

Leveraging biased noise for more efficient quantum error correction at the circuit level with two-level qubits

Josu Etxezarreta Martinez^{1,2,*} Paul Schnabl^{1,3} Javier Oliva del Moral^{1,4} Reza Dastbasteh¹
Pedro M. Crespo¹ and Ruben M. Otxoa⁵

¹Department of Basic Sciences, Tecnun—University of Navarra, 20018 San Sebastian, Spain

²Cavendish Laboratory, Department of Physics, University of Cambridge, Cambridge CB3 0HE, United Kingdom

³Institute for Theoretical Physics, University of Innsbruck, A-6020 Innsbruck, Austria

⁴Donostia International Physics Center, 20018 San Sebastian, Spain

⁵Hitachi Cambridge Laboratory, J. J. Thomson Avenue, Cambridge CB3 0HE, United Kingdom

 (Received 26 May 2025; revised 14 December 2025; accepted 16 December 2025; published 9 January 2026)

Tailoring quantum-error-correction codes (QECCs) to biased noise has demonstrated significant benefits. However, most of the prior research on this topic has focused on code-capacity noise models. Furthermore, a no-go theorem prevents the construction of controlled-NOT (CNOT) gates for two-level qubits in a bias-preserving manner which may, in principle, imply that noise bias cannot be leveraged in such systems. In this work, we show that a residual bias up to $\eta \sim 5$ can be maintained in CNOT gates under certain conditions. Moreover, we employ controlled-Z (CZ) gates in syndrome-extraction circuits and show how to natively implement these in a bias-preserving manner for a broad class of qubit platforms. This motivates the introduction of what we call a *hybrid biased-depolarizing* (HBD) circuit-level noise model that captures these features. We numerically study the performance of the $XZZX$ surface code and observe that bias-preserving CZ gates are critical for leveraging biased noise. Accounting for the residual bias present in the CNOT gates, we observe an increase in the code threshold up to a 1.27% physical error rate, representing a 90% improvement. Additionally, we find that the required qubit footprint can be reduced by up to 75% at relevant physical error rates.

DOI: [10.1103/q7w6-nljp](https://doi.org/10.1103/q7w6-nljp)

I. INTRODUCTION

The extremely fragile nature of quantum information implies that quantum-error-correction (QEC) techniques are going to be necessary to obtain the full benefits that quantum computing offers theoretically [1,2]. In the quest of fault tolerance, many technologies are being explored as candidates to host the error-corrected quantum computers of the future. Superconducting junctions, neutral atoms, ion traps, or silicon spin quantum dots are some examples of these technologies [3].

Interestingly, some of these qubit technologies present a strong bias toward dephasing errors, i.e., the probability of experiencing Pauli-Z (phase-flip) errors is much higher than of experiencing X (bit-flip) and Y (bit- and phase-flip)

errors. This occurs for qubits with much longer relaxation times than dephasing times, $T_1 \gg T_2$. The bias or degree of asymmetry is usually quantified by the parameter $\eta = p_z/(p_x + p_y) \approx T_1/T_2 - 1/2$ [4–6]. This is, e.g., the case for trapped-ion qubits [7–10], silicon spin qubits [7,11–14], nitrogen-vacancy- (N-V) center qubits [15,16], or certain superconducting-qubit architectures [5,7,17–19]. Intuitively, errors that are strongly biased toward dephasing should be helpful for QEC, since the entropy of the error source is smaller than for the symmetric or depolarizing case [4,20]. Indeed, bias tailoring of QECCs has resulted in great performance improvements when the noise experienced by qubits presents such a feature [4,5,21–24]. The $XZZX$ surface code is a notable example of such bias tailoring, exhibiting a threshold of approximately 48% at a finite and experimentally relevant bias of $\eta = 1000$ [4]. This significantly surpasses the 10.9% threshold of this code [and for the Calderbank-Shor-Steane (CSS) surface code] for the depolarizing case.

However, most of the bias tailoring of QECCs has been done considering code-capacity noise, i.e., considering that only the data qubits of the code experience errors and that the syndrome-extraction circuits are ideal [2]. In

*Contact author: jetxezarreta@unav.es

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

reality, all elements involved in the syndrome extraction are faulty and introduce additional errors. This noise model is referred to as *circuit-level noise*. For this model, each noisy operation in the syndrome-extraction circuits is modeled as an ideal operation followed by a noise channel, i.e., if $\tilde{U}$ is a noisy gate, then it is modeled as $\mathcal{N} \circ \mathcal{U}$. While one may naively think that in a system with biased noise all the noisy operations can be modeled as ideal gates followed by biased-noise channels, i.e., as to be *bias preserving*, this assumption does not hold in general. A no-go theorem in Ref. [25] demonstrates that a bias-preserving controlled-NOT (CNOT) gate cannot be implemented between two qubits encoded in finite-dimensional Hilbert spaces (we will refer to these as *two-level qubits* throughout the text, as in Ref. [26]). This result is a huge impediment for bias tailoring QECCs due to the fundamental role of the CNOT gate in syndrome-extraction circuits.

In this context, several alternatives to circumvent the no-go theorem in Ref. [25] have been proposed in groundbreaking results for bias-preserving CNOT gates in superconducting cat qubits [27] and for certain neutral-atom qubit processors [28]. As a result, cat-qubit technologies have received special attention [29]. Furthermore, huge qubit-overhead reductions have been reported theoretically as a result of the combination of bias tailoring with quantum low-density-parity-check codes (qLDPC) over cat qubits [30,31]. For example, the calculations in Ref. [31] predict that around 10^5 cat qubits can be enough to run Shor's algorithm to factor 2048 RSA (Rivest, Shamir and Adleman) integers, while surface-code-based QEC over a platform with symmetric noise would require almost 20×10^6 physical qubits [32]. That is a 200 times overhead reduction. In the quest for constructing the first *Megaquop machine*, or a quantum computer able to realize millions of error-free operations, such overhead savings may be crucial for certain technologies to surpass others [33]. The construction of a Megaquop machine consisting of around 100 logical qubits is considered to be the next milestone for the field of quantum computing, posing the end of the so-called noisy intermediate-scale quantum (NISQ) era.

Nevertheless, a question remains open: can we leverage the natural bias toward dephasing noise present in certain platforms based on two-level qubits (or qubits encoded in finite-dimensional systems) even if the CNOT gates implemented in those are not bias preserving? This has lost some of the attention of the community beyond cat qubits, mainly due to the limitation of constructing bias-preserving CNOT gates in such platforms, leading to the study of QECCs over depolarizing circuit-level noise models whenever two-level qubit platforms are considered [26,30,31,34,35]. Furthermore, it is also important to state that while possible, the experimental implementation of bias-preserving CNOT gates on cat-qubit platforms is very challenging and remains an open question [36]. In fact, note that the recent experimental demonstration of a

repetition code in a hybrid cat-transmon architecture uses CNOT gates with high bias (higher than 25), but which are not bias preserving [37].

In this work, we aim to answer to such a question, focusing on which elements of the syndrome-extraction circuit can, in fact, be executed in a bias-preserving manner so that dephasing noise can still be predominant enough in the system that bias tailoring can still be successful. Specifically, we begin by defining what bias-preserving gates are, and then discuss which of the usual gates used for syndrome extraction are bias preserving and which are not. We study the residual or remainder bias of CNOT [and controlled-Y (CY)] gates by means of a Lindblad equation and determine that while most of the bias is destroyed, there is still a certain bias in their noise channels. Also, we observe that Hadamard gates fully depolarize the noise bias. We then focus on the importance of controlled-Z (CZ) gates in syndrome-extraction protocols and how those gates can be constructed in a bias-preserving manner across several quantum technologies. Furthermore, we discuss the importance of implementing the CZ gate in a bias-preserving manner at the native-hardware level. As a result, we propose a *hybrid biased-depolarizing circuit-level noise model* (HBD) that captures the noise features of the qubit platforms under discussion in a generic way. In this model, we consider that CZ gates are bias preserving while the rest of the gates strictly introduce fully symmetric, i.e., depolarizing, noise. In order to determine if tailored codes can leverage noise bias under this noise model, we numerically study the performance of the *XZZX* variant of the rotated surface code. We compare the obtained performance with the case in which the CZ gates are not natively constructed in a bias-preserving manner. We observe almost no improvement of the threshold for the latter case as a function of bias. This starkly contrasts with the bias-preserving case, with notable improvement in performance (around 40%). We consider this generic HBD circuit-level noise model to be a valuable tool for determining whether certain bias-tailored codes can exhibit performance improvements when implemented on a two-level qubit technology.

Moreover, motivated by our analysis of the remainder bias noise present in CNOT gates, we numerically study the rotated *XZZX* surface code over a more fine-grained HBD circuit-level noise model for which the CNOT gates do exhibit this residual bias. The main rationale to do so resides in the fact that a small bias does result in significant threshold increase whenever biased code-capacity models are considered [4]. These results confirm our intuition, with a code threshold surpassing 1.2% whenever $\eta \geq 100$, which is a 90% boost relative to the standard depolarizing (SD) case. We continue our study by numerically exploring the mega-, giga-, and teraquop footprints for two experimentally relevant physical error rates, $p = 0.003$ and $p = 0.001$. We observe that as a result of the bias

tailoring, footprint reductions ranging from 33% up to 75% are obtained for all relevant quantum operation regimes. We envision that these results are relevant for practical QEC since huge physical-qubit savings can be obtained.

II. SYNDROME-EXTRACTION CIRCUITS AND BIAS-PRESERVING GATES

QECCs rely on extraction circuits to obtain the syndrome that the decoder will use to estimate the error that has occurred and therefore propose a recovery operator [2]. Those extraction circuits refer to measuring the stabilizers defining the code and are generally conducted by means of two-qubit entangling gates between the check qubit and the data qubits, the parity of which is captured by such an ancilla. In general, a stabilizer measurement is done by means of a Hadamard test in which the controlled unitary refers to the stabilizer being measured by the check qubit in question [6]. Since the stabilizers are Pauli strings, the controlled stabilizer operation is split into CNOT (X -component), CZ (Z -component), and controlled- Y (CY) (Y -component) two-qubit gates entangling the associated data qubits and the respective check. It is well known that not all gates can be executed in a bias-preserving manner (for a formal definition of bias-preserving gates, see Appendix A). As said above, the ubiquitous CNOT gate in syndrome-extraction circuits cannot be implemented in a bias-preserving manner for two-level qubits [25,27]. As a result, it is a general conception that the benefits of bias tailoring are effectively lost for two-level qubits with natural biased noise. In the following, we discuss this issue and analyze the scenario for CZ gates. As we discuss in Appendix B, our analysis here assumes continuous noisy time evolution with a static Hamiltonian under the Born-Markov and uncorrelated (white-) noise approximation with a specific structure.

A. CNOT gates (and CY)

CNOT gates are essential gates in order to measure X stabilizers and, thus, detect possible phase-flip errors. However, these gates cannot be constructed in a bias-preserving manner for two-level qubits [25]. As explained in Ref. [27], the continuous evolution required to implement a CNOT gate can convert phase errors into Y errors [26], resulting in a reduction of the bias of the noise toward dephasing. Specifically, a CNOT gate can be implemented with an interaction described by the Hamiltonian [27]

$$H_{\text{CNOT}} = V \left[\frac{I+Z}{2} \otimes I + \frac{I-Z}{2} \otimes X \right], \quad (1)$$

where V refers to the strength of the interaction and the CNOT gate is realized at time $VT = \pi/2$, up to a global phase. During the continuous noisy evolution described by this interaction, some of the phase flips occurring in




FIG. 1. The remainder bias toward dephasing errors when implementing a CNOT gate with an interaction in Eq. (1). We set the coupling rates of the dissipators so that the gate fidelity is around 99.7% and change the system bias η_{sys} . We name the residual bias after the CNOT gate as η_{CNOT} .

the target qubit are transformed due to the anticommutation between the target term and the predominant Z errors, implying that a CNOT gate cannot be implemented in a bias-preserving manner, i.e., $\tilde{\mathcal{U}}_{\text{CNOT}} \neq \Lambda_P(\eta_{\text{sys}}) \circ \mathcal{U}_{\text{CNOT}}$, where η_{sys} refers to the bias of the system. While there are alternatives to implementing a CNOT gate, no bias-preserving CNOT gate can be constructed for two-level qubits [25].

However, the fact that some phase flips may be transformed into bit flips does not imply that the noise is fully depolarized, i.e., that all possible Pauli errors are equiprobable. To elucidate this, we analyze the underlying structure that the noise presents after the application of the noisy CNOT gate implemented by means of the interaction in Eq. (1).

We present the remainder or residual noise bias toward dephasing errors in Fig. 1 (for details on the numerical simulations, see Appendix B). The results show that the actual bias of the noise toward dephasing for the CNOT gate is reduced very significantly, rapidly saturating at $\eta_{\text{CNOT}} \approx 5$ for $\eta_{\text{sys}} > 1000$. While the bias is reduced by orders of magnitude, it is noteworthy to recall that, e.g., the $XZZX$ code has shown almost a 25% increase of the threshold in the code-capacity model for a $\eta = 5$ bias [4]. We will further discuss how this moderate bias could be very beneficial later on.

Some other code proposals tailored to deal with biased noise involve Y stabilizers, which are susceptible to both bit-flip and phase-flip errors [34,38]. In order to measure such stabilizers, CY gates are employed to entangle the data and check qubits. These gates can be implemented with an interaction similar to Eq. (1), but with a Y operator on the target qubit instead of the X operator. Note that for these gates, the predominant dephasing noise in the system will

also not commute with such an operator of the target qubit, indicating that the bias is also similarly lowered.

B. Hadamard gates

Hadamard gates also play a pivotal role in syndrome-extraction circuits and therefore it becomes crucial to elucidate what happens with the bias when applying these single-qubit gates. A Hadamard gate is a π rotation over the $(X + Z)/\sqrt{2}$ axis in the Bloch sphere. Hence, it can already be seen that, since the interaction does not commute with the predominant dephasing noise, it will not result in a bias-preserving gate.

Furthermore, we have conducted a similar analysis as for the CNOT and CY gates in order to quantify the residual noise after the execution of a Hadamard gate (for details on the numerical simulations, see Appendix B). For doing so, we perform the rotation by means of the interaction

$$H_{\text{Had}} = V \left[\frac{X + Z}{\sqrt{2}} \right], \quad (2)$$

which results in the desired gate at time $VT = \pi/2$ up to a global phase. The rest of the setting is similar, i.e., the resulting gate fidelity is around 99.7%, with the slight difference that there are three dissipators, $\{X, Y, Z\}$. We have observed that for this case, the noise bias is completely depolarized, i.e., the structure of the resulting noise channel is almost symmetric independent of the system noise bias [39]. Thus, we conclude that the Hadamard gates will introduce depolarizing errors, at least if implemented by the interaction in Eq. (2).

C. CZ gates

CZ gates are also ubiquitous in syndrome-extraction circuits, with the aim of measuring Z stabilizers susceptible to bit-flip errors. The essential difference between this entangling gate and the previously discussed ones is that it is a rotation around the ZZ basis. Hence, this gate commutes with the predominant error channel, i.e., pure dephasing errors, and is a bias-preserving gate in the sense of Definition A1 (see Appendix A). Similar to the CY gate, the CZ gate can be realized by means of substituting the X term in the control qubit of the Hamiltonian in Eq. (1) by a Z term. Furthermore, as explained in Ref. [27], it can also be achieved by means of an interaction described by the Hamiltonian

$$H_{ZZ} = -V[Z \otimes Z] \quad (3)$$

and some local Pauli- Z rotations, i.e., also commuting with the predominant noise [40].

To sum up, the main point is that CZ gates can be implemented in a bias-preserving manner, as has been pointed out before in Refs. [7,26,27]. As we will discuss later,

TABLE I. Certain native interactions to enable CZ gates for several qubit technologies. For descriptions of these gates, see Appendix C.

Platform	Interaction	Bias-preserving
Silicon spin qubits	Exchange (strong limit)	No
	Exchange (weak limit)	Yes
Trapped ions	Mølmer–Sørensen (XX)	No
	Liebfried–Sørensen	Yes
Neutral atoms	Rydberg blockade	Yes
N- V centers	Magnetic dipolar interaction	Yes
	Hyperfine interaction	Yes
Superconducting	Cross-resonance gate	No
	STC coupler	Yes

this fact is critical in order to leverage biased noise when working with two-level qubits.

D. Native implementations of the CZ gate

In the previous section, we have seen that a CZ gate can be enabled in a bias-preserving manner for two-level qubits. However, this does not imply that using this gate for syndrome-extraction circuits will result unavoidably in biased noise. We emphasize here the term *native gate*, which refers to the implementation of a specific quantum operation in a physical hardware without requiring further decomposition into other gates. The interaction between qubits, as well as interactions between them and an external electromagnetic field, depends on the underlying hardware architecture. Achieving a desired quantum operation requires precise control over these interactions. In Table I, we summarize ways of enabling CZ gates [41] in several technologies exhibiting biased noise as well as if those would be bias preserving. For a description of how those native CZ gates are implemented, see Appendix C.

III. QUANTUM ERROR CORRECTION FOR TWO-LEVEL QUBITS EXHIBITING BIASED NOISE

A pertinent question is whether the noise bias can still be leveraged for more efficient QEC even if some of the gates do not preserve the bias. In this section, we answer this question by studying the $XZZX$ surface code [4] (for a description of the rotated $XZZX$ code, see Appendix D). Specifically, we first study a *hybrid biased-depolarizing circuit-level noise model* as a generic tool for studying QEC tailored to biased noise when dealing with two-level qubits and study the performance of the rotated $XZZX$ code over such a model (for our noise-model proposal, see Appendix E). We then proceed to study a more tailored hybrid circuit-level noise model considering our previous discussion on how some residual bias still exists when implementing CNOT gates.




FIG. 2. The estimated threshold of the rotated $XZZX$ code as a function of the system bias for the HBD circuit-level noise model with and without bias-preserving CZ gates. (a) Bias-preserving CZ gates. The threshold has improved from around 0.66% for the SD noise model up to 0.93% when $\eta \geq 100$. This is a 40% improvement. Furthermore, a saturation of the threshold can be observed. (b) CZ gates that are not bias preserving, modeled here as introducing two-qubit depolarizing errors. There is no significant improvement for this case. In fact, thresholds saturate early on, around the value obtained for $\eta = 1$ with bias-preserving CZ gates.

A. The hybrid biased-depolarizing circuit-level noise model

In Fig. 2(a), we show the estimated thresholds for the $XZZX$ variant of the rotated surface code when operating over the proposed HBD circuit-level noise model. The threshold values are summarized in Table II. Note that here we refer to the SD circuit-level noise model with an $\eta = 1/2$. Note also that the thresholds for the horizontal (H) memory are slightly higher than those for the vertical (V) memory (for a description of these memories, see Appendix G). However, we have observed that their subthreshold, i.e., at physical error rates substantially lower than the threshold, the logical error rates are very similar. Specifically, even if the threshold is higher, the logical error rates of the H memory are slightly worse (almost insignificantly) than those for the V memory for the deep-subthreshold physical error rates. The cause for this effect is probably that for the H memory there are $(d+1)/2$ logical strings consisting of $(d+1)/2$ Z operators and $(d-1)/2$ X operators, while for the V memory there are $(d-1)/2$ of these, i.e., there is one more for the H memory. Since the considered distances are odd and the noise is predominantly biased, this makes the logical error probability slightly higher for the H memory.

TABLE II. Improvement of the threshold as a function of the bias for the HBD circuit-level noise model.

η	$p_{\text{threshold}}$ (%)	Improvement (%)
1/2 (SD)	0.66	...
1	0.69	~ 4.5
10	0.85	~ 29
100	0.92	~ 40
1000	0.93	~ 40
10 000	0.93	~ 40

Furthermore, Fig. 2(b) shows the estimated thresholds for the $XZZX$ code if the CZ gates used for syndrome extraction are not executed in a bias-preserving manner, i.e., they introduce depolarizing noise. For such a case, it can be seen that the threshold barely increases and rapidly saturates around 0.7%. This is very close to the threshold of the $\eta = 1$ case, in which the CZ gates are done so that they preserve bias [i.e., in Fig. 2(a)], indicating that there is no real performance gain for this case. Hence, it is very important to enable the CZ interactions in such a way (see Sec. II) as to leverage the bias present in certain two-level qubits. Otherwise, no performance gain will be obtained.

B. A more tailored hybrid circuit-level noise model with residually biased CNOT gates

If we recall Sec. II, CNOT gates do show some small residual bias in the structure of the noise. In this section, we will study the importance of such residual bias.

The noise model considered in this section is the HBD circuit-level noise model, with the difference that the CNOT gates present some residual biased noise. Specifically, the model considers the tuple $(\eta_{\text{sys}}, \eta_{\text{CNOT}})$ with the values obtained in Sec. II, in which the CNOT gate is enabled via the interaction in Eq. (1). For clarity, we will use the value $\eta = \eta_{\text{sys}}$ for the plots and discussion, but it should be clear that this does not imply that the CNOT gates preserve this bias, i.e., for each η , the CNOT will have its own residual bias. Finally, we will consider that the CZ gates are constructed by means of native interactions that preserve the system bias.

1. Threshold

In Fig. 3, we present the threshold improvement as a result of the system bias for the studied code. As can be




FIG. 3. The estimated threshold of the rotated $XZZX$ code as a function of the system bias considering residual CNOT bias according to Fig. 1. We use the residual bias values as studied in Sec. II. The threshold has improved from around 0.66% for the SD noise model up to 1.27%. This is a 93% improvement. Furthermore, a saturation of the threshold can still be observed for high bias values.

observed, the actual threshold improvement is significantly higher than for the HBD model without considering CNOT gates with a residual bias. In fact, the threshold improves up to approximately 1.27% for system biases exceeding 1000. It is noteworthy to comment that this is a 91% improvement with respect to the standard depolarizing circuit-level noise case.

Thus, the fact that the CNOT gates present residual bias is very important for leveraging the bias toward dephasing, even for two-level qubits. Furthermore, the obtained threshold improvements are significant considering moderate residual values saturating at $\eta_{\text{CNOT}} = 5$, as presented in Sec. II. To understand this, it is important to recall that this is only possible due to the

TABLE III. Improvement of the threshold as a function of the bias for the HBD circuit-level noise model and CNOT gates exhibiting residual bias as studied in Sec. II.

η	$p_{\text{threshold}}$ (%)	Improvement (%)
1/2 (SD)	0.66	...
1	0.71	~ 7.5
10	0.96	~ 45.5
100	1.2	~ 81.9
1000	1.26	~ 91
10 000	1.27	~ 91

presence of bias-preserving CZ gates, indicating that any improvement in the residual bias on the CNOT results in significantly increasing the overall bias in the whole noisy extraction circuits. As seen for code-capacity models in Ref. [4], even a slight increase in the bias has resulted in huge threshold improvements and, here, we are observing the same effect for the residual bias in the noisy CNOT gates used for extraction.

We quantify the overall threshold improvements for different system bias values in Table III. Note how the gain does also saturate for this case, as happened for the HBD model. However, it saturates for higher biases. This makes sense since the residual bias for the CNOT gates also presents such a similar saturation tendency, as seen in Fig. 1.

2. Footprints

We now discuss the required qubit footprints to reach a certain regime of error-free quantum operations as a function of the bias (for a description of QEC footprints, see Appendix F).

We have numerically computed the logical error rates per round, p_L , as a function of the code distance for $d \in$




FIG. 4. The logical error rates per round versus the rotated $XZZX$ surface-code distances with bias-preserving CZ and CNOT gates with the residual biases discussed in Fig. 1. The logical error rates in these figures are for a horizontal (H) memory experiment. Both memories behave similarly deep below threshold, with the H memory showing a slightly worse logical error probability, p_L , and thus it is the limiting case. (a) Physical error rate $p = 0.003$. (b) Physical error rate $p = 0.001$.

TABLE IV. The qubit footprints and relative decrease in percentage with respect to the symmetric SD circuit-level noise model for each bias value, considered at $p = 0.003$. The footprints are computed by selecting the first odd distance value below the target logical error probability in Fig. 4(a) and using $d^2 + (d^2 - 1)$ as the number of qubits (data and check qubits) required for such a distance.

η	Megaquop footprint	Decrease (%)	Gigaquop footprint	Decrease (%)	Teraquop footprint	Decrease (%)
1/2 (SD)	1681	...	4417	...	8977	...
10	721	~ 57	1921	~ 57	3361	~ 63
100	449	~ 74	1249	~ 72	2449	~ 73
1000	449	~ 74	1249	~ 72	2177	~ 76
10 000	449	~ 74	1249	~ 72	2177	~ 76

{5, 7, 9, 11, 13, 15} and used the values to project further. Simulations have been done for the same circuit-level noise models as in Sec. III C 1 (considering residual bias in CNOT gates). We have done so for two experimentally relevant physical error rates: $p = 0.003$, representing state-of-the-art CZ gate error rates reported for the Google Quantum AI Willow processor [42]; and $p = 0.001$, representing the usual physical error rate for practical QEC, i.e., the 99.9% gate fidelity [43]. The results are shown in Fig. 4 (details on the numerical methods used can be found in Appendix G) and convincingly illustrate that the required distances to reach certain logical error probabilities are significantly reduced when a biased system is considered. Furthermore, it can be observed that the reduction in logical error rate provided by the bias shows a similar saturation effect as for the code threshold for $\eta \geq 100$. Table IV quantifies the qubit footprints required to reach the Mega-, Giga-, and Teraquop regimes at a physical error rate of $p = 0.003$. The reduction in qubit footprint ranges from around 60% for a moderate bias of $\eta = 10$ to around 75% for biases exceeding the value of a 100 compared to an SD circuit-level noise model. Furthermore, we quantify the same data for the $p = 0.001$ physical error rate case in Table V. Inspecting the table, we can observe that for this case, the decrease in qubit footprint ranges from 33% to around 50%.

The main conclusion that can be drawn from these results is that noise bias toward dephasing errors can still be leveraged for two-level qubits for which bias-preserving CNOT gates cannot be constructed, resulting in a significant reduction in the required number of qubits

to achieve certain fault-tolerant quantum computational regimes of interest. It is noteworthy to state that the footprint reductions discussed here come “for free” for those qubit platforms if the CZ gates are enabled natively in a bias-preserving manner.

C. Compiling the XZZX code with CZ gates

In many qubit platforms, CNOT gates are not usually enabled in a direct manner by means of pulses such as the one described in Sec. II. The CNOT gates are compiled by means of a CZ gate sandwiched by Hadamard gates on the target qubit. This is, e.g., the case for silicon spin qubits or neutral atoms, among others [28,44]. Thus, it is relevant to compile the extraction circuit by means of those gates. Here, we numerically study how this compilation does allow leveraging of the system bias. For our numerical evaluation, we consider the HBD circuit-level noise model proposed in this work.

1. Threshold

In Fig. 5, we show the threshold of the XZZX surface code as a function of the system bias for the extraction-circuit compilation using exclusively CZ gates as entangling gates. The threshold here improves from a physical error rate around 0.53% up to 0.8% for system biases above $\eta = 100$. This is a 50% improvement in threshold. Note that, here, the threshold values both for the standard depolarizing and biased cases are lower than for the previous sections. This occurs due to the extra error locations introduced by the layers of Hadamard gates on the

TABLE V. The qubit footprints and relative decrease in percentage with respect to the symmetric SD circuit-level noise model for each bias value, considered at $p = 0.001$. The footprints are computed by selecting the first odd distance value below the target logical error probability in Fig. 4(b) and using $d^2 + (d^2 - 1)$ as the number of qubits (data and check qubits) required for such a distance.

η	Megaquop footprint	Decrease (%)	Gigaquop footprint	Decrease (%)	Teraquop footprint	Decrease (%)
1/2 (SD)	241	...	721	...	1249	...
10	161	~ 34	449	~ 38	881	~ 30
100	161	~ 34	449	~ 38	721	~ 43
1000	161	~ 34	337	~ 54	721	~ 43
10 000	161	~ 34	337	~ 54	721	~ 43




FIG. 5. The estimated threshold of the rotated $XZZX$ code over the hybrid biased-depolarizing (HBD) model with a syndrome-extraction circuit with only CZ gates as entangling gates. The threshold has improved from around 0.53% for the standard depolarizing (SD) noise model up to 0.8%. This is a 51% improvement. Furthermore, a saturation of the threshold can still be observed for high bias values.

data qubits required for X -stabilizer extractions. Since the hybrid bias depolarizing (HBD) model that we consider assumes all error locations to fail with the same error rate, the errors introduced by said Hadamard gates have a considerable impact on the code performance. Note, however, that in realistic hardware, single-qubit gates present error rates around an order of magnitude lower than entangling gates. Thus, the performance will be better when considering such nonidentically distributed noise scenarios [17,42,43,45,46].

2. Footprints

In Tables VII and VIII, we present the qubit footprints required to reach the three quantum computational regimes

TABLE VI. Improvement of the threshold as a function of the bias for the HBD circuit-level noise model and CZ compilation.

η	$p_{\text{threshold}}$ (%)	Improvement (%)
1/2 (SD)	0.53	...
10	0.73	~ 37.7
100	0.79	~ 49
1000	0.8	~ 51
10 000	0.8	~ 51

discussed before, for physical error rates of $p = 0.003$ and $p = 0.001$, respectively. The logical error rates per round as a function of the code distance obtained by numerical simulations are presented in Fig. 6. Similar to Sec. III B, we observe that the required footprints reduce in a significant manner, ranging from 34% up to 79%. As discussed before, these reductions in qubit footprints come with no other cost than constructing CZ gates that preserve the bias of the system. Importantly, note that the footprints obtained in this case using only CZ entangling gates result in higher qubit numbers than in Sec. III B. This arises from the extra error locations introduced by the layers of Hadamard gates on the data qubits required for X -stabilizer extractions, as explained for the code thresholds.

IV. DISCUSSION AND CONCLUSIONS

In this paper, we have discussed how to leverage the bias toward dephasing noise that certain two-level qubit platforms exhibit. Our results indicate that by leveraging the system bias toward dephasing, one can reach practical surface-code-based QEC with a significantly lower amount of resources than for systems the noise of which is symmetric. First, the obtained thresholds for the rotated $XZZX$ surface code with $\eta \geq 100$ stand above 1.2%. This implies that technologies exhibiting such amounts of bias, such as,




FIG. 6. The logical error rates per round versus the rotated $XZZX$ surface-code distances with the extraction circuit using only bias-preserving CZ gates as entangling gates discussed in Sec. II. The logical error rates in these figures are for a horizontal (H) memory experiment. Both memories behave similarly deep below threshold, with the H memory showing a slightly worse logical error probability, p_L , and thus it is the limiting case. (a) Physical error rate $p = 0.003$. (b) Physical error rate $p = 0.001$.

TABLE VII. The qubit footprints and relative decrease in percentage with respect to the SD circuit-level noise model for each bias value, considered at $p = 0.003$ with the compilation only using CZ gates. The footprints are computed by selecting the first odd distance value below the target logical error probability in Fig. 6(a) and using $d^2 + (d^2 - 1)$ as the number of qubits (data and check qubits) required for such a distance, where d is the code distance.

η	Megaquop footprint	Decrease (%)	Gigaquop footprint	Decrease (%)	Teraquop footprint	Decrease (%)
1/2 (SD)	3697	...	10 657	...	21 217	...
10	881	~ 76	3041	~ 72	6049	~ 72
100	881	~ 76	2449	~ 77	4417	~ 79
1000	881	~ 76	2449	~ 77	4417	~ 79
10 000	881	~ 76	2449	~ 77	4417	~ 79

e.g., silicon spin qubits or ion traps, could operate well below threshold right now [47]. Importantly, it is common for the CNOT gates to be compiled by means of a CZ gate sandwiched by Hadamard gates on the target qubit in many technologies such as silicon spin qubits or neutral atoms, among others [28,44]. We have presented the results of the $XZZX$ code with extraction circuits following such a compilation, reaching similar conclusions.

The methods presented provide a clear improvement to qubit technologies exhibiting bias toward dephasing. The next milestone toward the goal of fault-tolerant quantum computing that would actually make the field go beyond the NISQ era stands in the construction of the Megaquop machine [33]. Currently, many technologies are competing to reach such a milestone as fast as possible and, in an era in which each physical qubit counts, reducing the required footprint by up to 75% may give the edge to technologies chasing the predominant superconducting qubits. Also, this is especially relevant since, as we have studied in this paper, qubit savings are much more significant for the physical error rates around 0.3% that are common at this point in time.

Furthermore, it is also important to comment not only on the space (footprint) savings but also on the time savings that are obtained. In this sense, we will discuss the implications that the distance reductions imply for the logical clock rate in Litinski's surface-code lattice-surgery-based framework for the construction of a fault-tolerant quantum computer [48]. In such a setting, the time cost to

perform fault-tolerant operations essentially defines the logical clock rate of a fault-tolerant quantum computer and it is related to the number of syndrome-extraction rounds [48]. Thus, our results indicate that a reduction in the required code distance also results in a higher logical clock speed. Note that while fully understanding this would also require stability experiments, those are much more forgiving than memory experiments, as pointed out by Gidney in Ref. [49]. The overall space-time cost, in terms of (physical qubits) $\times$ (code cycles), of a computation depends on how the logical operations are enabled, but the leading order has a cubic dependence on the code distance [48]. Therefore, we can determine the space-time cost savings between equivalent protocols (i.e., the logical operations are done similarly, indicating that the constants of the cost are the same) whenever biased noise is leveraged. We summarize these space-time cost savings in Tables IX and X. The overall space-time cost reductions are relevant and undoubtedly convenient to make fault-tolerant quantum machines in a significantly more efficient manner.

One of the key points of the obtained results is that the improvement in the QEC performance is obtained with no additional overhead. We have seen that the key element to leverage bias toward dephasing for two-level qubit technologies for which a bias-preserving CNOT cannot be constructed is to obtain tailored codes that use CZ gates to extract the syndrome and to natively enable those entangling gates to preserve the bias. The implementation of native bias-preserving CZ gates is rather straightforward

TABLE VIII. The qubit footprints and relative decrease in percentage with respect to the symmetric SD circuit-level noise model for each bias value, considered at $p = 0.001$ with the compilation only using CZ gates. The footprints are computed by selecting the first odd distance value below the target logical error probability in Fig. 6(b) and using $d^2 + (d^2 - 1)$ as the number of qubits (data and check qubits) required for such a distance.

η	Megaquop footprint	Decrease (%)	Gigaquop footprint	Decrease (%)	Teraquop footprint	Decrease (%)
1/2 (SD)	337	...	881	...	1681	...
10	241	~ 29	577	~ 35	1057	~ 37
100	161	~ 52	449	~ 49	881	~ 48
1000	161	~ 52	449	~ 49	881	~ 48
10 000	161	~ 52	449	~ 49	881	~ 48

TABLE IX. The overall space-time cost decrease as a result of bias tailoring at $p = 0.003$. The decreases are computed by selecting the first odd distance value below the target logical error probability in Fig. 4(a) and using the leading-order space-time cost d^3 [48] to compute the decrease when compared to the standard depolarizing case.

η	Megaquop (%)	Gigaquop (%)	Teraquop (%)
10	72	72	77
100	86	85	86
1000	86	85	88
10 000	86	85	88

for many technologies, as discussed in Sec. II. Additionally, we have observed how the remainder bias present in the CNOT gates required for syndrome extraction boosts the performance of the tailored code even further. It is noteworthy to state that our analysis considers a generic Hamiltonian and error model to enable the CNOT gate. However, the residual bias of CNOT gates with other pulses and error models may be larger and, thus, the tailored code is expected to perform even better. Hence, even if a bias-preserving CNOT gate cannot be executed for two-level qubits, trying to maximize the residual bias is an important future line of research, as it will make the QEC gains even bigger. In this line, optimal pulse control has proven to be a powerful technique to reduce the noise of certain tasks in NISQ devices [50] and, thus, such pulse engineering can potentially be used to increase the residual bias presented by CNOT gates. Additionally, experimentalists and engineers can aim to develop entangling gate protocols with the same aim.

Furthermore, note that the considered HBD circuit-level noise models, both with depolarizing and with residual bias CNOT gates, consider that each element of the circuit fails at the same rate. In reality, the errors will be nonidentically distributed, i.e., each element will present a different error rate [17,42,43,45,46]. This means that some operations will be more fault prone than others. For example, it is very typical that single-qubit gates, e.g., Hadamard gates, are an order of magnitude less prone to experiencing errors than

entangling gates such as CNOT or CZ gates [43,45,46]. This aspect is important for our discussion since the Hadamard gates used in the syndrome-extraction circuits are the only gates exhibiting no residual bias, i.e., they introduce depolarizing noise. As for real systems, these elements will introduce errors with orders of magnitude less frequency than the entangling gates, this will make overall noise in the system more predominantly biased. Therefore, we expect certain performance increases whenever such more fine-grained noise models with bias are considered. This becomes especially relevant for the compilation just using CZ gates that involves two extra layers of Hadamard gates per round on the data qubits (see Sec. III C).

It is also important to discuss the way in which we have modeled state preparation and measurement (SPAM) errors. Note how in our HBD models, state preparation errors imply flipping to the orthogonal state and measurement errors flipping to the contrary measurement, each of them with probability p . This contrasts with some other circuit-level noise models including bias that are present in the literature [26,34,35]. Those models consider that the probability of preparing or measuring a state in the Z basis is η times smaller than for doing so in the X basis. This implies that those operations on the Z basis would fail orders of magnitude less frequently than the other elements in the system. This contrasts with the way in which bias is included for the gates in our model, i.e., the fidelity remains the same, but the frequency in which each of the Pauli errors occurs changes. Moreover, the physics behind SPAM operations is different to that regarding quantum gates, indicating that adding such bias is not as clear as for the other elements that we have discussed in this paper [11]. Studying how to realistically model bias in SPAM for different qubit technologies is considered as a future work.

Another important consequence of the distance reduction achieved by the tailoring to leverage bias toward dephasing relates to the decoder. Specifically, decoders must be fast enough in order to be able to implement non-Clifford operations in a fault-tolerant manner [2]. We have used the SPARSE BLOSSOM implementation of the minimum-weight perfect-matching (MWPM) decoder to aim to correct the errors [51]. This implementation can decode distance-17 surface-code syndromes in less than a microsecond per round at the target physical error rate of 0.1%, which is enough to avoid the backlog problem [51]. As can be seen in Table V, the required footprint for all relevant regimes when the bias exceeds 100 does not involve a distance beyond 17, indicating that the decoder used should be able to operate below the microsecond per round. The microsecond per round target is not necessary for all qubit technologies, since neutral-atom or ion-trap systems require much more relaxed latencies [52]. However, technologies with fast gates, such as spin qubits, will have stringent decoding latencies of around 1 μ s. Moreover, more precise decoding algorithms could also

TABLE X. The overall space-time cost decrease as a result of bias tailoring at $p = 0.001$. The decreases are computed by selecting the first odd distance value below the target logical error probability in Fig. 4(b) and using the leading-order space-time cost d^3 [48] to compute the decrease when compared to the standard depolarizing case.

η	Megaquop (%)	Gigaquop (%)	Teraquop (%)
10	45	51	41
100	45	51	56
1000	45	68	56
10 000	45	68	56

be used in order to reduce the footprints even more. For example, the use of belief propagation as a predecoder has been shown to be beneficial for biased-noise models [34], so using belief matching or the almost-linear-time belief-propagation plus ordered-Tanner-forest (BP+OTF) decoder may improve performance and reductions even further [34,53].

Regarding code construction, the present work also provides insights and future paths to explore. First, the present paper has been centered around the rotated $XZZX$ variant of the surface code [4,46], mainly due to it being one of the strongest proposals to leverage bias and its planar nature, making it easier to implement experimentally in many platforms. However, there exist many bias-tailored code proposals in the literature considering code-capacity noise [4,5,21–24,34,38]. The impossibility of constructing a bias-preserving CNOT gate in two-level qubits has made the community think that bias cannot be leveraged for such platforms that naturally exhibit such a feature. We have refuted such a statement in the present paper, noting that if the code involves CZ gates, constructed in a bias-preserving manner, substantial improvement in error-correction performance can be obtained. This indicates that bias-tailored codes based on using CY gates, such as the XY or the $ZZZY$ surface codes, may not be good candidates to exploit bias in two-level qubits [34,38]. The main reason for that is that these proposals rely on using Y stabilizers instead of Z ones and, as we have discussed in Sec. II, the CY gates used to perform said measurements cannot be enabled in a bias-preserving manner, similar to CNOT gates for X -stabilizer measurements. Hence, the syndrome-extraction circuits for those codes will mainly consist of gates not preserving bias and, similar to what happens for CZ gates constructed in a non-bias-preserving way, will not be able to leverage bias at the circuit level. Anyway, there are many other code-tailoring proposals involving X and Z stabilizers, such as the domain-wall color codes [24], $XZZX$ toric codes, or even general bias-tailored qLDPC codes [5]. Furthermore, qLDPC codes have shown the ability to reduce the required qubit footprints by up to a factor of 10 for circuit-level noise models [54], so tailoring those codes and circuits with our findings may result in further qubit savings for qubit technologies with noise biased toward dephasing, making those closer to fault tolerance. Note also that this reduction in footprint and required code distance would also simplify the decoding problem for such codes.

Furthermore, we can envision that our study can be relevant to make certain quantum error mitigation (QEM) protocols more efficient. In this context, a recent work addresses exploiting the extremely biased nature of cat qubits for realizing probabilistic error cancellation (PEC) in a much more efficient manner [55]. The authors observe that the sampling cost of the aforementioned QEM technique can be exponentially suppressed, making the protocol much more amenable for practical applications in

cat-qubit-based NISQ devices. Therefore, the QEM community can make use of our QEC framework to aim for more efficient PEC and other QEM protocols when biased two-level qubit NISQ processors are considered.

To sum up, in this paper we have discussed how to leverage bias for QEC at the circuit level even when bias-preserving CNOT gates cannot be constructed. In a time in which qubits are expensive resources, reducing the required physical qubit per logical qubit rates is critical in order to make the first fault-tolerant realizations of quantum algorithms a reality. Here, we recover the case that noise bias toward dephasing can give technologies that present such characteristics an edge in such a quest. Importantly, this physical-qubit footprint reduction comes for free. We hope that bias tailoring for two-level qubits will flourish again as an efficient path toward fault-tolerance, as it has done over the past five years, during which time QEC theorists have mostly considered code-capacity models. We hope that by following this path, Megaquop machines consisting of 100 logical qubits can be constructed with physical-qubit counts in the order of thousands rather than several tens of thousands, making such first milestone and further ones closer to current times.

ACKNOWLEDGMENTS

We would like to thank Antonio de Marti i Oliu for many useful discussions and recommendations. Furthermore, we also thank other members of the Cavendish Laboratory, the Hitachi-Cambridge Laboratory, and the Quantum Information Group at Tecnum for their support and many useful discussions. This work was supported by the Spanish Ministry of Economy and Competitiveness through the MADDIE project (Grant No. PID2022-137099NBC44), by the Spanish Ministry of Science and Innovation through the project “Few-qubit quantum hardware, algorithms and codes, on photonic and solid-state systems” (PLEC2021-008251); by the Diputación Foral de Gipuzkoa through the “Biased quantum error mitigation and applications of quantum computing to gene regulatory networks” project (2024-QUAN-000020); and by the Ministry for Digital Transformation and of Civil Service of the Spanish Government through the QUANTUM ENIA project call—QUANTUM SPAIN project, and by the European Union through the Recovery, Transformation and Resilience Plan—NextGenerationEU within the framework of the Digital Spain 2026 Agenda. J.E.M. is funded by the Spanish Ministry of Science, Innovation and Universities through a José Castillejo mobility grant for his stay at the Cavendish Laboratory of the University of Cambridge.

DATA AVAILABILITY

The data that support the findings of this paper are openly available [56].

APPENDIX A: BIAS-PRESERVING GATES

For code-capacity noise models, the elements in the syndrome-extraction circuits, i.e., gates and SPAM; are considered to be ideal and noiseless [2]. However, each of these is actually imperfect and, in reality, introduces errors. This results in the so-called *circuit-level noise* error model, considered to be realistic enough to represent what happens in actual hardware. As discussed before, stabilizer measurements are split into sequences of two-qubit gates that cannot be done simultaneously and, thus, an extraction schedule should be defined [57]. Many extraction schedules are valid, but generally a good schedule should aim to parallelize as many stabilizer operations as possible while trying to maximize the *circuit-level distance* [54]. The circuit-level distance refers to the minimum number of circuit errors that result in an undetectable logical error. This value is upper bounded by the code distance and depends on the actual extraction schedule proposed.

Biased noise has shown the possibility of considerably improving the performance of QECCs when those are tailored to operate over models with an asymmetric noise structure [4,5,21–24]. This, however, has been mainly done from the perspective of code-capacity noise models. The extension of this noise asymmetry to the circuit-level noise picture is far from trivial, since the noisy-circuit operations should in principle maintain the noise bias, $\eta = p_z/(p_x + p_y)$, for leveraging such bias tailoring.

Definition A1 (Bias-preserving gate). A bias-preserving gate is a noisy gate, $\tilde{\mathcal{U}}$, that can be modeled as the perfect gate, $\mathcal{U}$, followed by a noise channel maintaining the noise bias, $\Lambda_P(\eta)$, present in the system, i.e.,

$$\tilde{\mathcal{U}} = \Lambda_P(\eta) \circ \mathcal{U}. \quad (\text{A1})$$

Note that the subscript P in the noise channel refers to the fact that we are considering Pauli noise [58]. It is important to comment that there are other conceptions of bias-preserving gates, related to how Pauli operators propagate through those gates [55,59]. While such a conception might be of interest for studies related to quantum error mitigation, in the context of QEC, a bias-preserving gate refers to the one in Definition A1. Importantly, it is known that not all gates can be constructed in a bias-preserving manner, i.e., $\tilde{\mathcal{U}} \neq \Lambda_P(\eta) \circ \mathcal{U}$.

APPENDIX B: NUMERICAL SIMULATIONS OF ERROR CHANNELS

In Sec. II, we have presented a numerical study of the error channels resulting from the application of certain gates in order to determine the remainder bias after applying such gates. For doing so, we have performed master-equation simulations of the gates in question using the QuTiP software package [60]. Specifically, we numerically

simulate the gates in question defined by a Hamiltonian, H , under Markovian uncorrelated (white) noise. This results in solving the following Lindblad equation for the time at which the gate of interest is realized:

$$\dot{\rho} = -\frac{i}{\hbar}[H, \rho] + \sum_i \lambda_i \mathcal{L}_i\{\rho\}, \quad (\text{B1})$$

where ρ is the density matrix; the dissipators $\mathcal{L}_i$ are given by the Pauli matrices, $\{X, Y, Z\}$ for single-qubit gates and $\{I, X, Y, Z\}^{\otimes 2}/I^{\otimes 2}$ for two-qubit gates; and the λ_i are their coupling rates. The error channels for each gate are then determined by means of the so-called Pauli transfer matrix, defined as [27]

$$R_{ij} = \frac{1}{2^n} \text{Tr}[P_i \Lambda(P_j)], \quad (\text{B2})$$

where n is the number of qubits, Λ represents the noise map, and $P_i, P_j \in \{I, X, Y, Z\}^{\otimes n}$. In this way, the Pauli transfer matrix describes the fidelity that each of the elements of the density matrix shows after experiencing the noise channel. Thus, in order to extract the noise Pauli transfer matrix, we simulate Eq. (B1) to obtain the $R_{\text{noisy}}^{\text{gate}}$ transfer matrix and its associated ideal evolution, $R_{\text{ideal}}^{\text{gate}}$, by means of the von Neumann equation. The noise transfer matrix is then obtained as $R_{\text{noise}} = R_{\text{noisy}}^{\text{gate}} (R_{\text{ideal}}^{\text{gate}})^{-1}$, i.e., obtaining just the part referring to the noise in the evolution [27]. Finally, since the diagonal elements of the Pauli transfer matrix are related to fidelities, we use a symplectic Walsh-Hadamard transform to convert those into the error probabilities in an operator-sum representation [61,62].

For the simulations using the Hamiltonian in Eq. (1), we set $V = 1$ and the coupling rates of the dissipators of the Lindblad equation so that the combined rates of the $\{ZI, IZ, ZZ\}$ elements are η_{sys} times higher than those of the rest of the nontrivial Pauli matrices. Furthermore, we set $\lambda_{ZI} + \lambda_{IZ} + \lambda_{ZZ} = 0.002$, which results in an overall gate fidelity of around 99.7%, i.e., an experimentally reasonable value.

It is important to point out that the numerical study assumes a continuous noisy time evolution with a static Hamiltonian under the Born-Markov and uncorrelated (white-)noise approximation with a specific structure. In this sense, each qubit technology presents specific characteristics both in terms of gate implementation and noise behavior. For example, it is typical that gates are implemented as combinations of several types of pulses, i.e., in a time-dependent manner. Furthermore, $1/f$ correlated noise is important for spin and superconducting qubits. Consequently, even if the implemented CNOT gates will have a qualitatively similar behavior, the actual residual bias may change. Hence, it is important for pulse engineers and experimentalists focusing on specific implementations to understand all of those to obtain conclusions based on

our bias-tailoring framework. In this sense, our analysis here is intended to show that the impossibility of constructing CNOT gates that fully preserve bias in two-level qubits does not imply that those gates cannot show a level of bias that cannot be exploited by means of code tailoring.

APPENDIX C: NATIVE CZ GATES FOR SEVERAL QUBIT TECHNOLOGIES

In this appendix, we review several native implementations of CZ gates in different physical platforms and analyze whether they preserve noise bias. First, it is known that a CZ gate can be implemented conjugating the control of a CNOT gates with Hadamard gates. Nevertheless, this realization of a CZ gate will not preserve the noise bias, since CNOT and Hadamard gates do not commute with the dephasing noise. Hence, CZ gate compilations using native CNOT gates enabled by means of ZX interactions will not preserve bias, as explained in the main text.

Spin qubits are encoded in the spin state of an electron bounded to a quantum dot (QD) generated on a semiconductor and present a high bias toward Z errors [7,11–14]. In these devices, the spin state (s_i) of the electron interacts with an external magnetic field $\mathbf{B}(t)$ through a Zeeman Hamiltonian $H_{\text{Zeeman}} = \mathbf{B}(t) \cdot \mathbf{s}_i$ and with other electron spin qubits s_j through the Heisenberg exchange interaction $H_{\text{exc}} = J_{\varepsilon}(t) \mathbf{s}_i \cdot \mathbf{s}_j$, the coupling strength of which depends on the voltage detuning $\varepsilon(t)$ [63]. The control is executed by means of electron spin resonance (ESR) using oscillating magnetic fields [64] or by means of an electron dipole spin resonator (EDSR) using micro-magnets [65,66]. The relationship between the coupling strength, $J_{\varepsilon}(t)$, and the Zeeman gradient, $\delta B_z = B_{z_1} - B_{z_2}$, defines two regimes in which CZ gates can be enabled. In the strong-exchange regime ($J_{\varepsilon}(t) \gg \delta B_z$), the exchange interaction generates the mixed-singlet and neutral-triplet states, i.e., it allows the transition $|\uparrow\downarrow\rangle$ to $|\downarrow\uparrow\rangle$, which does not commute with the dephasing noise. CZ gates are usually compiled as sequences of native $\sqrt{\text{SWAP}}$ gates [44], which, as explained, will not preserve bias. In the weak-exchange limit ($J_{\varepsilon}(t) \ll \delta B_z$), however, these transitions in the antiparallel spin subspace can be treated as a perturbation, resulting in a dipole-dipole interaction that is diagonal in the Z basis [44]. Letting the system evolve for a certain time results in a CZ gate up to single-qubit Z rotations [44,67–70]. Thus, the resulting gate will be bias preserving.

Trapped-ion qubits are encoded in the hyperfine energy levels of the ground state of a trapped ion. Qubits are manipulated using external lasers that couple to the spin and motion states of the ion [71]. The Mølmer-Sørensen (MS) two-qubit gate evolution operator is defined as $\hat{U}(\theta) = \exp\left(-i(\theta/2)\hat{\sigma}_1^{\phi_S}\hat{\sigma}_2^{\phi_S}\right)$ with $\hat{\sigma}_i^{\phi_S} = X_i \cos \phi_S + Y_i \sin \phi_S$, where ϕ_S is the effective spin rotation angle. The MS gate does not preserve the bias since

the Hamiltonian rotates the state around the X and Y axes [72]. However, the Liebfried-Sørensen (LS) gate [73], a variation of the MS gate, performs a rotation defined by the Hamiltonian $H_I = \hbar\chi Z_i Z_j$, where $\chi = \eta^2 \Omega^2 / (N\delta)$, η being the Lamb-Dicke parameter, Ω is the Rabi frequency of the oscillations between two internal states of the ion, N is the number of ions, and δ is the detuning between the laser and the ion transitions [74]. The evolution described by this operator is a CZ gate up to single-qubit Z rotations and, thus, preserves the bias.

In neutral- or Rydberg-atom devices, qubits are encoded in the hyperfine ground states of the atoms and are manipulated using laser beams. The electrons bound to an atom can be excited to a high-energy state, known as the Rydberg state. When an atom is excited to this state, it induces a shift in the energy levels of nearby atoms through dipole-dipole interactions, preventing multiatom excitation. This phenomenon, known as Rydberg blockade, is a fundamental mechanism for implementing two-qubit gates [75,76]. The CZ gate can be executed by applying three different laser pulses that couple the state $|1\rangle$ with the Rydberg state $|r\rangle$: the first one performs a π rotation on the control atom, then, a 2π rotation is applied to the target atom; and, finally, another π rotation is applied to the control atom [77]. The Rydberg CZ gate is bias preserving after appropriate conversion of leakage errors to Z-type errors [28].

Nitrogen-vacancy (N-V) centers in diamond store quantum information in an electron bound to the defect. In its ground state, the N-V center hosts an electronic spin system with total spin $S = 1$ and spin states ($m_S = \{-1, 0, 1\}$) [78]. There are two ways to perform a two-qubit gate. The first one is performed by using the magnetic dipolar interaction between two N-V centers. The damping operators of the dipolar interaction can be neglected, since the dipolar coupling v_{dip} is much weaker than the energy gap between the qubit states, so the interaction Hamiltonian can be expressed as $H_{\text{dip}} \approx \frac{1}{2} \hbar v_{\text{dip}} Z_A Z_B$ [79]. The evolution obtained is a CZ gate up to single-qubit rotations over the Z axis. The second way is by using the hyperfine interaction between a N-V center and a nitrogen or carbonous nucleus. Usually, a ^{14}N atom is used with the information encoded in the states $|m_n = +1/2\rangle \equiv |0\rangle$ and $|m_n = -1/2\rangle \equiv |1\rangle$. The interaction Hamiltonian is $H_I = \mathbf{S} \cdot \mathbf{A} \mathbf{I}$, where $\mathbf{S}$ and $\mathbf{I}$ are the electron spin and nuclear spin operators, respectively. $\mathbf{A}$ is the hyperfine interaction coupling tensor, where the perpendicular hyperfine interaction $\mathbf{A}_{\perp} = (A_x, A_y, 0)$ can be neglected [80–83]. Therefore, the CZ gate, which preserves the noise bias, can be implemented natively using both interactions.

Superconducting qubits generally do not exhibit highly biased noise when used as two-level systems, with bias values around $\eta \approx 10$ [6]. However, there are recent proposals based on two-level superconducting qubits that show high bias values [84]. Entangling gates can be

enabled by means of the cross-resonance (CR) gate [85,86]. This gate results from a ZX interaction Hamiltonian that results in a native CNOT gate, as shown in the main text. As explained before, the CZ gate using this interaction will not be bias preserving. Furthermore, a native CZ gate can also be enabled for transmon qubits by means of single-frequency-tunable (STC) transmon couplers that enable ZZ interactions between qubits. As seen before, this interaction results in CZ gates up to single-qubit Z rotations, indicating that they preserve bias [84,87,88].

APPENDIX D: THE XZZX SURFACE CODE

The $XZZX$ variant of the surface code can be obtained by means of a Clifford deformation of the CSS surface code [2,89]. Specifically, a Hadamard deformation is applied to every other qubit of the CSS surface code [4,46,89]. Up to the boundaries, the stabilizers of the resulting code are all of the $S_i = X_a Z_b Z_c X_d$ type, where the labels refer to the nearest-neighbor qubits. Specifically, we will consider the $XZZX$ code on a rotated layout with open boundary conditions, yielding parameters $[[d^2, 1, d]]$ [45], where d is the code distance. In contrast to the standard CSS surface code, every check qubit performs Z- and X-check measurements and, thus, it is a non-CSS code. The layout for this rotated $XZZX$ code is shown in Fig. 7. While the code is locally equivalent to the CSS surface code (inheriting its properties), it has an extremely good performance when




FIG. 7. The layout of the rotated $XZZX$ surface code. The white circles represent data qubits, while the brown circles represent check qubits. The green and orange squares represent Z and X stabilizers, respectively. It can be seen that the check qubits measure $XZZX$ patterns.

operating over biased noise, reaching a threshold of 50% at the infinite-bias limit for code-capacity noise [4]. The main reasons for such a performance boost comes from the fact that for the $XZZX$ code there is a single logical operator purely consisting of Z errors, significantly reducing the probability of it happening [2]; and because the defects created by the X and Z errors have particular directions (with perpendicular among them), implying that weighting a decoder for leveraging the fact that the noise is biased is very beneficial for decoding [26]. Furthermore, under biased noise, one can refine the weight of each stabilizer operator depending on the amount of noise. This refined notion of weight is called the *effective weight* of an operator [21]. For instance, under depolarizing noise, as all the errors X, Y, and Z are equally likely, the effective weight of an operator is simply the number of nontrivial Pauli operators in that operator. This coincides with the conventional definition of the Hamming weight. However, in the case of infinite biased noise ($P_Z \gg P_X$), the effective weight of each operator is ∞ if it contains at least one X or Y Pauli and, otherwise, its effective weight is the number of Z Pauli matrices in its decomposition. Therefore, the set of likely errors is reduced to those consisting of only Z errors, also depending on their effective weights. Under this scenario, as there exists only one logical operator in the rotated $XZZX$ code consisting solely of Z errors (the main diagonal error), the probability of a logical error in such a code is significantly reduced. Under other bias rates, using a similar argument, one can see that the probability of likely logical errors happening is reduced (in comparison to depolarizing noise), which can boost the performance of such codes.

A fundamental step for a QEC protocol is the syndrome-extraction circuit that is used in order to measure the stabilizers. The first important decision is which actual gates will be used to measure the stabilizers. For this rotated $XZZX$ code under consideration, we will use sequences of CNOT and CZ gates to measure the stabilizers by means of the check qubits, which will be prepared and measured in the Z basis and converted to the X basis by means of Hadamard gates. An $XZZX$ stabilizer-check measurement is depicted in Fig. 8. Furthermore, the operations involved in the stabilizer measurements of each check should be parallelized as much as possible to reduce the time steps in which the qubits are idling and, thus, subjected to more errors. This should be done in such a way that the circuit-level distance is as high as possible, since certain errors at the circuit level can reduce the distance that the code should show in theory, e.g., hook errors [54,57]. Thus, it is necessary to define a syndrome-extraction schedule. As explained in Ref. [57], the order of the physical data qubits in the circuits of a check measurement, as in Fig. 7, should follow S or Z patterns since, otherwise, commutation of nearby stabilizers is lost. Thus, the possible schedules involve $XZZX$ or $ZXXZ$ patterns. For our circuits, we




FIG. 8. An $XZZX$ -check measurement. The check qubit is colored in brown and the four data qubits involved in such a parity measurement are in white. The check qubit is prepared in $|0\rangle$ and mapped to the X basis by means of a Hadamard gate. Check measurement involves the same basis transformation to then be measured in the Z basis. The CNOT and CZ gates represent the $XZZX$ pattern.

specially select an $XZZX$ pattern since, as pointed out by Darmawan *et al.* in Ref. [26], an X error after the second entangling gate propagates the next two data qubits and aligns with the logical error strings. In Ref. [26], the authors discuss that since their CNOT gates are bias preserving, such a hook error is a low-rate error and, thus, is benign. In our circuits, the CNOT gates are not bias preserving, implying that in the $ZXXZ$ extraction pattern, the hook event will not be rare. Thus, we use the $XZZX$ pattern, as the CZ gate can be executed in a bias-preserving manner and such a hook error is again a low-rate event.

Furthermore, we discuss the compilation of the $XZZX$ surface code by means of CZ gates and Hadamard gates. For this case, the CNOT gates used for extracting the X parities are enabled by means of a CZ gate and two Hadamard gates sandwiching the target qubit. For this case, the extra Hadamards are included in the data qubits. Furthermore, the scheduling used to extract the syndrome allows the cancelation of pairs of Hadamard gates in the bulk of the extraction circuits, leading to a sequence of CZ gates, a layer of Hadamard gates over all data qubits, two sequences of CZ gates, another layer of Hadamard gates over all data qubits, and a last sequence of CZ gates. The resulting extraction circuit is the same as the one used by Google Quantum AI in Refs. [42,90]. The specific schedule can be found in the Supplementary Material of Ref. [90]. The superconducting processor used by Google in those works does not present bias toward dephasing noise, but they have compiled the $XZZX$ code using only CZ gates because those are native to their technology. Note that the gate pattern is the one that we have also used

for the $XZZX$ surface codes compiled with CNOT gates. Finally, note that in this case, all entangling gates will be bias preserving, indicating that the harmful hook errors commented on before will be very low-rate events for high bias values.

APPENDIX E: THE HYBRID BIASED-DEPOLARIZING CIRCUIT-LEVEL NOISE MODEL

Generic noise models that capture the most important characteristics of the noise in certain systems are invaluable for quantum-error-correction theorists to evaluate the methods they propose without having to dig into the specific details of a certain implementation. In this sense, the *standard depolarizing* (SD) circuit-level noise model is the most studied model considering imperfect syndrome-extraction circuits and it is extremely useful to determine if a certain QEC protocol can work whenever noise at the circuit level is considered [1,2,11,32,34,35,45,51,53,54,91–95]. Biased circuit-level noise models have also been considered in the literature; however, those models are specific for qubit platforms for which bias-preserving single- and two-qubit gates can be constructed [26,30,31,34,35]. From our discussion in Sec. II, none of those models capture the specific characteristics of two-level qubits, the noise of which is biased toward dephasing. Thus, we will propose a noise model that is generic enough so that theoretical evaluations of bias-tailored QEC over those platforms can be done while still capturing the essential properties of the noise.

The *hybrid biased-depolarizing circuit-level noise model* is defined by the tuple of parameters physical error rate, p , and bias, η , as follows:

(a) Hadamard gates are followed by single-qubit depolarizing errors at rate p , i.e., each X , Y , or Z error occurs with probability $p/3$ [96].

(b) CNOT (or CY) gates are followed by two-qubit depolarizing errors at rate p , i.e., each error from the set $\{I, X, Y, Z\}^{\otimes 2}/I^{\otimes 2}$ occurs with probability $p/15$.

(c) CZ gates are followed by two-qubit biased Pauli errors with bias η at rate p . Specifically, each pure dephasing error $\{ZI, IZ, ZZ\}$ occurs with probability $\eta p/(3(1 + \eta))$, while the rest occur with probability $p/(12(1 + \eta))$. By this encoding, the overall error rate of the gate is p independent of the bias, and the bias just changes how each of the errors contributes. The bias quantifies how much more probable pure dephasing errors are compared to the other ones.

(d) Preparing a state in the $|0\rangle$ ($|+\rangle$) state instead results in preparing the orthogonal state $|1\rangle$ ($|-\rangle$) with probability p .

(e) Measurement results in the X or Z basis experience a flip reporting an incorrect outcome with probability p .

(f) Idling steps experience single-qubit biased Pauli errors with bias η at rate p . Specifically, Z errors occur with probability $\eta p/(1 + \eta)$, while X and Y errors each occur with probability $p/(2(1 + \eta))$.

A slight difference of this model compared to previous biased circuit-level noise models is the way in which we define the probabilities of the biased and unbiased errors. In Refs. [26,30,31,34,35], the authors have used a noise model in which the biased errors occur with probability $p/(4^n - 1)$, while the rest occur at $p/((4^n - 1)\eta)$, where n is the number of qubits involved in the gate. Although this may be convenient for some cases, the overall gate error probability changes as a function of the bias, η . In fact, a higher bias implies a lower gate error rate. On the contrary, in our model, the overall gate error rate is the same for all bias values, with the change being in how the probabilities of each individual event are distributed [97]. We consider this to be a fairer comparison between depolarizing and biased cases, since the probability that any element fails will be the same, independent of the bias.

Note that for this model we have ignored the possibility that some of the gates may exhibit some residual bias, a case that is analyzed in Sec. III B. We do so because such residual bias will probably depend on the interaction used to form those gates and the goal is to have some generic circuit-level noise model considering bias in two-level qubits. Since the depolarizing case is the most demanding one (any residual bias will only improve the tailoring), then it is useful to study the performance of QECCs while being agnostic to the way in which gates that

do not preserve bias are constructed. The main feature is that the CZ gates are actually executed in a bias-preserving manner, i.e., by means of a native interaction that exhibits such a feature.

APPENDIX F: QUANTUM-ERROR-CORRECTION FOOTPRINTS

While the threshold is an interesting metric to study, since it basically provides the maximum physical error rate for which the QEC method is effective, the required physical-qubit footprints to be able to implement those are actually what matters in practice. Footprints refer to the space overhead in terms of physical qubits per logical qubit in order to execute quantum circuits of a certain size with a probability of approximately $1/e$ [43,98]. In this context, the *Megaquop* (10^6 quantum operations), *Gigaquop* (10^9 quantum operations), and *Teraquop* (10^{12} quantum operations) regimes are usually considered as the interesting regimes for fault-tolerant quantum computing. These relate to logical error rates per round of the order of 10^{-6} , 10^{-9} , and 10^{-12} , respectively. In order to make a fair comparison in terms of numbers, we will consider that the required distance to reach a certain regime is the first one exactly below the target error rates presented before. However, this should be taken with a grain of salt since, as we comment in the main text, there are cases for which a smaller distance is very close to the actual error rate.

APPENDIX G: NUMERICAL SIMULATIONS

Extensive Monte Carlo simulations of the *XZZX* code have been performed in order to estimate the performance of the codes considering the circuit-level noise model described. We have implemented the noisy extraction circuits in order to perform the sampling of the errors by using the STIM software package [99]. STIM considers the check measurements upon a set of syndrome extractions together with a final measurement of the data qubits. In order to conduct the memory experiments of the *XZZX* rotated surface code, two different memory experiments are conducted depending on which initial state is intended to be preserved: the horizontal (H) memory experiment, in which data qubits are initialized in an interchanging pattern of $|+\rangle$ and $|0\rangle$ states starting from $|+\rangle$ for the top-left data qubit and finally measured in their corresponding bases; and the vertical (V) memory experiment, in which data qubits are initialized in an interchanging pattern of $|0\rangle$ and $|+\rangle$ states starting from $|0\rangle$ for the top-left data qubit and finally measured in the corresponding basis. The horizontal- and vertical-memory names come from the direction of the strings that form logical operators in each of the cases. Note that the results of memory experiments are also applicable to understand logical operations performed by means of lattice surgery [45,100]. To fully

understand these protocols, stability experiments would also be required [49].

The operational figure of merit that we use to evaluate the performance is the logical error probability, p_L , per extraction round. We have run $3d$ rounds of syndrome extraction to reduce time-boundary edge effects coming from the fact that the first and last extraction rounds are less noisy than the ones in the bulk [45]. Once the circuits are run to collect the samples, we use the *PyMatching* implementation of the MWPM decoder in order to aim recovery and determine whether or not a logical error has occurred [2,51,95]. In order to have good enough statistical accuracy from the Monte Carlo simulation estimations, we have begun to collect samples with a ceiling of 20×10^6 circuit shots and 10^5 logical errors. For the lowest logical error probability points, i.e., $d = 13, 15$ at $p = 10^{-3}$, we have increased the shot ceiling up to 10×10^9 shots and 5×10^5 logical errors. We highlight regions showing p_L values for which the conditional probabilities $P(p_L|k)$ are within a factor of 1000 of the maximum-likelihood estimation, $p_L = k/n$, assuming a binomial distribution, only for the numerically estimated points—some are too small to be observed [45].

-
- [1] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).
 - [2] A. deMarti iOlius, P. Fuentes, R. Orús, P. M. Crespo, and J. Etzezarreta Martinez, Decoding algorithms for surface codes, *Quantum* **8**, 1498 (2024).
 - [3] S. S. Gill, A. Kumar, H. Singh, M. Singh, K. Kaur, M. Usman, and R. Buyya, Quantum computing: A taxonomy, systematic review and future directions, *Softw.: Pract. Exp.* **52**, 66 (2022).
 - [4] J. P. Bonilla Ataides, D. K. Tuckett, S. D. Bartlett, S. T. Flammia, and B. J. Brown, The XZZX surface code, *Nat. Commun.* **12**, 2172 (2021).
 - [5] J. Roffe, L. Z. Cohen, A. O. Quintavalle, D. Chandra, and E. T. Campbell, Bias-tailored quantum LDPC codes, *Quantum* **7**, 1005 (2023).
 - [6] J. Etzezarreta Martinez, P. Fuentes, P. M. Crespo, and J. Garcia-Frias, Approximating decoherence processes for the design and simulation of quantum error correction codes on classical computers, *IEEE Access* **8**, 172623 (2020).
 - [7] P. Aliferis and J. Preskill, Fault-tolerant quantum computation against biased noise, *Phys. Rev. A* **78**, 052331 (2008).
 - [8] Y. Seis, B. J. Brown, A. S. Sørensen, and J. F. Goodwin, Improving trapped-ion-qubit memories via code-mediated error-channel balancing, *Phys. Rev. A* **107**, 052417 (2023).
 - [9] M. A. Sepiol, A. C. Hughes, J. E. Tarlton, D. P. Nadlinger, T. G. Ballance, C. J. Ballance, T. P. Harty, A. M. Steane, J. F. Goodwin, and D. M. Lucas, Probing qubit memory errors at the part-per-million level, *Phys. Rev. Lett.* **123**, 110503 (2019).
 - [10] T. R. Tan, J. P. Gaebler, Y. Lin, Y. Wan, R. Bowler, D. Leibfried, and D. J. Wineland, Multi-element logic gates for trapped-ion qubits, *Nature* **528**, 380 (2015).
 - [11] B. Hetényi and J. R. Wootton, Tailoring quantum error correction to spin qubits, *Phys. Rev. A* **109**, 032433 (2024).
 - [12] K. Takeda, A. Noiri, T. Nakajima, T. Kobayashi, and S. Tarucha, Quantum error correction with silicon spin qubits, *Nature* **608**, 682 (2022).
 - [13] A. Noiri, K. Takeda, T. Nakajima, T. Kobayashi, A. Sammak, G. Scappucci, and S. Tarucha, Fast universal quantum gate above the fault-tolerance threshold in silicon, *Nature* **601**, 338 (2022).
 - [14] P. Steinacker *et al.*, Industry-compatible silicon spin-qubit unit cells exceeding 99% fidelity, *Nature* **646**, 81 (2025).
 - [15] N. Wang, C.-F. Liu, J.-W. Fan, X. Feng, W.-H. Leong, A. Finkler, A. Denisenko, J. Wrachtrup, Q. Li, and R.-B. Liu, Zero-field magnetometry using hyperfine-biased nitrogen-vacancy centers near diamond surfaces, *Phys. Rev. Res.* **4**, 013098 (2022).
 - [16] J. E. March, B. D. Wood, C. J. Stephen, L. D. Ferverza, B. G. Breeze, S. Mandal, A. M. Edmonds, D. J. Twitchen, M. L. Markham, O. A. Williams, and G. W. Morley, Long spin coherence and relaxation times in nanodiamonds milled from polycrystalline ^{12}C diamond, *Phys. Rev. Appl.* **20**, 044045 (2023).
 - [17] A. deMarti iOlius, J. Etzezarreta Martinez, P. Fuentes, P. M. Crespo, and J. Garcia-Frias, Performance of surface codes in realistic quantum hardware, *Phys. Rev. A* **106**, 062428 (2022).
 - [18] R. Lescanne, M. Villiers, T. Peronnin, A. Sarlette, M. Delbecq, B. Huard, T. Kontos, M. Mirrahimi, and Z. Leghtas, Exponential suppression of bit-flips in a qubit encoded in an oscillator, *Nat. Phys.* **16**, 509 (2020).
 - [19] C. Berdou *et al.*, One hundred second bit-flip time in a two-photon dissipative oscillator, *PRX Quantum* **4**, 020350 (2023).
 - [20] J. Fern and K. B. Whaley, Lower bounds on the nonzero capacity of Pauli channels, *Phys. Rev. A* **78**, 062335 (2008).
 - [21] Q. Xu, N. Mannucci, A. Seif, A. Kubica, S. T. Flammia, and L. Jiang, Tailored XZZX codes for biased noise, *Phys. Rev. Res.* **5**, 013035 (2023).
 - [22] D. K. Tuckett, A. S. Darmawan, C. T. Chubb, S. Bravyi, S. D. Bartlett, and S. T. Flammia, Tailoring surface codes for highly biased noise, *Phys. Rev. X* **9**, 041031 (2019).
 - [23] D. K. Tuckett, S. D. Bartlett, and S. T. Flammia, Ultrahigh error threshold for surface codes with biased noise, *Phys. Rev. Lett.* **120**, 050505 (2018).
 - [24] K. Tiurev, A. Pesah, Peter Jan H. S. Derks, J. Roffe, J. Eisert, M. S. Kesselring, and J.-M. Reiner, Domain wall color code, *Phys. Rev. Lett.* **133**, 110601 (2024).
 - [25] J. Guillaud and M. Mirrahimi, Repetition cat qubits for fault-tolerant quantum computation, *Phys. Rev. X* **9**, 041053 (2019).
 - [26] A. S. Darmawan, B. J. Brown, A. L. Grimsmo, D. K. Tuckett, and S. Puri, Practical quantum error correction

- with the XZZX code and Kerr-Cat qubits, *PRX Quantum* **2**, 030345 (2021).
- [27] S. Puri, L. St-Jean, J. A. Gross, A. Grimm, N. E. Frattini, P. S. Iyer, A. Krishna, S. Touzard, L. Jiang, A. Blais, S. T. Flammia, and S. M. Girvin, Bias-preserving gates with stabilized cat qubits, *Sci. Adv.* **6**, eaay5901 (2020).
- [28] I. Cong, H. Levine, A. Keesling, D. Bluvstein, S.-T. Wang, and M. D. Lukin, Hardware-efficient, fault-tolerant quantum computation with Rydberg atoms, *Phys. Rev. X* **12**, 021049 (2022).
- [29] J. Preskill, Crossing the quantum chasm: From NISQ to fault tolerance, <http://theory.caltech.edu/preskill/talks/Preskill-Q2B-2023> (2023), accessed October 26, 2024.
- [30] E. Gouzien, D. Ruiz, F.-M. Le Régent, J. Guillaud, and N. Sangouard, Performance analysis of a repetition cat code architecture: Computing 256-bit elliptic curve logarithm in 9 hours with 126 133 cat qubits, *Phys. Rev. Lett.* **131**, 040602 (2023).
- [31] D. Ruiz, J. Guillaud, A. Leverrier, M. Mirrahimi, and C. Vuillot, LDPC-cat codes for low-overhead quantum computing in 2D, *Nat. Commun.* **16**, 1040 (2025).
- [32] C. Gidney and M. Ekerå, How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits, *Quantum* **5**, 433 (2021).
- [33] J. Preskill, Beyond NISQ: The megaquop machine, *ACM Trans. Quantum Comput.* **6**, 1 (2025).
- [34] O. Higgott, T. C. Bohdanowicz, A. Kubica, S. T. Flammia, and E. T. Campbell, Improved decoding of circuit noise and fragile boundaries of tailored surface codes, *Phys. Rev. X* **13**, 031007 (2023).
- [35] C. Chamberland and E. T. Campbell, Universal quantum computing with twist-free and temporally encoded lattice surgery, *PRX Quantum* **3**, 010331 (2022).
- [36] C. T. Hann, K. Noh, H. Putterman, M. H. Matheny, J. K. Iverson, M. T. Fang, C. Chamberland, O. Painter, and F. G. S. L. Brandão, Hybrid cat-transmon architecture for scalable, hardware-efficient quantum error correction, *PRX Quantum* **6**, 030305 (2025).
- [37] H. Putterman *et al.*, Hardware-efficient quantum error correction via concatenated bosonic qubits, *Nature* **638**, 927 (2025).
- [38] D. Forlivesi, L. Valentini, and M. Chiani, Quantum codes for asymmetric channels: ZZZY surface codes, *IEEE Commun. Lett.* **28**, 2233 (2024).
- [39] We observe that the output bias is in the range $\eta_{\text{Had}} \in 0.5, 0.6$, i.e., almost depolarizing for every case. Recall that for the single-qubit case, the bias is quantified as $\eta = p_z / (p_x + p_y)$, with $1/2$ being the depolarizing case.
- [40] T. Müller, T. Stollenwerk, D. Headley, M. Epping, and F. K. Wilhelm, Coherent and non-unitary errors in ZZ-generated gates, *Quantum Sci. Technol.* **10**, 015058 (2024).
- [41] Note that there can be many other ways of enabling such gates in those platforms. Our intention here is to show that there are ways that preserve bias and ways that do not.
- [42] Google Quantum AI and Collaborators, Quantum error correction below the surface code threshold, *Nature* **638**, 920 (2024).
- [43] C. Gidney, M. Newman, A. Fowler, and M. Broughton, A fault-tolerant honeycomb memory, *Quantum* **5**, 605 (2021).
- [44] Z. Cai, M. A. Fogarty, S. Schaal, S. Patomäki, S. C. Benjamin, and J. J. L. Morton, A silicon surface code architecture resilient against leakage errors, *Quantum* **3**, 212 (2019).
- [45] A. R. O'Rourke and S. Devitt, Compare the pair: Rotated versus unrotated surface codes at equal logical error rates, *Phys. Rev. Res.* **7**, 033074 (2025).
- [46] R. M. Otxoa, J. Etxezarreta Martinez, P. Schnabl, N. Mertig, C. Smith, and F. Martins, SpinHex: A low-crosstalk, spin-qubit architecture based on multi-electron couplers, *arXiv:2504.03149*.
- [47] Riverlane, The quantum error correction report, <https://www.riverlane.com/quantum-error-correction-report-2024> (2024), accessed January 5, 2025.
- [48] D. Litinski, A game of surface codes: Large-scale quantum computing with lattice surgery, *Quantum* **3**, 128 (2019).
- [49] C. Gidney, Stability experiments: The overlooked dual of memory experiments, *Quantum* **6**, 786 (2022).
- [50] C. K. Long, N. J. Mayhall, S. E. Economou, E. Barnes, C. H. W. Barnes, F. Martins, D. R. M. Arvidsson-Shukur, and N. Mertig, Minimal state-preparation times for silicon spin qubits, *npj Quantum Inf.* **11**, 113 (2025).
- [51] O. Higgott and C. Gidney, SPARSE BLOSSOM: Correcting a million errors per core second with minimum-weight matching, *Quantum* **9**, 1600 (2025).
- [52] S. Wolanski and B. Barber, Ambiguity clustering: An accurate and efficient decoder for qLDPC codes, *arXiv:2406.14527*.
- [53] A. deMarti iOlius, I. Etxezarreta Martinez, J. Roffe, and J. Etxezarreta Martinez, An almost-linear time decoding algorithm for quantum LDPC codes under circuit-level noise, *arXiv:2409.01440*.
- [54] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, *Nature* **627**, 778 (2024).
- [55] M. Rennela and H. Ollivier, Low bit-flip rate probabilistic error cancellation, *Adv. Quantum Technol.* **8**, e2500029 (2025).
- [56] J. Etxezarreta Martinez, <https://github.com/jetxezarreta/qec-two-level-qubits-circuit-noise-bias> qec-two-level-qubits-circuit-noise-bias (2025).
- [57] Y. Tomita and K. M. Svore, Low-distance surface codes under realistic quantum noise, *Phys. Rev. A* **90**, 062320 (2014).
- [58] This is sufficient since the noise can be transformed into stochastic Pauli noise by means of Pauli twirling [2,6]. In fact, Pauli noise is the one generally considered for theoretical QEC analysis [99].
- [59] M. Fellous-Asiani, M. Naseri, C. Datta, A. Streltsov, and M. Ozmaniec, Scalable noisy quantum circuits for biased-noise qubits, *npj Quantum Inf.* **11**, 145 (2025).
- [60] J. Johansson, P. Nation, and F. Nori, QuTiP: An open-source PYTHON framework for the dynamics of open quantum systems, *Comput. Phys. Commun.* **183**, 1760 (2012).
- [61] S. T. Flammia and J. J. Wallman, Efficient estimation of Pauli channels, *ACM Trans. Quantum Comput.* **1**, 1 (2020).
- [62] E. van den Berg, Z. K. Mineev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse Pauli-Lindblad

- models on noisy quantum processors, *Nat. Phys.* **19**, 1116 (2023).
- [63] G. Burkard, T. D. Ladd, A. Pan, J. M. Nichol, and J. R. Petta, Semiconductor spin qubits, *Rev. Mod. Phys.* **95**, 025003 (2023).
- [64] M. Veldhorst, C. H. Yang, J. C. C. Hwang, W. Huang, J. P. Dehollain, J. T. Muhonen, S. Simmons, A. Laucht, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak, A two-qubit logic gate in silicon, *Nature* **526**, 410 (2015).
- [65] T. F. Watson, S. G. J. Philips, E. Kawakami, D. R. Ward, P. Scarlino, M. Veldhorst, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen, A programmable two-qubit quantum processor in silicon, *Nature* **555**, 633 (2018).
- [66] D. M. Zajac, A. J. Sigillito, M. Russ, F. Borjans, J. M. Taylor, G. Burkard, and J. R. Petta, Resonantly driven CNOT gate for electron spins, *Science* **359**, 439 (2018).
- [67] M. Russ, D. M. Zajac, A. J. Sigillito, F. Borjans, J. M. Taylor, J. R. Petta, and G. Burkard, High-fidelity quantum gates in Si/SiGe double quantum dots, *Phys. Rev. B* **97**, 085421 (2018).
- [68] T. Meunier, V. E. Calado, and L. M. K. Vandersypen, Efficient controlled-phase gate for single-spin qubits in quantum dots, *Phys. Rev. B* **83**, 121403(R) (2011).
- [69] M. Veldhorst, C. H. Yang, J. C. C. Hwang, W. Huang, J. P. Dehollain, J. T. Muhonen, S. Simmons, A. Laucht, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak, A two-qubit logic gate in silicon, *Nature* **526**, 410 (2015).
- [70] J. Y. Huang *et al.*, High-fidelity spin qubit operation and algorithmic initialization above 1 K, *Nature* **627**, 772 (2024).
- [71] D. Jonathan, M. B. Plenio, and P. L. Knight, Fast quantum gates for cold trapped ions, *Phys. Rev. A* **62**, 042307 (2000).
- [72] T. Manovitz, Y. Shapira, L. Gazit, N. Akerman, and R. Ozeri, Trapped-ion quantum computer with robust entangling gates and quantum coherent feedback, *PRX Quantum* **3**, 010347 (2022).
- [73] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, C. Jelenković, Band Langer, T. Rosenband, and D. J. Wineland, Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate, *Nature* **422**, 412 (2003).
- [74] G. Milburn, S. Schneider, and D. James, Ion trap quantum computing with warm ions, *Fortschr. Phys.* **48**, 801 (2000).
- [75] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, Fast quantum gates for neutral atoms, *Phys. Rev. Lett.* **85**, 2208 (2000).
- [76] M. D. Lukin, M. Fleischhauer, R. Cote, L. M. Duan, D. Jaksch, J. I. Cirac, and P. Zoller, Dipole blockade and quantum information processing in mesoscopic atomic ensembles, *Phys. Rev. Lett.* **87**, 037901 (2001).
- [77] M. Saffman, T. G. Walker, and K. Mølmer, Quantum information with Rydberg atoms, *Rev. Mod. Phys.* **82**, 2313 (2010).
- [78] L. Childress and R. Hanson, Diamond NV centers for quantum computing and quantum networks, *MRS Bull.* **38**, 134 (2013).
- [79] F. Dolde, I. Jakobi, B. Naydenov, N. Zhao, S. Pezzagna, C. Trautmann, J. Meijer, P. Neumann, F. Jelezko, and J. Wrachtrup, Room-temperature entanglement between single defect spins in diamond, *Nat. Phys.* **9**, 139 (2013).
- [80] R. Finsterhoelzl, W.-R. Hanne, and G. Burkard, High-fidelity entangling gates for electron and nuclear spin qubits in diamond, *Phys. Rev. B* **111**, 214104 (2025).
- [81] M. S. Everitt, S. Devitt, W. J. Munro, and K. Nemoto, High-fidelity gate operations with the coupled nuclear and electron spins of a nitrogen-vacancy center in diamond, *Phys. Rev. A* **89**, 052317 (2014).
- [82] J. H. Shim, I. Niemeyer, J. Zhang, and D. Suter, Room-temperature high-speed nuclear-spin quantum memory in diamond, *Phys. Rev. A* **87**, 012301 (2013).
- [83] Z.-T. Liang, X. Yue, Q. Lv, Y.-X. Du, W. Huang, H. Yan, and S.-L. Zhu, Proposal for implementing universal superadiabatic geometric quantum gates in nitrogen-vacancy centers, *Phys. Rev. A* **93**, 040305(R) (2016).
- [84] M. Hays, J. Kim, and W. D. Oliver, Nondegenerate noise-resilient superconducting qubit, *PRX Quantum* **6**, 040321 (2025).
- [85] C. Rigetti and M. Devoret, Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies, *Phys. Rev. B* **81**, 134507 (2010).
- [86] J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, J. A. Smolin, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, Simple all-microwave entangling gate for fixed-frequency superconducting qubits, *Phys. Rev. Lett.* **107**, 080502 (2011).
- [87] F. Yan, P. Krantz, Y. Sung, M. Kjaergaard, D. L. Campbell, T. P. Orlando, S. Gustavsson, and W. D. Oliver, Tunable coupling scheme for implementing high-fidelity two-qubit gates, *Phys. Rev. Appl.* **10**, 054062 (2018).
- [88] R. Li, K. Kubo, Y. Ho, Z. Yan, Y. Nakamura, and H. Goto, Realization of high-fidelity CZ gate based on a double-transmon coupler, *Phys. Rev. X* **14**, 041050 (2024).
- [89] A. Dua, A. Kubica, L. Jiang, S. T. Flammia, and M. J. Gullans, Clifford-deformed surface codes, *PRX Quantum* **5**, 010347 (2024).
- [90] R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, D. Bacon *et al.*, Suppressing quantum errors by scaling a surface code logical qubit, *Nature* **614**, 676 (2023).
- [91] A. Gong, S. Cammerer, and J. M. Renes, Toward low-latency iterative decoding of QLDPC codes under circuit-level noise, [arXiv:2403.18901](https://arxiv.org/abs/2403.18901).
- [92] A. Siegel, A. Strikis, and M. Fogarty, Towards early fault tolerance on a $2 \times n$ array of qubits equipped with shuttling, *PRX Quantum* **5**, 040328 (2024).
- [93] A. deMarti iOlus and J. Etxezarreta Martinez, The closed-branch decoder for quantum LDPC codes, [arXiv:2402.01532](https://arxiv.org/abs/2402.01532).
- [94] T. Hillmann, L. Berent, A. O. Quintavalle, J. Eisert, R. Wille, and J. Roffe, Localized statistics decoding for quantum low-density parity-check codes, *Nat. Commun.* **16**, 8214 (2025).
- [95] O. Higgott, PyMatching, <https://github.com/oscarhiggott/PyMatching>, accessed July 26, 2024.

- [96] We only discuss Hadamard gates for the noise model since those are generally the ones appearing in QEC syndrome-extraction circuits. In a more generic case, single-qubit gates related to rotations involving the X or Y basis will be followed by depolarizing noise, while rotations strictly over the Z axis can be executed in a bias-preserving manner.
- [97] Note that the standard depolarizing case does not occur at $\eta = 1/2$ since at such a value, the two-qubit gate errors will occur with probability $p/9$ for the biased errors and with $p/18$ for the rest. Thus, the standard depolarizing case has to be defined on its own.
- [98] S. J. S. Tan, C. A. Pattison, M. McEwen, and J. Preskill, Resilience of the surface code to error bursts, [arXiv:2406.18897](#).
- [99] C. Gidney, STIM: A fast stabilizer circuit simulator, [Quantum](#) **5**, 497 (2021).
- [100] D. Horsman, A. G. Fowler, S. Devitt, and R. V. Meter, Surface code quantum computing by lattice surgery, [New J. Phys.](#) **14**, 123011 (2012).