Completeness for Fault Equivalence of Clifford ZX Diagrams

Maximilian Rüsch, Aleks Kissinger, Benjamin Rodatz
University of Oxford, Oxford, UK

Two circuits are considered to be equivalent under noise if the effect of faults on one circuit is no worse than the effect of faults on the other circuit. We call this relationship *fault equivalence*. Fault equivalence offers a way to transform circuits while provably preserving their fault-tolerant properties, enabling a framework for fault-tolerant circuit synthesis and optimisation that is correct by construction. The ZX calculus offers a diagrammatic way to represent and reason about quantum circuits and is a useful tool for manipulating circuits while preserving fault equivalence. For this, the usual set of ZX rewrites has to be restricted to not only preserve the underlying linear map represented by the diagram, but also fault equivalence.

In this work, we provide a set of ZX rewrites that are sound and complete for fault equivalence of Clifford ZX diagrams. This means that any equivalence that can be derived using the proposed rules is certain to be correct, and any correct equivalence can be derived using only these rules. For this, we utilise diagrammatic constructions called fault gadgets to reason about arbitrary, possibly correlated Pauli faults in ZX diagrams. Fault gadgets allow us to separate the diagram into a fault-free part, which captures the noise-free behaviour of a diagram, and a noisy part that enumerates the effects of all possible faults. Using this, we provide a unique normal form for ZX diagrams under noise and show that any diagram can be brought into this normal form using our proposed rule set.

The full technical manuscript of this work can be found at https://arxiv.org/abs/2510.08477.

Introduction

Synthesis and optimisation of quantum circuits is a field of increasing importance, in particular as computations grow larger and significant time savings are possible. In the classical as well as in the quantum realm, equational *calculi* offer formal ways to rewrite and transform computations while preserving key properties such as their semantics. The usefulness and versatility of these calculi are usually determined by establishing two core properties: soundness and completeness. Soundness guarantees that any relationship derived using the equational rules of the calculus is true, and is therefore often viewed as a prerequisite to even using it. Conversely, completeness guarantees that any relationship that is true can be derived using only the fixed ruleset of the calculus. As such, completeness is a really useful theoretical guarantee as well: It makes the calculus as powerful as any other calculus focused on that relationship.

An emerging calculus for transformation of quantum computations is the ZX calculus [4] built on ZX diagrams representing computations. There exist rule sets for the ZX calculus that have been proven to be sound and complete w.r.t. the diagram semantics, i.e. preserving the underlying linear map [1, 11].

Recently [10] introduced a new equivalence relationship on Clifford ZX diagrams. This relationship not only considers the computation a ZX diagram represents but also its behaviour under noise. Two diagrams are considered *fault-equivalent* if for every undetectable fault on one diagram there exists an equivalent fault on the other diagram that is at least as likely. This guarantees that rewriting circuits does not introduce new, badly propagating faults (e.g. hook errors). Rodatz et al. [10] and follow-up work [6, 7, 9] show how fault equivalence can be used to create new, state-of-the-art fault-tolerant circuits. However, while these works provide sets of fault-equivalent (and thus sound) rewrites, completeness remained an open question.

In this work, we provide a set of rewrites that we prove sound and complete for fault equivalence of Clifford ZX diagrams. The rules are displayed in Appendix A. For this, we will construct a new unique

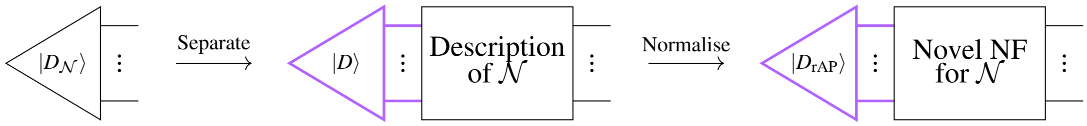

Figure 1: A description of the two distinct proof stages: first, we push all the faults to the boundary, then we bring the fault-free computation and the description of the noise model into unique normal forms. We use purple to indicate fault-free subdiagrams.

normal form for Clifford ZX diagrams under Pauli noise and show that the rewrite system is sufficiently expressive to reduce any diagram to the unique representative of its equivalence class. As rewrites are bidirectional, this means that if two diagrams D_1, D_2 are fault-equivalent, we can rewrite D_1 into the unique representative and then undo the rewrites in the other direction to reach D_2, yielding completeness.

Overview of the Proof

We consider Pauli noise on Clifford ZX diagrams. The noise model assumed on a diagram is expressed as a collection of independent error mechanisms (similar to [5]) called *atomic faults*. The set of all possible faults on a diagram is the group generated by the atomic faults. Faults are assigned a positive integer *weight* indicating how likely a fault is, with higher weight faults being treated as less likely.

To show that two fault-equivalent ZX diagrams D_1, D_2 with respective noise models $\mathcal{N}_1, \mathcal{N}_2$ can be rewritten into one another, we use *fault gadgets* [10] to diagrammatically capture the ZX diagram and noise model in a unified figure. That is, we construct diagrams $D_{1,\mathcal{N}_1} \mathrel{\widehat{=}} D_{2,\mathcal{N}_2}$ that respectively faithfully visualize D_1, D_2 and $\mathcal{N}_1, \mathcal{N}_2$. Since fault gadgets simplify representations of noise models using multi-edge faults to only single-edge Z fault noise models (so-called edge-flip noise), we conclude that these simpler noise models are *universal* for arbitrarily weighted Pauli noise. Furthermore, the simplicity of edge-flip noise enables the use of local equational reasoning to transform noise models, both in terms of amending atomic faults and amending the weight assignments. As we prove our equational ruleset to be sound, every such transformation implies fault equivalence to the original diagram and noise model.

For two diagrams $D_{1,\mathcal{N}_1}, D_{2,\mathcal{N}_2}$ that are in the same fault equivalence class, we propose a unique third member of that class and prove that both diagrams can be transformed into it using our ruleset. This procedure, visualized in Figure 1, can be divided into two distinct stages:

1. Push all atomic faults to the boundary of the base diagram, to **separate** the diagrammatic description of the noise model $\mathcal{N}_1$ and the base diagram describing the fault-free semantics of the computation.
2. For both parts of the diagram, **normalise** them individually to make them provably unique in their equivalence classes (for fault equivalence and semantic equivalence respectively).

Stage 1: Separation

To separate the semantic part of the diagram from the description of the noise, we observe that certain faults are indistinguishable and can therefore be treated as equivalent. For example, an X fault before the control of a CNOT is the same as an XX fault after the CNOT. We define:

Definition (Inconsequential faults, spacetime equivalence)**.** A fault F on a non-zero D is inconsequential if $D^F = D$, i.e. F happening does not change D. We say two faults F_1, F_2 are spacetime-equivalent if there exists some inconsequential fault F such that $F_1 F = F_2$.

We show that, if two faults F_1 and F_2 are spacetime-equivalent, we can, using our complete rule set, rewrite a fault gadget representing F_1 into a fault gadget representing F_2. As fault equivalence cannot distinguish between spacetime-equivalent faults, this operation is fault-equivalent.

To better understand faults in spacetime, we show that a fault's spacetime equivalence class is fully characterised by its commutation relationship with the Pauli webs [3] of the diagram. That is, two faults are spacetime-equivalent if and only if they anticommute with the same Pauli webs. Leveraging this observation, we show that all faults on some diagram D can be equivalently expressed as flipping some subset of webs not supported on the boundary (so-called *detecting regions*), along with Paulis on the boundary of D. Thus, $D_{1,\mathcal{N}_1}$ can be rewritten into a fault-equivalent diagram $D_{1,\mathcal{N}'_1}$ where faults only act on detecting regions and the boundary of D, allowing us to fully separate the faults from the rest of the diagram.

Stage 2: Normalisation

We now take on the two independent parts of the diagram and normalise them. For the part describing the fault-free semantics (thus being a fully idealised Clifford ZX diagram), we can reuse existing normal forms such as the *reduced AP form* [8]. Our ruleset in Figure 2 includes fully idealised variants of the standard Clifford ZX rules, which are proven to be complete for semantic equivalence [1, 2].

The part describing the noise model is normalised in four stages, based upon fault equivalence handling *undetectable* faults *up to spacetime equivalence*: (1) Remove all detectable faults while guaranteeing the group of generated undetectable faults stays the same, (2) diagrammatically construct at least one instance for every undetectable spacetime equivalence class, (3) reduce all members of each class to a single representative, keeping the minimum weight encountered in each class, and finally (4) canonically and uniquely order the representatives on the qubit world-lines. For each step, we guarantee soundness while also guaranteeing that the result is unique for all equivalent noise models. Recomposed with the diagram part describing the fault-free semantics, we transform the original diagram into the unique member of the fault equivalence class, and by reversing the rewrite chain for the second diagram, we obtain completeness. For example, we can verify one of the rewrites proposed in [10], for the sake of simplicity focusing only on X faults:

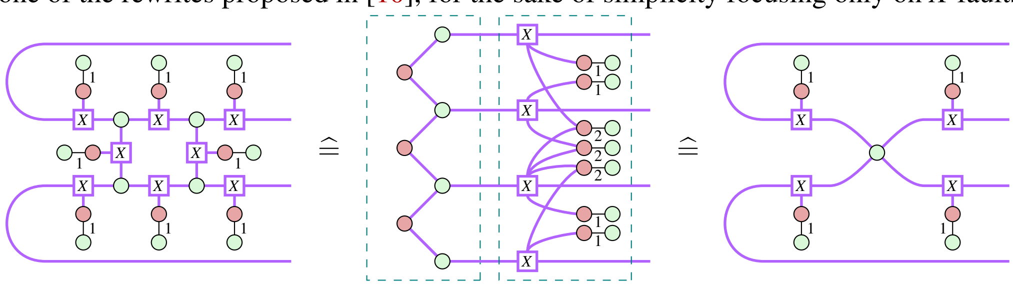

$D_{1,\mathcal{N}_1}$ $\quad\hat{=}\quad$ $D_{\mathrm{rAP},\mathcal{N}_{\mathrm{NF}}}$ $\quad\hat{=}\quad$ $D_{2,\mathcal{N}_2}$

Beyond fault equivalence, we consider soundness and completeness for three generalisations of the notion: (1) Limiting the guarantee of fault equivalence to a certain weight w, (2) dropping the symmetry requirement of fault equivalence, and (3) a combination of both. Based on the normal form for fault equivalence, we show that adding one rule each is enough to yield rulesets that are complete for generalisations (1) and (2), and adding both completes (3) (see the rules displayed in Figure 3).

Conclusion

This work provides a ruleset for the Clifford ZX calculus that is provably sound and complete for fault equivalence, as well as additional rules that extend the completeness to a few more general notions. As the proof is constructive, it enables an almost verbatim implementation for checking fault equivalence, which eases practical adaptation of the notion and integration of automated methods leveraging it[1]. Beyond these results, this work offers insights like spacetime equivalence and a more rigorous understanding of fault equivalence, extending and improving the set of tools associated with fault equivalence.

[1] An OSS implementation can be found at https://github.com/paritea/paritea

References

[1] Miriam Backens. "The ZX-calculus Is Complete for Stabilizer Quantum Mechanics". In: *New Journal of Physics* 16.9 (Sept. 17, 2014), page 093021. ISSN: 1367-2630. DOI: 10.1088/1367-2630/16/9/093021. arXiv: 1307.7025 [quant-ph].

[2] Miriam Backens, Simon Perdrix, and Quanlong Wang. "A Simplified Stabilizer ZX-calculus". In: *Electronic Proceedings in Theoretical Computer Science* 236 (Jan. 2017), pages 1–20. ISSN: 2075-2180. DOI: 10.4204/eptcs.236.1.

[3] Hector Bombin, Daniel Litinski, Naomi Nickerson, Fernando Pastawski, and Sam Roberts. "Unifying Flavors of Fault Tolerance with the ZX Calculus". In: *Quantum* 8 (June 18, 2024), page 1379. ISSN: 2521-327X. DOI: 10.22331/q-2024-06-18-1379. arXiv: 2303.08829 [quant-ph].

[4] Bob Coecke and Ross Duncan. "Interacting Quantum Observables: Categorical Algebra and Diagrammatics". In: *New Journal of Physics* 13.4 (Apr. 14, 2011), page 043016. ISSN: 1367-2630. DOI: 10.1088/1367-2630/13/4/043016. arXiv: 0906.4725 [quant-ph].

[5] Craig Gidney. "Stim: a fast stabilizer circuit simulator". In: *Quantum* 5 (July 2021), page 497. ISSN: 2521-327X. DOI: 10.22331/q-2021-07-06-497.

[6] Andrey Boris Khesin, Sarah Meng Li, Boldizsár Poór, Benjamin Rodatz, John van de Wetering, and Richie Yeung. *SpiderCat: Optimal Fault-Tolerant Cat State Preparation*. 2026. arXiv: 2603.05391 [quant-ph]. URL: https://arxiv.org/abs/2603.05391.

[7] Andrey Boris Khesin and Jonathan Z. Lu. *Mirror codes: High-threshold quantum LDPC codes beyond the CSS regime*. 2026. arXiv: 2603.05496 [quant-ph]. URL: https://arxiv.org/abs/2603.05496.

[8] Aleks Kissinger and John van de Wetering. *Picturing Quantum Software: An Introduction to the ZX-calculus and Quantum Compilation*. Preprint, 2024. URL: https://github.com/zxcalc/book.

[9] Boldizsár Poór, Benjamin Rodatz, and Aleks Kissinger. *Ultra Low Overhead Syndrome Extraction for the Steane code*. 2025. arXiv: 2511.13700 [quant-ph]. URL: https://arxiv.org/abs/2511.13700.

[10] Benjamin Rodatz, Boldizsár Poór, and Aleks Kissinger. *Fault Tolerance by Construction*. 2025. arXiv: 2506.17181 [quant-ph].

[11] Renaud Vilmart. "A Near-Minimal Axiomatisation of ZX-Calculus for Pure Qubit Quantum Mechanics". In: *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. 2019, pages 1–10. DOI: 10.1109/LICS.2019.8785765.

A The ruleset(s) complete for (generalisations of) fault equivalence

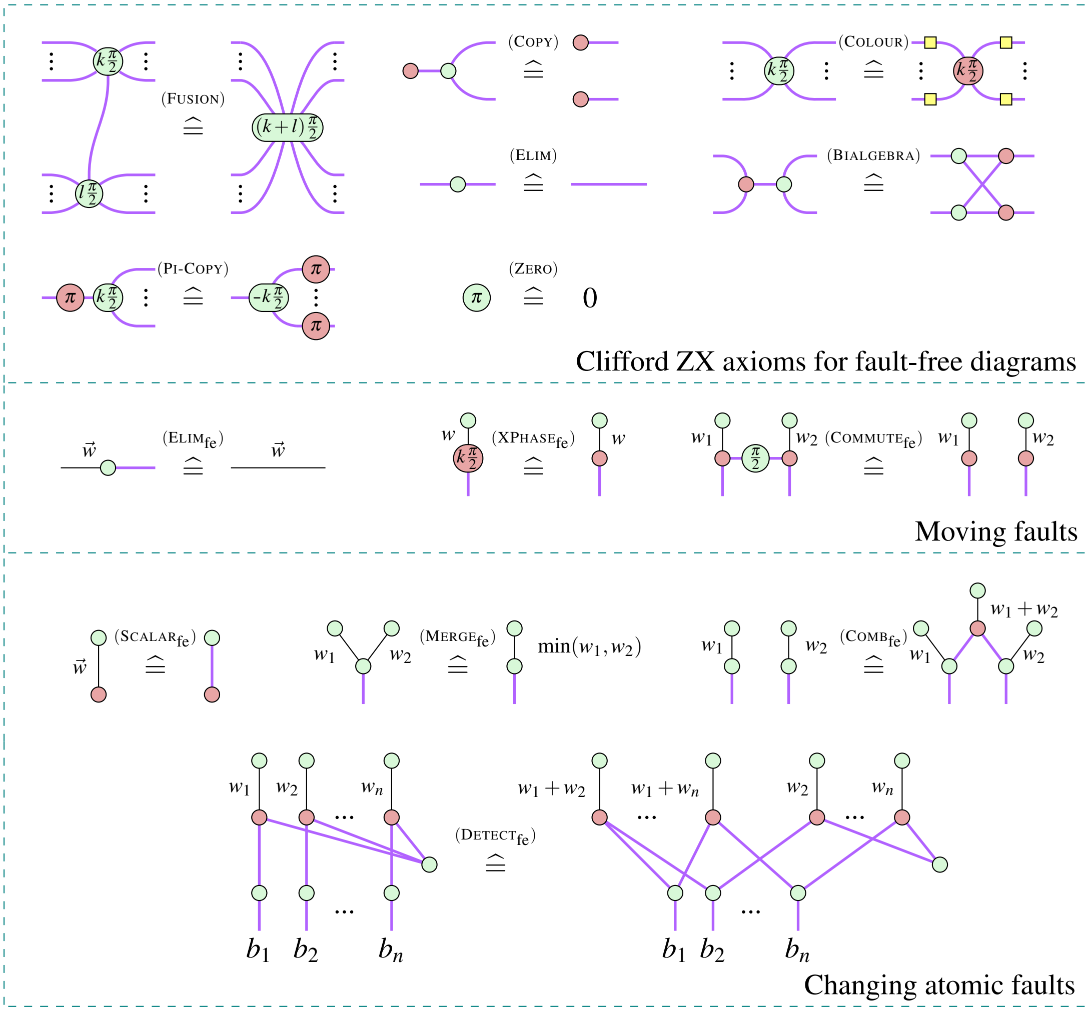

Figure 2: The ruleset that is complete for fault equivalence, divided into three groups: (i) the standard set from the Clifford ZX calculus that is complete for semantic equivalence, (ii) rules are required for transforming faults into spacetime-equivalent ones, (iii) rules for transforming noise models beyond spacetime equivalence.

Figure 3: Rules that, in addition to Figure 2, are required for a calculus that is complete for (a) w-fault equivalence and (b) fault boundedness. Include both to receive a ruleset complete for w-fault boundedness.