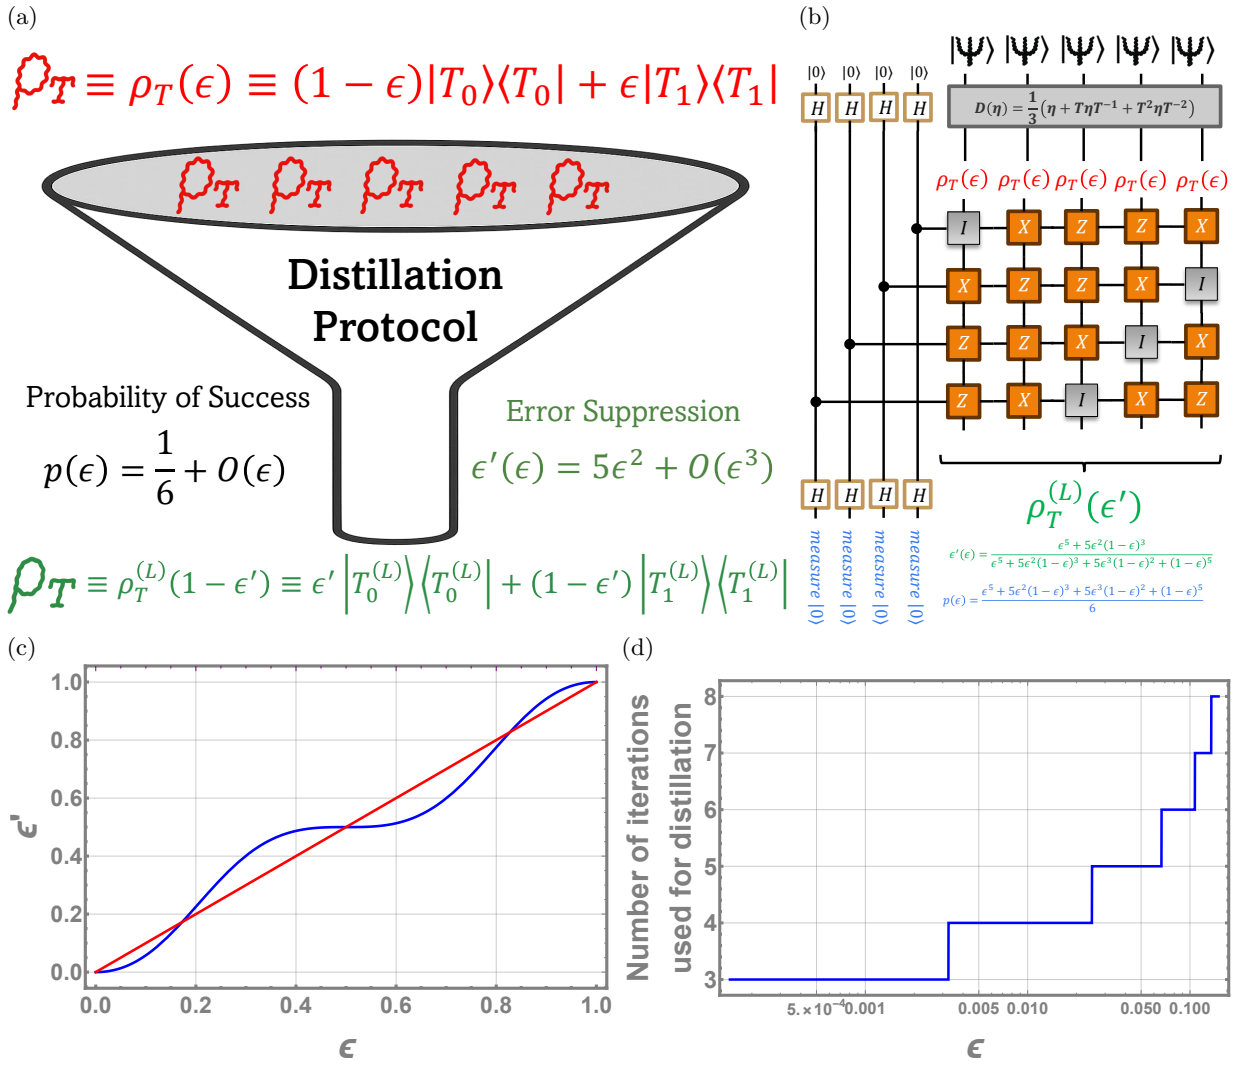

FIG. 6. Five-qubit T-state distillation protocol. (a) Conceptual schematic of a single round of the T-magic-state distillation protocol. One round consumes five noisy copies of $\rho_T(\epsilon) \equiv (1-\epsilon)\,|T_0\rangle\langle T_0| + \epsilon\,|T_1\rangle\langle T_1|$ (denoted by ρ_T with heavily wavy symbols to emphasize noise). With probability $p(\epsilon) = \frac{1}{6} + O(\epsilon)$, the protocol outputs the logical state $\rho_T^{(L)}(\epsilon')$, whose error is quadratically suppressed, $\epsilon'(\epsilon) = 5\epsilon^2 + O(\epsilon^3)$. We represent this improved state using symbols with reduced waviness, reflecting its lower noise. (b) Circuit-level details of the distillation protocol (see Appendix A for a full derivation). The input state is the coherent superposition $|\psi\rangle = \sqrt{1-\epsilon}\,|T_0\rangle + e^{i\phi}\sqrt{\epsilon}\,|T_1\rangle$, which, after dephasing via random application of I, T, or T^2 (each with probability 1/3), reduces to the mixed state $\rho_T(\epsilon)$. (c) Distillation curve $\epsilon'(\epsilon)$: the output error rate as a function of the input error rate, showing quadratic suppression after one round. (d) Required number of recursive distillation rounds to achieve an error below the 10^{-15} fault-tolerance threshold, illustrating the exponential scaling of qubit overhead with the initial error ϵ.

1. Superconducting Qubits

In superconducting qubits (e.g., transmons), quantum states are manipulated via resonant microwave drives. A typical driven Hamiltonian in the rotating frame is

$$H(t) = \frac{\Omega(t)}{2}\left(\cos\phi\, X + \sin\phi\, Y\right), \tag{B1}$$

where $\Omega(t)$ is the Rabi frequency envelope (determined by pulse amplitude), and ϕ is the microwave phase. The resulting unitary is a rotation:

$$U(\theta,\phi) = \exp\left(-i\frac{\theta}{2}(\cos\phi\, X + \sin\phi\, Y)\right), \tag{B2}$$

with total angle $\theta = \int \Omega(t)dt$. To implement $\mathcal{T}$, one chooses $\theta = 2\beta$ and $\phi = 3\pi/4$ to realize a rotation in the equatorial plane with the desired complex phase. The $\mathcal{H}$ gate, involving a rotation by $\pi/4$, is typically synthesized through a sequence of calibrated pulses, e.g., using a combination of $R_x(\theta)$ and $R_z(\phi)$ rotations.

Typical errors in the physical realization of quantum gates include amplitude errors ($\theta \to \theta + \delta\theta$), which are caused by inaccuracies in pulse duration or power and result in over- or under-rotation; phase errors ($\phi \to \phi + \delta\phi$), which arise from miscalibration of the phase reference frame or IQ imbalance in the control electronics; and leakage errors, which correspond to unintended transitions to noncomputational energy

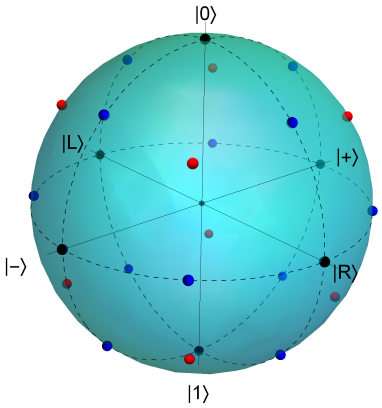

FIG. 7. Representation of stabilizer and magic states for a single qubit plotted on the Bloch sphere. Black points denote stabilizer states, which form the vertices of the stabilizer polytope. Red points correspond to the non-stabilizer T-states, and blue points represent the H-states, both of which are essential resources for universal quantum computation. These magic states lie outside the stabilizer polytope and enable non-Clifford operations via state injection.

levels outside the qubit subspace. Composite pulse schemes like BB1 or SK1 suppress systematic amplitude errors. For example, the BB1 sequence corrects a rotation $R_\phi(\theta)$ using a sequence of four additional rotations:

$$\begin{aligned} U_{\text{BB1}} =& R_\phi(\theta/2) R_{\phi+\arccos(-\theta/4\pi)}(\pi) \\ & R_{\phi+3\arccos(-\theta/4\pi)}(2\pi) R_{\phi+\arccos(-\theta/4\pi)}(\pi) R_\phi(\theta/2). \end{aligned} \tag{B3}$$

The amplitude error is largely global across composite gates (same $\delta\theta/\theta$), but phase errors may be gate-specific depending on the control electronics.

2. Trapped Ions

In trapped ion systems, qubits are encoded in the internal energy levels of ions (e.g., $^{40}Ca^+$ or $^{171}Yb^+$). Laser pulses drive Rabi oscillations, enabling high-fidelity single-qubit gates. The effective Hamiltonian for a Raman transition is:

$$H = \frac{\Omega}{2}\left(e^{i\phi}\left|0\right\rangle\left\langle 1\right| + e^{-i\phi}\left|1\right\rangle\left\langle 0\right|\right), \tag{B4}$$

leading to evolution:

$$U(\theta,\phi) = \begin{pmatrix} \cos(\theta/2) & -ie^{i\phi}\sin(\theta/2) \\ -ie^{-i\phi}\sin(\theta/2) & \cos(\theta/2) \end{pmatrix}. \tag{B5}$$

Here, the laser phase ϕ and pulse area θ control the rotation axis and angle. By tuning the laser parameters-detuning, amplitude, and phase- rotations like those required for $\mathcal{T}$ and $\mathcal{H}$ can be implemented. A Raman beam pair can create an effective Hamiltonian that induces the desired unitary operation with a phase corresponding to $e^{i\pi/4}$, thus implementing $\mathcal{T}$. The $\mathcal{H}$ gate is often realized via composite pulse sequences or through direct single-pulse methods.

Typical errors in the physical realization of quantum gates in trapped-ion systems include laser intensity fluctuations, where variations in laser power modify the effective Rabi frequency Ω, resulting in rotation angle errors ($\theta \to \theta + \delta\theta$); laser detuning (off-resonant driving), where frequency mismatch between the laser and the qubit transition introduces over- or under-rotation and undesired phase accumulation, often modeled as spurious Z rotations; and AC Stark shifts, where off-resonant couplings to auxiliary levels induce systematic shifts in the qubit energy levels, effectively contributing additional unwanted phases to the gate evolution. Robust sequences like CORPSE or SCROFULOUS cancel detuning and amplitude errors. For instance, CORPSE decomposes a π pulse into three rotations with over- and under-rotation to cancel detuning effects. Errors like detuning are typically global to all gate components, but intensity noise may vary slightly across segments, especially in pulse-shaping scenarios.

3. Neutral Atoms

Neutral atoms trapped in optical tweezers or lattices encode qubits in their hyperfine ground states. Microwave or Raman transitions are used to implement single-qubit gates. The Hamiltonian has a similar form to trapped ions:

$$H = \frac{\Omega(t)}{2}\left(\cos\phi\, X + \sin\phi\, Y\right), \tag{B6}$$

where the control parameters are again laser phase and pulse area. Arbitrary gates are synthesized using:

$$U = R_z(\gamma) R_y(\theta) R_z(\delta), \tag{B7}$$

where $R_y(\theta)$ and $R_z(\phi)$ are physical operations: R_y by microwave pulses, and R_z often via virtual frame updates. Arbitrary unitaries like $\mathcal{T}$ and $\mathcal{H}$ are realized using laser pulses with precisely controlled phase and amplitude. The $\mathcal{T}$ gate can be realized by choosing angles to match a rotation axis with complex coefficients (e.g., a rotation in the XY-plane with global phase $e^{i\pi/4}$).

Typical errors in the physical realization of quantum gates with neutral atom qubits include intensity fluctuations, where variations in laser or microwave power alter the effective pulse area and lead to rotation angle errors ($\theta \to \theta + \delta\theta$); frequency drift, where temporal fluctuations in the driving field frequency result in detuning from the qubit transition, effectively introducing unwanted Z rotations and dephasing; and crosstalk, where imperfect spatial addressing causes unintentional excitation of neighboring atoms, leading to correlated errors and decoherence. To mitigate systematic amplitude and detuning errors, composite pulse sequences such as BB1, CORPSE, and Walsh-modulated schemes are commonly

employed. These techniques are designed under the assumption that the dominant errors are coherent and approximately global across the entire pulse train, thereby enabling robust error cancellation through constructive interference of control segments.

4. Integrated Photonics and Directional Couplers

In integrated photonic quantum computing, qubits are frequently encoded in the path degree of freedom of single photons using planar waveguide structures. A central component in such systems is the *directional coupler*, which enables coherent interference between two waveguides via evanescent coupling.

Under coupled-mode theory, the evolution of the optical field amplitudes $E_1(t)$ and $E_2(t)$ in a directional coupler of fixed cross-section is governed by the unitary propagator:

$$\begin{pmatrix} E_1(t) \\ E_2(t) \end{pmatrix} = e^{-i\frac{t}{2}(\Omega X - \Delta Z)} \begin{pmatrix} E_1(0) \\ E_2(0) \end{pmatrix}, \tag{B8}$$

where Ω is the coupling coefficient (primarily dependent on the inter-waveguide spacing), Δ denotes the detuning or phase mismatch (determined by differences in propagation constants), and t is the propagation length along the coupler.

The parameters Ω and Δ are real-valued functions determined by the physical geometry. For silicon-on-insulator rib waveguides, they depend on the widths of the individual waveguides (w_1, w_2), their heights (H_1, H_2), the gap between the waveguides (g), the etch depths (h_1, h_2), and the coupling length (t). When the heights, etch depths, and gap are fixed and symmetric, the dependence simplifies to:

$$\Omega = \Omega(w_1, w_2), \tag{B9a}$$
$$\Delta = \Delta(w_1, w_2). \tag{B9b}$$

This correlated and nontrivial dependence makes traditional composite pulse schemes, which typically assume independent control over rotation axes, nontrivial to apply. Nevertheless, photonic implementations employ alternative strategies to enhance robustness. Post-fabrication tuning using integrated heaters or phase shifters can dynamically adjust phase and coupling to correct static errors. Symmetric design, such as balanced Mach-Zehnder interferometers (MZIs), reduces sensitivity to fabrication deviations. Robust architectures incorporate circuit redundancy and heralding techniques to mitigate the impact of parameter variations and losses. Finally, programmable photonics employs configurable interferometers with adjustable elements to enable adaptive calibration and real-time error correction.

Composite directional coupler schemes offer an alternative approach by improving error resilience through careful static design, without relying on active control. Although these designs increase circuit footprint and complexity, they are well-suited for passive platforms where dynamic reconfiguration is limited or undesired.

Crucially, due to the generator structure in Eq. (B8), only X and Z appear in the effective Hamiltonian. Implementing gates involving Y, such as the $\mathcal{T}$ gate, thus requires at least two segments with varying parameters. In Section VI, we introduce composite designs for the $\mathcal{T}$ gate that is robust against fully correlated width-dependent errors, following the construction developed in [16].

C. Plots for the T-Magic Errors in the Composite Schemes

To provide a deeper understanding of the role composite pulse sequences play in improving non-Clifford gate performance, we include here additional figures that focus specifically on the behavior of the T-magic error. These plots offer a more direct and operationally meaningful view of gate quality, particularly in relation to magic state distillation thresholds.

Figures 8, 9, and 10 demonstrate the effectiveness of composite pulse schemes in enhancing the fidelity of the $\mathcal{T}$ gate under global systematic errors. Each plot shows the T-magic error as a function of the global physical error ϵ, both on logarithmic scales. The figures correspond to the designs developed in Sections IV, V, and VI, respectively, with color schemes matched to the design-specific curves introduced in each section. In all cases, the composite pulse sequences exhibit significantly improved error suppression compared to the single-pulse implementation, especially in the low-error regime.

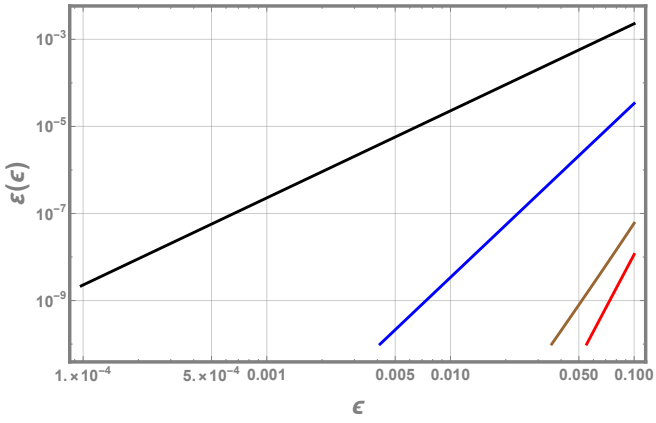

FIG. 8. T-magic error for the $\mathcal{T}$ gate as a function of the global physical error ϵ in X-Y control systems with global Rabi frequency errors. Each curve corresponds to a composite sequence with a different number of pulses: single-pulse (black), three-pulse (blue), five-pulse (red), and seven-pulse (brown). All values are plotted on logarithmic scales. Composite pulse sequences significantly reduce the T-magic error, particularly in the low-error regime, thereby improving gate fidelity before any distillation is applied.

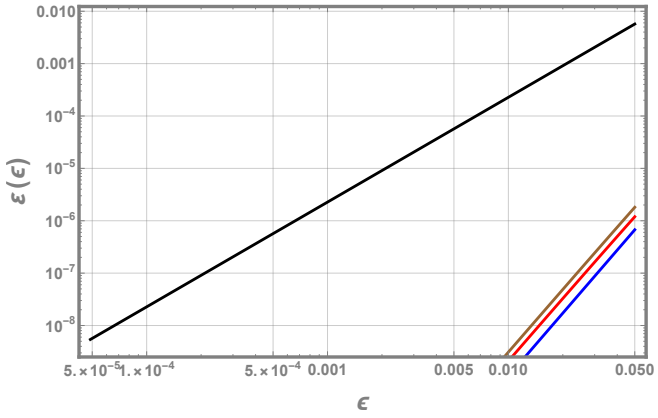

FIG. 9. T-magic error as a function of the global physical error ϵ under correlated Z errors in X-Z Hamiltonians. The black curve represents the two-segment baseline, while the red, blue, and brown curves correspond to three distinct three-pulse composite designs labeled (a), (b), and (c). Composite sequences achieve notable suppression of the T-magic error across a broad range of error values, outperforming the baseline design.

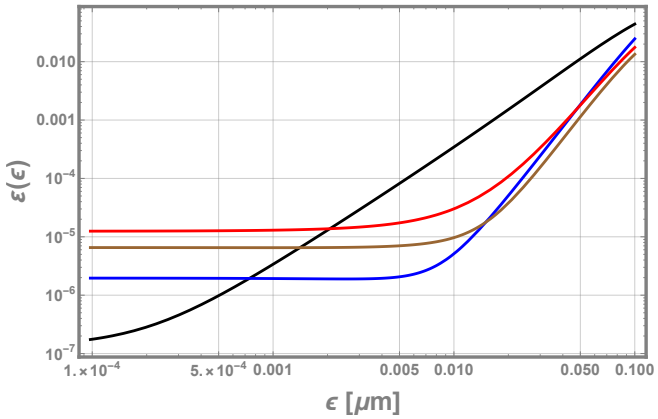

FIG. 10. T-magic error plotted versus the global physical error ϵ for directional coupler implementations in integrated photonics, where imperfections in width introduce global systematic errors. The black curve shows the performance of the two-segment design, and the blue, red, and brown curves represent four-segment composite schemes (a), (b), and (c), respectively. Although the composite schemes may slightly underperform the baseline at very small ϵ, they demonstrate superior robustness across the physically relevant error range.

Figures 11, 12, and 13 provide a complementary perspective by plotting the base-10 logarithm of the $\mathcal{T}$-gate T-magic error (see Eq. (22)) against the base-10 logarithm of the error associated with the corresponding single-pulse implementation. The steeper slopes observed for the composite pulse curves indicate superior error scaling, underscoring the effectiveness of composite control in improving gate robustness- a critical requirement for fault-tolerant quantum computation.

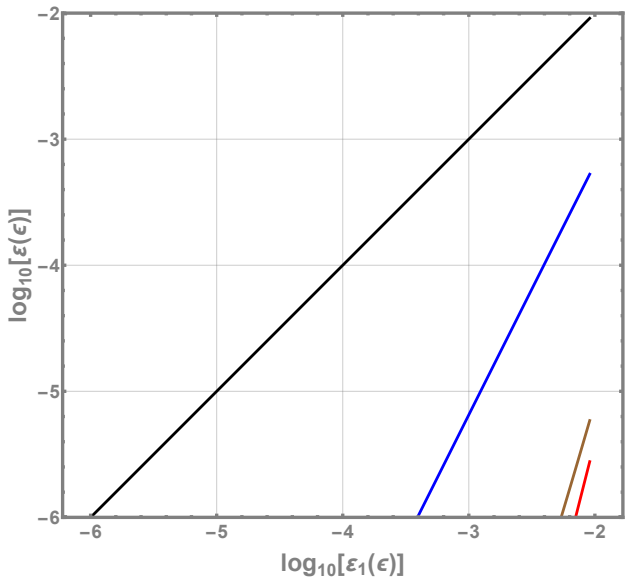

FIG. 11. Log-log plot comparing the T-magic error of composite $\mathcal{T}$ gate implementations to that of the single-pulse gate in X-Y control systems with global Rabi frequency errors. The x-axis shows the base-10 logarithm of the error for the single-pulse implementation, and the y-axis shows the logarithm of the T-magic error for each design. The steeper slopes of the multi-pulse curves indicate improved error scaling and robustness, essential for reducing distillation overhead in practical settings.

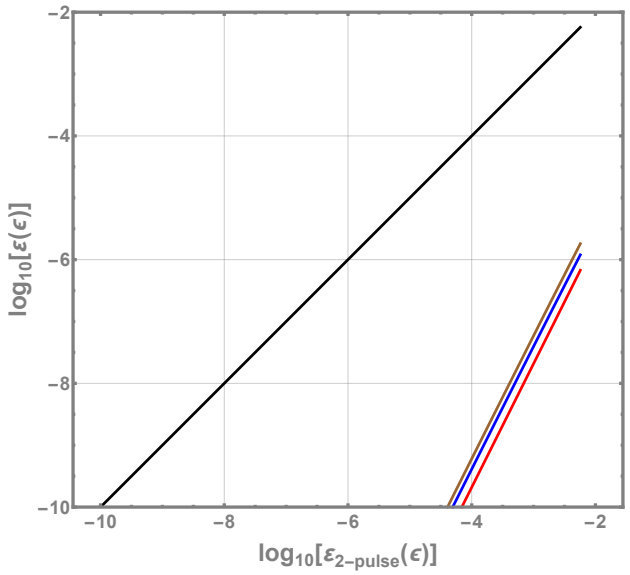

FIG. 12. Log-log comparison of the T-magic error for three-pulse composite sequences versus the baseline two-segment gate in X-Z systems with global Z errors. Each point reflects the composite scheme's T-magic error plotted against that of the corresponding single-pulse gate. The improved slope indicates enhanced scaling behavior of the composite designs, enabling a lower distillation burden in quantum error correction protocols.

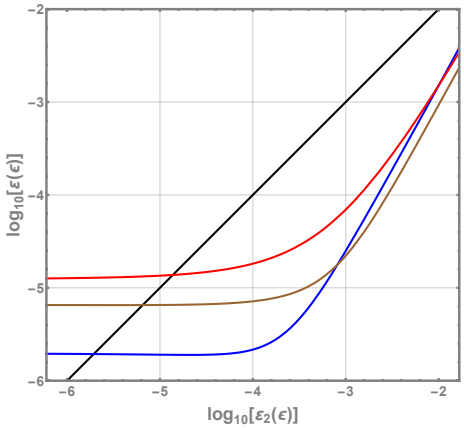

FIG. 13. Logarithmic comparison of T-magic error scaling for $\mathcal{T}$ gate implementations in integrated photonic circuits with directional couplers. The plot compares composite pulse schemes (a, b, c) to the baseline two-segment design in terms of T-magic error suppression relative to the single-pulse gate error. Despite similar performance in the low-error limit, the composite schemes (especially scheme (c)) show better scaling and reduced distillation demands across the full error range as discussed in the main text.

These six figures complement the results presented in the main text, where we plotted the Frobenius fidelity, the trace fidelity, and the corresponding reduction in the number of distillation iterations required to reach a target gate error of 10^{-15}. That distillation overhead is directly determined by the T-magic error, which is the quantity plotted in the figures shown here.

D. Plots for the Frobenius Fidelity and the T-Magic Fidelity

In this appendix, we illustrate the behavior of the newly introduced fidelity measure, the *T-magic fidelity*, by comparing it to the conventional Frobenius fidelity across a broad range of error values. This comparison is intended to provide the reader with intuition for how this operationally motivated measure behaves and how it differs from standard fidelity metrics. Notably, the T-magic fidelity does not necessarily exhibit monotonic behavior with respect to traditional gate fidelities, highlighting its distinct functional character and the specific insights it offers. Recall that the T-magic fidelity is designed to quantify how effectively a given implementation of the $\mathcal{T}$ gate prepares the target magic state when acting on the initial state $|0\rangle$ for MSD.

Figures 14, 15, and 16 show this comparison under global systematic errors for the composite designs discussed in Sections IV, V, and VI, respectively. In each case, solid lines represent the Frobenius fidelity, while dashed lines correspond to the T-magic fidelity, with color coding consistent with the composite pulse schemes introduced in the respective sections. These plots clearly reveal the non-trivial relationship between the two measures and underscore the unique role of the T-magic fidelity in evaluating the gate's suitability for MSD.

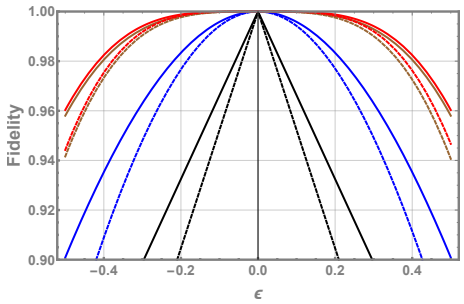

FIG. 14. Comparison between the Frobenius fidelity (solid curves) and the T-magic fidelity (dashed curves) for $\mathcal{T}$ gate implementations in X-Y control systems with global Rabi frequency errors. Each curve corresponds to a composite sequence with a different number of pulses: single-pulse (black), three-pulse (blue), five-pulse (red), and seven-pulse (brown).

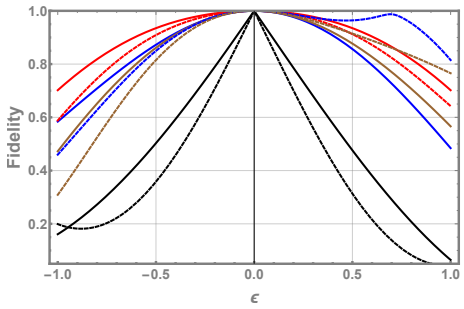

FIG. 15. Comparison of Frobenius fidelity (solid lines) and T-magic fidelity (dashed lines) for $\mathcal{T}$ gate implementations under correlated systematic Z errors in X-Z Hamiltonians. The black curve represents the two-segment baseline, while the red, blue, and brown curves correspond to three distinct three-pulse composite designs labeled (a), (b), and (c).

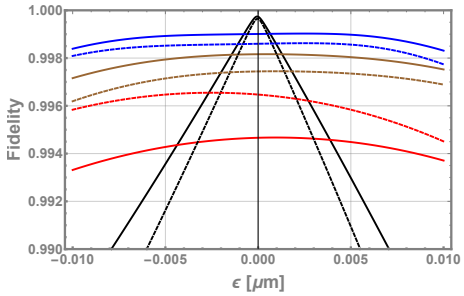

FIG. 16. Frobenius fidelity (solid) versus T-magic fidelity (dashed) for $\mathcal{T}$ gate implementations using directional couplers in integrated photonics, where global systematic width errors dominate. The black curve shows the performance of the two-segment design, and the blue, red, and brown curves represent four-segment composite schemes (a), (b), and (c), respectively.