Effective Error Channels for Magic State Distillation

2025 Global Quantum Leap IRTE Project

Aiden Karpf
Physics and Astronomy Department, Pomona College
arka2021@mymail.pomona.edu

Advisors:
Principal Investigator: Dr. Markus Müller
Research Mentors: Berat Yenilen, Dr. Manuel Rispler

Abstract

Quantum error correction (QEC) is necessary to close the gap between the logical error rates required by quantum algorithms and those achievable on experimental hardware. Universal quantum computation requires non-Clifford gates, which are costly to implement fault-tolerantly and make classical simulation computationally difficult. As a result, simulating and researching QEC-encoded circuits that include many qubits and non-Clifford elements (e.g., fault-tolerant magic-state distillation) quickly becomes infeasible. In this work, we investigate whether effective error channels (EECs) can be defined at the logical level to model the physical Clifford noise associated with QEC-encoded circuits. Through quantum process tomography, we construct EECs for a Steane-encoded magic-state distillation circuit [1] and simulate the logical circuit with these EECs. We obtain estimates for how physical Clifford noise (modeled by the EECs) alters the acceptance probability and undetected error probability of the logical circuit. These results serve as a proof of concept suggesting that EECs can help researchers explore difficult-to-simulate non-Clifford circuits.

This work was funded by the National Science Foundation.

RWTH Aachen — Aachen, Germany — Summer 2025

1 Introduction

Given the inherent sensitivity of quantum systems to noise, it is necessary to encode redundancy into quantum information using Quantum Error Correction (QEC). To do this, many common QEC codes—e.g., stabilizer codes—are constructed with quantum gates from the Clifford group,

$$\mathcal{C}_n = \langle H,\ S,\ \text{CNOT} \rangle,$$

the group generated by the Hadamard gate H, phase gate S, and controlled-not gate CNOT.[1] One reason that these stabilizer codes are the leading class of QEC codes in the literature is that they can be simulated efficiently on a classical computer (Gottesman-Knill theorem) [2] and thus are easy to study. However, for universal quantum computation, at least one non-Clifford gate ($\notin \mathcal{C}_n$) must be included within the set of generating gates to perform an arbitrary quantum operation.

In order to prevent the uncorrectable spread of errors, ideal quantum computation also requires fault-tolerance on the encoded level. The easiest way to implement logical gates fault-tolerantly is by performing them transversally (no entangling gates within single logical qubits). This prevents the spread of errors within each codeblock of physical qubits that contains the information for a single logical qubit. However, the Eastin-Knill theorem tells us that no QEC code can implement a universal gate set transversally [3]. Thus, as most QEC codes are designed to implement the Clifford gates transversally, this makes the logical implementation of non-Clifford gates noisy and difficult.

1.1 Magic States: Injection and Distillation

To get around these difficulties, researchers often turn to "magic" states—non-stabilizer states which can be injected to implement a logical non-Clifford gate. To make this a bit more concrete, let's turn to an example:

$$|H\rangle = Y\left(\frac{\pi}{4}\right)|0\rangle = \cos\left(\frac{\pi}{8}\right)|0\rangle + \sin\left(\frac{\pi}{8}\right)|1\rangle\,.$$

This is the $|H\rangle$ magic state, an eigenstate of the Hadamard operator, and it is defined as the $|0\rangle$ state rotated by $\frac{\pi}{4}$ radians around the y-axis of the Bloch sphere.[2] The purpose of this $|H\rangle$ magic state is that it can be injected to realize a non-Clifford $Y\left(\pm\frac{\pi}{4}\right)$ rotation via a circuit containing only Clifford gates:

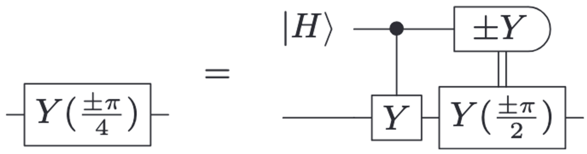

Figure 1: $|H\rangle$ magic state injection circuit to realize a non-Clifford $Y\left(\pm\frac{\pi}{4}\right)$ rotation.

Here, the bottom qubit is the data qubit, $|\psi\rangle$, for which we want to apply a $Y\left(\pm\frac{\pi}{4}\right)$ rotation, and the top qubit is an ancilla qubit for which we initialize our $|H\rangle$ magic state. By entangling the qubits with a CY-gate and then applying a classically-conditioned $Y\left(\pm\frac{\pi}{2}\right)$ correction rotation, we realize the state $Y(\pm\frac{\pi}{4})|\psi\rangle$ as desired. This is extremely useful as it allows us to implement a non-Clifford gate using only Clifford gates (which are easily implemented fault-tolerantly) and a single magic state. Thus, the bottleneck becomes the preparation of these magic states, which is often difficult and noisy.

This is where distillation comes in. Magic-state distillation schemes take many lower-fidelity magic states as input and output fewer higher-fidelity magic states. For example, the distillation circuit which is the main subject of this work is displayed in Figure 2. This circuit takes 8 noisy $|H\rangle$ magic states as input and produces a single, higher-fidelity $|\text{Toffoli}\rangle$ state as output. This $|\text{Toffoli}\rangle$ state can then be injected into another logical circuit to realize a Toffoli gate, which is a commonly used gate in quantum algorithms.[3]

[1] Formally, the Clifford group is the normalizer of the Pauli group:

$$\mathcal{C}_n \;=\; \{\, U \in U(2^n) \mid U\mathcal{P}_n U^\dagger = \mathcal{P}_n \,\},$$

but defining the group by a set of generating gates can be more straightforward pedagogically.

[2] For those familiar with the T-gate, this $Y\left(\frac{\pi}{4}\right)$-gate is essentially a T-gate in the y-basis.

[3] As we are primarily concerned with the present *distillation* scheme, see [1] for more details on the $|\text{Toffoli}\rangle$ injection.

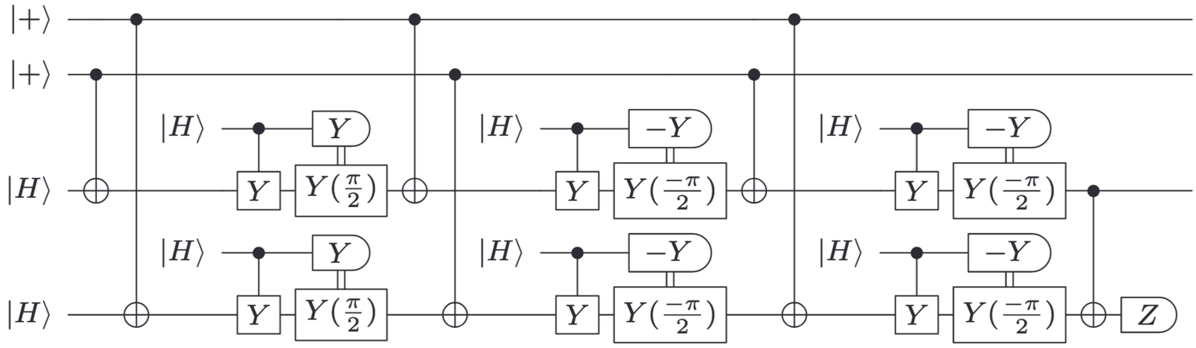

Figure 2: H-to-Toffoli distillation circuit from Eastin [1].

1.2 Research Question

Recall that the Gottesman-Knill theorem says that stabilizer (Clifford-only) circuits can be simulated efficiently on a classical computer. The distillation circuit of Figure 2, however, contains non-Clifford $|H\rangle$ magic states, and thus we are unable to simulate it efficiently.[4] At the logical level, this is not so problematic as our circuit only requires 6 qubits (4 data and 2 ancilla) and thus a 2^6-dimensional state vector. That being said, what we ultimately care about is the QEC-encoded version of this circuit (containing many qubits), which is very difficult to simulate classically. If researchers want to truly investigate the error rates of a distillation scheme (or any non-Clifford circuit), they would need to simulate the fully-encoded version of the circuit since the noise will depend on the underlying QEC code. Unfortunately, this is often computationally infeasible given the simulation's exponential scaling.

One possible approach in addressing this problem is to define "effective error channels" that model the noise of components encoded with the underlying QEC code. These effective error channels (EECs) can then be applied to the logical circuit, ideally reproducing the same logical error rates present in the fully-encoded version of the circuit. In particular, this approach would significantly aid further study into magic state distillation and injection. It is not clear, however, if this EEC approach is equivalent to simulating the full QEC-encoded circuit, and this has not been investigated to our knowledge in the literature.[5]

Thus, we are motivated to investigate the following question: can effective error channels be defined to accurately model the noise in QEC-encoded components of a logical circuit? In other words—by applying these EECs to a *logical* distillation circuit, can we precisely reproduce the error rates of an entire *QEC-encoded* distillation circuit? A "cartoon" depiction of our research question is displayed in Figure 3.

In the following section, we describe and reproduce results for the logical H-to-Toffoli circuit that is the focus of this work. Next, we will encode this logical circuit with the Steane code and construct effective error channels to model the noise of these encoded circuit components. We then apply these effective error channels to the logical circuit and examine how the behavior of the logical circuit is adjusted. Finally, we will discuss the outlook of simulating the full QEC-encoded circuit for comparison, and summarize the project as a whole.

2 H-to-Toffoli Distillation Circuit

In this work, we are primarily concerned with the H-to-Toffoli distillation circuit from Figure 2. To reiterate, this distillation scheme takes noisy $|H\rangle$ states as input, and outputs a single, higher-fidelity $|\text{Toffoli}\rangle = (|000\rangle + |100\rangle + |010\rangle + |111\rangle)/2$ state that can be injected to realize a Toffoli gate. It can be shown that—given a Pauli noise model—the $|H\rangle$ states inputted in this scheme can be assumed to only suffer Y errors with probability p [1]. This serves as the main source of error for the distillation scheme, and—when Eastin analyzes his circuit—he assumes that all Clifford gates are noiseless.

[4] To the best of our knowledge; technically, "quantum advantage" (and the classical hardness of simulating non-Clifford circuits) remains an open question. That being said, it is widely believed that non-Clifford circuits cannot be efficiently simulated on a classical computer, as such an ability would eliminate any computational advantage of quantum computers.

[5] This lack of investigation in the literature is almost certainly due to the extreme difficulty of simulating the entire QEC-encoded circuit for comparison with most distillation schemes.

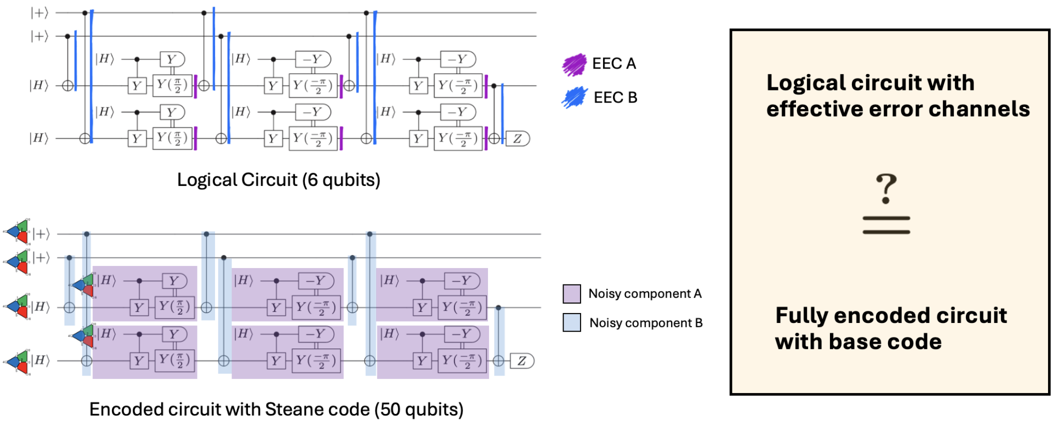

Figure 3: Cartoon depicting our research question. Displayed are the two circuits we are interested in comparing: the logical circuit with defined effective error channels (Top) and the QEC-encoded circuit (Bottom).

Notice that our outputted $|\text{Toffoli}\rangle$ state consists of only three qubits despite our distillation circuit containing four logical qubits—this is because the bottom qubit of Figure 2 serves as a "flag" that can detect errors. If the flag measurement yields a nontrivial outcome, the outputted three-qubit state is discarded and the distillation circuit is run again. This reduces the error rate of our accepted $|\text{Toffoli}\rangle$ state. In fact, if we assume perfect Cliffords and only errors on the $|H\rangle$ states occurring with probability p, we can detect and discard all weight-one errors.[6]

Through exhaustive counting, the probability that we accept the output of this circuit can be shown to be [1]:

$$a(p) = 1 - 8p + 56p^2 - 224p^3 + 560p^4 - 896p^5 + 896p^6 - 512p^7 + 128p^8. \tag{1}$$

A bit of intuition for this result to linear order: there are eight possible locations for error events, and all of these single weight-one errors are detected and discarded by our flag.

Another useful result from the Eastin paper is the undetected error probability, $a(p)e(p)$. This simply represents the probability of a logical error on our output state given we accept the output:

$$a(p)e(p) = 28p^2 - 168p^3 + 476p^4 - 784p^5 + 784p^6 - 448p^7 + 112p^8. \tag{2}$$

Some basic intuition for this result to leading order: there are $\binom{8}{2} = 28$ combinations of two weight-one errors (occurring with probability p^2) which will not be detected by our flag measurement and result in a logical error.

[6]See that in Figure 2, these Y errors would propagate straight through each injection, and then through the final CNOT, flipping the measurement.

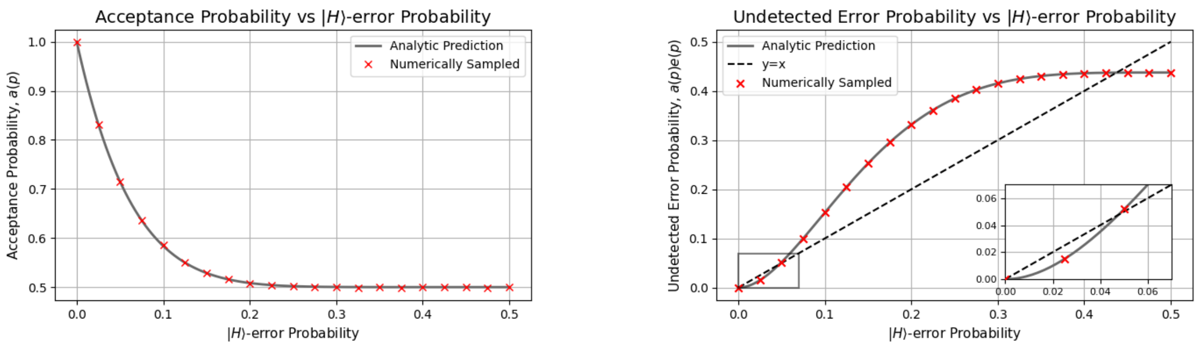

(a) $a(p)$

(b) $a(p)e(p)$

Figure 4: Reproduction of analytic results for the logical H-to-Toffoli distillation circuit assuming perfect Cliffords (as done by Eastin).

As an initial exercise, I numerically reproduced these two analytic results from the Eastin paper for the logical circuit. My simulated results for $a(p)$ and $a(p)e(p)$ and their corresponding analytic curves are displayed in Figure 4. The numerics agree well with Eastin's analytic results (1) and (2), which provides us with confidence that the logical circuit has been implemented correctly.

Additionally, notice the inset in Figure 4b, which displays the region at which our distillation scheme is actually effective. The line $y = x$ serves as a "threshold line" because without distillation, our output error probability would be equivalent to our input error probability; thus, a distillation scheme improves the fidelity only when the $a(p)e(p)$ curve is below this line. From the plot, it is clear that Eastin's H-to-Toffoli distillation scheme is effective below an $|H\rangle$-error rate of $\approx 4.7\%$. Keep in mind that this is assuming only errors on the $|H\rangle$ magic states and that all Clifford gates are perfect—an assumption we aim to relax by defining EECs as described in the next section.

3 Constructing Effective Error Channels

Now that we have successfully assembled the logical circuit, we can move forward with the task of constructing effective error channels (EECs) for our circuit components. As we plan to perform rounds of QEC between each injection sub-circuit and transversal CNOT, it makes sense to construct EECs for each of these two components. This allows us to only consider the probability that the components yield *uncorrectable* errors, since correctable errors would be detected and corrected by our QEC cycles.

Thus, we decided to construct two different effective error channels for our circuit—one for each injection sub-circuit, and one for each transversal CNOT. To do this, we turn to quantum process tomography, a procedure used to reconstruct full quantum error channels describing the noise of components under arbitrary input states.

Importantly, when performing process tomography to construct these EECs, we make two key assumptions:

1. Our stabilizer measurements are perfect (thus requiring only one round of QEC for each cycle).
2. All initial states are perfect (prepared without noise).[7]

These assumptions are both fairly straightforward to relax (and can be done in future work), but for now they make the construction of our EECs significantly easier.

The results of our EEC construction at different physical error rates are displayed in Figure 5. Here, we can see how the outputted logical error rate of each component (the sum of all error cases in the error channel) scales with the physical Clifford error rate within the component. For our CNOT EEC (Figure 5b), we observe quadratic scaling, which makes sense as all weight-one errors are detected and

[7] Note that with this latter assumption, we also assume that $|H\rangle_L$ is perfect when constructing our EECs (the Cliffords involved in the $|H\rangle$-encoding circuit taken from [4] are noiseless). Importantly, this does not mean that $|H\rangle_L$ is assumed to be perfect in our final, logical-level simulation—in fact, we consider the $|H\rangle$-error probability to be a separate parameter altogether (as this non-Clifford error will occur more frequently than the Clifford errors in our encoding), and simulate for a range of $|H\rangle$-error rates (see Figure 6).

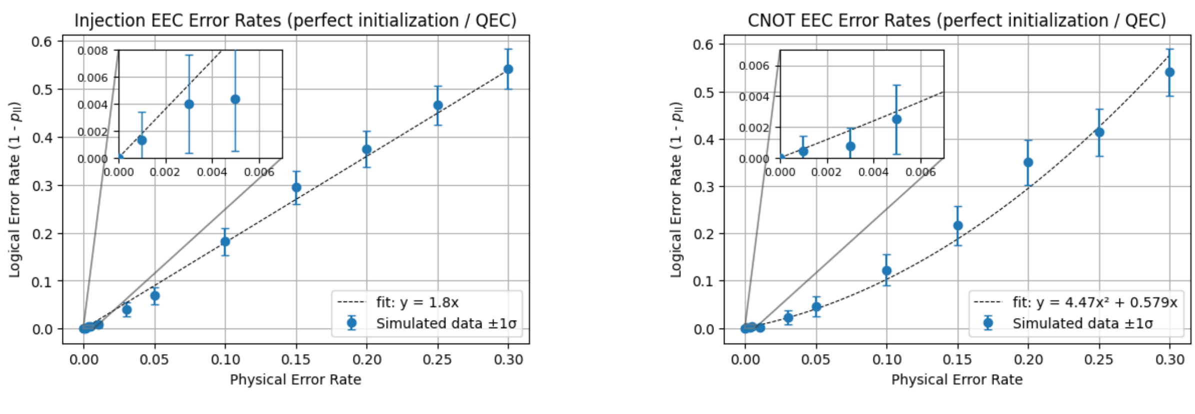

(a) Injection EEC

(b) CNOT EEC

Figure 5: Plots of each component's output logical error rate vs the physical Clifford error rate within the component. Each inset displays the low physical error rates of relevance to our distillation scheme.

corrected by our QEC cycles. For our Injection EEC (Figure 5a), we observe linear scaling, which is almost certainly an artifact of how this EEC was constructed—error correction was performed on only the data qubit and not the ancilla. If revised to perform QEC on the ancilla qubit as well (currently in progress), we would expect quadratic scaling similar to the CNOT EEC.

The insets of Figure 5 highlight the low physical error rates of relevance to our distillation scheme. Higher physical error rates render the distillation scheme ineffective as their corresponding undetected error probability curves are entirely above the $y = x$ threshold line (as will be shown in Figure 6). We can see that with these low physical error rates, there are large error bars and significant uncertainty in the outputted logical error rates of the EEC—this is due to computational limitations in the number of shots that could be performed when constructing the EECs. We have identified several approaches to further increase the number of shots and thus improve the quality of these error channels; however, this implementation is left to future work.

4 Logical Circuit with EECs

After constructing our EECs for each component, we resimulated the logical circuit with these EECs. This allowed us to specify physical Clifford error rates (denoted q) and determine how they affected the logical error rates, significantly extending the scope of the Eastin paper (where perfect Cliffords are assumed). Curves displaying the undetected error probability for different physical error rates, q, are displayed in Figure 6 below:

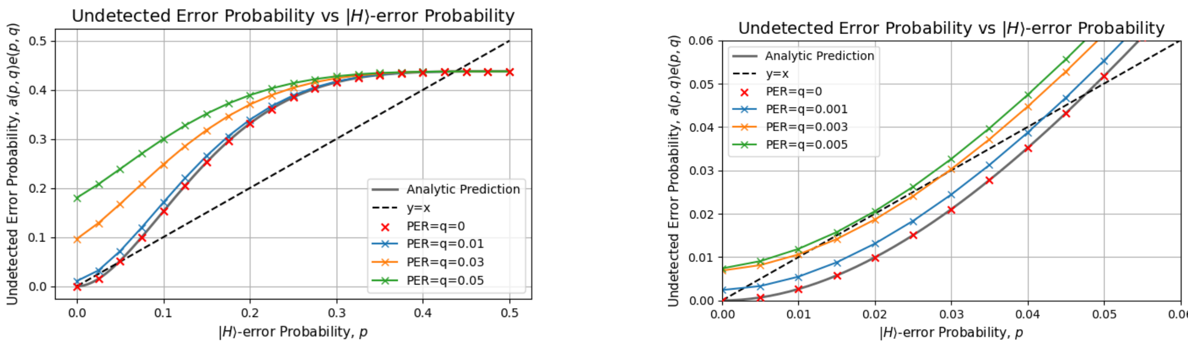

(a) $a(p, q)e(p, q)$ for large q

(b) $a(p, q)e(p, q)$ for small q

Figure 6: Undetected error probability for the H-to-Toffoli distillation circuit with additional curves representing non-zero physical Clifford error rates q. Figure 5a displays larger values of q, and Figure 5b displays smaller values of q at which the distillation scheme remains effective for a non-zero range of H-error probabilities. The simulated curves from Figure 5b use the EECs from the insets of Figure 5.

Both subfigures depict the same analytic curve from Figure 4b (perfect Cliffords), now with additional curves representing non-zero physical Clifford error rates q. From Figure 6a, it is clear that physical error rates on the order of 10^{-2} render our distillation scheme completely ineffective, as the curves remain above the $y = x$ threshold line for all $|H\rangle$-error probabilities. In Figure 6b, we can see that our distillation scheme remains effective for physical error rates on the order of 10^{-3}, given $|H\rangle$-error probabilities $p \sim 0.01 - 0.03$. Specifically, the distillation scheme appears to become ineffective at $q \approx 5 \times 10^{-3}$. That being said, as these curves correspond to simulations using the EECs displayed in the insets of Figures 5a and 5b, the statistical uncertainty of these low q simulations is substantial. Further optimization of the code used to construct the EECs (and thus more shots) would allow us to determine the regime where the distillation scheme remains effective with more confidence, and enable exploration of even smaller values of q.

5 Conclusion and Outlook

In this work, we have constructed effective error channels to model the Clifford noise associated with a Steane-encoded distillation circuit. Furthermore, we have extended the undetected error probability results for Eastin's H-to-Toffoli distillation scheme by simulating the circuit for non-zero physical Clifford error rates. While the current results have too much uncertainty to quantitatively characterize the circuit with high accuracy, this work serves as a proof of concept: constructing EECs and applying them at the logical level can help researchers characterize non-Clifford circuits that are difficult to simulate.

Recall that our primary aim is to determine if applying EECs to a logical circuit truly reproduces the error rates of the entire QEC-encoded version of the circuit. In this work, we have completed the first part of this task—we have successfully constructed EECs and applied them at the logical level. That being said, the task of simulating the entire QEC-encoded version of the circuit is left for future work. Originally, as our method for encoding Eastin's H-to-Toffoli distillation circuit with the Steane code requires 50 physical qubits,[8] we expected that the full simulation could be done using Qiskit's "Extended Stabilizer" simulator. This simulator scales computationally with the number of non-Clifford gates and can reportedly handle non-Clifford circuits up to 63 qubits [5]. Unfortunately, the Extended Stabilizer simulator does not currently support classically-conditioned operations, which are required for our injection sub-circuits (see Figure 1).[9] We have developed a few possible ideas on how to move forward with simulating the circuit given this issue, but a simple solution is to just limit the scope of our question—we could compare the logical circuit with EECs vs Steane-encoded circuit for just the injection sub-circuit, rather than the entire distillation scheme.

It is important to reiterate the future work that is more straightforward to implement. For one, our assumptions of perfect stabilizer measurements and perfect initializations can be relaxed when performing quantum process tomography to construct our EECs. Secondly, the code used to construct the EECs can be optimized, allowing us to take significantly more shots and achieve higher accuracy in our EEC error rates, especially for low physical Clifford error rates. These adjustments would provide a more complete quantitative understanding of the H-to-Toffoli distillation circuit.

[8] In fact, this particular circuit was chosen for our project as it is by far one of the simplest distillation schemes in the literature, requiring only 6 logical qubits (4 data qubits, 2 reused ancillas).

[9] An intuitive first workaround to this problem would be to somehow track this classically-conditioned operation through the circuit, applying it only during post-processing. Unfortunately, as our classically-conditioned $Y(\pm\frac{\pi}{2})$ rotation is Clifford and not Pauli, it is nontrivial to track it through our non-Clifford circuit. Modifying our injection circuit so that it contains a Pauli correction instead would allow us to use Qiskit's Extended Stabilizer simulator, a possible approach that we are currently exploring.

Acknowledgements

Firstly, I would like to thank the Global Quantum Leap program for making this life-changing research experience possible. In particular, I would like to thank Prof. Steve Koester for his incredible support throughout my participation in the program. I am additionally deeply grateful to all of the members of the RWTH Institut für Quanteninformation, who welcomed me into their truly one-of-a-kind community with open arms (including the post-lunch discussions over cake). I want to give a special thanks to Berat Yenilen, Dr. Manuel Rispler, and Dr. Markus Müller for all of the time they dedicated to my internship this summer and their incredible guidance throughout the entirety of the research project.

This project was funded by the National Science Foundation of the United States under grant number 2020174.

Appendix

A cleaned version of the code used in this project is available at `https://github.com/coalescion/EECs-for-Distillation.git`.

References

[1] Bryan Eastin. Distilling one-qubit magic states into Toffoli states. *Physical Review A*, 87(3):032321, March 2013. URL: `https://link.aps.org/doi/10.1103/PhysRevA.87.032321`, `doi:10.1103/PhysRevA.87.032321`.

[2] Daniel Gottesman. The heisenberg representation of quantum computers. arXiv preprint quant-ph/9807006, 1998.

[3] Bryan Eastin and Emanuel Knill. Restrictions on transversal encoded quantum gate sets. *Physical Review Letters*, 102(11):110502, March 2009. `doi:10.1103/PhysRevLett.102.110502`.

[4] Lukas Postler, Sascha Heußen, Ivan Pogorelov, Manuel Rispler, Thomas Feldker, Michael Meth, Christian D. Marciniak, Roman Stricker, Martin Ringbauer, Rainer Blatt, Philipp Schindler, Markus Müller, and Thomas Monz. Demonstration of fault-tolerant universal quantum gate operations. *Nature*, 605(7911):675–680, May 2022. `doi:10.1038/s41586-022-04721-1`.

[5] Qiskit Development Team. The extended stabilizer simulator, June 2025. Accessed: 2025-08-30. URL: `https://qiskit.github.io/qiskit-aer/tutorials/6_extended_stabilizer_tutorial.html`.