

Supplementary Information for *Unfolded distillation: very low-cost magic state preparation for biased-noise qubits*

Diego Ruiz,^{1,2,*} Jérémie Guillaud,¹ Christophe Vuillot,¹ and Mazyar Mirrahimi²

¹*Alice & Bob, 49 Bd du Général Martial Valin, 75015 Paris, France*

²*Laboratoire de Physique de l'École Normale Supérieure,
École Normale Supérieure, Centre Automatique et Systèmes,
Mines Paris, Université PSL, CNRS, Inria, Paris, France*


(Dated: November 17, 2025)

I. PRE- AND POSTSELECTION COST

In this section, we evaluate the contribution of preselection and postselection to the time overhead in unfolded distillation protocols. The results are summarized in Supplementary Figure 1. With the exception of the case where preselection is applied at a relatively high phase-flip error rate of $p_Z = 5 \times 10^{-3}$, the mean number of rounds is largely dominated by the base number of rounds. These correspond to the repeated X -type stabilizer measurements required to reliably establish their values. At $p_Z = 5 \times 10^{-3}$, however, the time overhead due to preselection exceeds that of the base rounds. By contrast, the preselection-free scheme described in Method ‘Unfolded distillation without just-in-time decoding or preselection’ exhibits a more favorable scaling with the physical error rate. We note that preselection could in principle be replaced by postselection on the final rounds of X -stabilizer measurements. However, this would incur a higher time cost, since the base rounds would need to be performed again if the postselection fails.

II. CNOT SCHEDULING


In the scheme described in Section ‘Distillation with the unfolded code’, it is necessary to determine a CNOT gate schedule such that the gates do not overlap and all stabilizers are measured within four timesteps. Additionally, it is advantageous for every qubit to which the $X^{-1/4}$ gate is applied to be involved in a CNOT interaction immediately beforehand. If this is not the case, the Z idling errors occurring between the CNOT and the $X^{-1/4}$ gate directly contribute to the physical input error rate of the distillation protocol. However, the layout shown in Figure 5 does not permit all required CNOTs between the ancilla and data qubits supporting the transversal $X^{1/4}$ gate to be performed in a single timestep. While in the scheme from Method ‘Unfolded distillation without just-in-time decoding or preselection’, the issue can be avoided by simply staggering the $X^{-1/4}$ operations across multiple timesteps, the scheme in Section ‘Distillation with the unfolded code’ requires a



Supplementary Figure 1. Time overhead decomposition for the unfolded distillation schemes listed in Table I (shown in the same order as in the table). Each bar indicates the contribution of the different categories to the mean number of rounds required to produce a magic state. The base number of rounds corresponds to the repeated rounds of X -type stabilizer measurements. Flag and complementary gap postselection are included only for the surface code scheme described in Section ‘Scheme with moderate noise bias’. The preselection scheme refers to the protocol described in Section ‘Distillation with the unfolded code’ while the preselection-free scheme refers to the protocol described in Method ‘Unfolded distillation without just-in-time decoding or preselection’.

different strategy, as $X^{-1/4}$ gates are only applied when the preselection condition is satisfied. One viable solution is to introduce an additional qubit and stabilizer, as illustrated in Supplementary Figure 2. Thus, by introducing only two additional qubits, it becomes possible to apply a CNOT gate to each qubit involved in the $X^{1/4}$ transversal gate immediately beforehand.

* diego.ruiz@alice-bob.com




Supplementary Figure 2. CNOT scheduling for the unfolding distillation protocol. In this protocol, an $X^{1/4}$ gate is applied transversally to all qubits shown in red. However, any phase-flip error occurring before this gate can be transformed into a bit-flip error afterward, potentially leading to a logical error in the distillation process. To minimize this possibility, it is crucial to reduce the time interval between the application of the $X^{-1/4}$ gate to a red qubit and the final CNOT involved in measuring the X -type stabilizer on that qubit. To achieve this, it is convenient to introduce an additional data qubit and an ancilla qubit, allowing all CNOTs on the qubits involved in the transversal $X^{1/4}$ gate to be performed simultaneously in the last step of the stabilizer measurements. This synchronization would not be possible without the addition of extra qubits.

III. UNFOLDED DISTILLATION WITH HIGHER PHYSICAL ERROR RATE

Unfolded distillation, unlike schemes based on Clifford measurements, remains robust even at relatively high physical error rates. This is illustrated in Supplementary Figure 3, which shows the logical error rate of a magic state prepared with a physical error rate of $p_Z = 5 \times 10^{-3}$. At this error rate, the repetition code distance must be increased from $d = 9$ to $d = 11$ to avoid limiting the logical fidelity of the magic state. Despite the high error rate, the logical error continues to follow the expected $35p^3$ scaling behavior characteristic of distillation. However, due to the need for preselection, the number of rounds required for preparation increases. In this case, approxi-


mately 18 rounds are necessary.



Supplementary Figure 3. Logical error probability of the $|X^{1/4}\rangle$ magic state prepared using the unfolded distillation protocol. A circuit-level physical phase-flip error rate of $p_Z = 5 \times 10^{-3}$ per operation is assumed. Due to the noise in X -type stabilizer measurements, a minimum number of rounds is required to achieve the expected distillation logical error rate of $35p^3$, where p is the error probability of the $X^{-1/4}$ gate. Preselection refers to repeatedly measuring the X -type stabilizers until, for a specified number of consecutive rounds, all outcomes are +1. The mean number of rounds includes both preselection and the final acceptance check based on the Z -type stabilizers of the Hadamard Reed-Muller code. At the assumed physical error rate, preselection becomes costly. As a result, two rounds of preselection and approximately 18 total rounds are required to reach the maximum achievable magic state fidelity. The inset shows how the mean number of rounds scales with the number of preselection rounds. This scaling is exponential, and at this level of physical noise, the scheme presented in Method ‘Unfolded distillation without just-in-time decoding or preselection’ becomes preferable.

In such high-error regimes, the alternative scheme described in Method ‘Unfolded distillation without just-in-time decoding or preselection’ is more efficient, as it avoids preselection altogether. Supplementary Figure 4 shows the resulting magic state fidelity under the same physical error rate $p_Z = 5 \times 10^{-3}$. Taking into account the additional factor of $3/2$ (see Method ‘Unfolded distillation without just-in-time decoding or preselection’) and the usual distillation postselection, the total number of rounds is given by 16.4, and the scheme also achieves a slightly lower logical error rate compared to the scheme presented in Supplementary Figure 3. However, the logical error rate as a function of the number of rounds does not converge to the $35p^3$ limit set by distillation. Empirical analysis of failing configurations reveals that logi-

cal failures result from higher-weight errors (greater than three), involving both $X^{-1/4}$ gates and errors occurring during the error correction cycles before and after the transversal $X^{1/4}$ gate. Although such high-weight errors are negligible when $p_Z = 10^{-3}$, as shown in Figure 14, they nonetheless become important at $p_Z = 5 \times 10^{-3}$.



Supplementary Figure 4. Logical error probability of a $|X^{1/4}\rangle$ magic state prepared using the unfolding distillation protocol presented in Method ‘Unfolded distillation without just-in-time decoding or preselection’. A circuit-level physical phase-flip error rate of $p_Z = 5 \times 10^{-3}$ per operation is assumed. A minimum number of X -type stabilizer rounds is required before and after the application of the transversal $X^{1/4}$ gate to reach a plateau in the logical error rate. This approach avoids preselection, and the total number of rounds must be multiplied by a factor of $3/2$ compared to the simulation, as explained in Method ‘Unfolded distillation without just-in-time decoding or preselection’. Although higher-order errors prevent reaching the $35p^3$ bound set by distillation (see Section III), where p denotes the error rate of the $X^{-1/4}$ gate, a magic state with a logical error rate of 2×10^{-5} can still be prepared using 15 rounds of error correction.