Linear Optical Fusion-Based Quantum Computation with Multi-Way Fusions

Sara Bartolucci

January 9, 2026

IMPERIAL COLLEGE LONDON

DEPARTMENT OF PHYSICS

IMPERIAL

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy in Physics of Imperial College London and the Diploma of Imperial College London.

Abstract

The realisation of a fault-tolerant universal quantum computer requires the ability to implement and control large scale quantum systems while maintaining physical noise levels sufficiently low to avoid errors in the computation. Fusion-Based Quantum Computation (FBQC) – a framework in which entangling fusion measurements are performed between constant-size entangled resource states – has emerged as a promising approach to achieve this, especially in the context of Linear Optical Quantum Computing (LOQC). Most FBQC schemes proposed so far use physical fusions between two qubits. In this thesis, we study the impact of using multi-way fusions, which act on more physical qubits simultaneously, on the error tolerance and resource overhead of linear optical FBQC architectures.

The main challenge of LOQC is that entangling operations are inherently non-deterministic, leading to large resource overheads. The success probability is known to decrease exponentially with the number of input qubits. We introduce novel fusion circuits which succeed at higher rates and/or entangle subsets of inputs upon failure. We present a framework to analyse the impact of noise on these circuits and describe the resulting errors as effective Pauli error channels on the input qubits.

We also examine the application of multi-way fusion circuits within a linear optical FBQC architecture. We find that they can be used to generate resource states with significantly reduced Pauli error rates and footprint compared to two-way fusions alone, in turn resulting in linear optical FBQC architectures with higher error tolerance and lower resource overheads. We also investigate their use as primitives of the computation, applied between different resource states, and find evidence that two-way fusions remain preferable in this context due to their lower erasure rate.

Acknowledgements

This thesis is the culmination of a journey that has lasted more than a decade. It has taken an immense number of small pieces all falling into the right place at the right time for me to get here, and I am infinitely grateful to everyone who has helped me put together this unique puzzle.

Thank you, first, to my supervisor, Terry, for believing in me from the very beginning, when I showed up at your door as a clueless undergrad, and for always looking for opportunities for me to be inspired and become a better scientist. I can still feel the excitement of that first summer spent learning about LOQC on the top floor of the EEE building, when I first saw what it meant to be a "researcher". Mercedes and Pete, thank you both for taking me under your wing and for sharing your infectious passion for your research with me. The chance you have given me to be part of the team working on one of the hardest technical challenges on the planet has forever changed the course of my life.

I have had the privilege of spending the better part of this ten-year journey at PsiQuantum, surrounded by unbelievably talented people. I am particularly grateful to the whole architecture team for their hard work of many years which laid very solid foundations for my research, and for generally being a great bunch to spend time with in and outside the office. Naomi, I consider myself extremely fortunate to have you as a role model for exceptional technical leadership and great camping skills. Patrick, Mihir and Sparrow, the depth and breadth of your knowledge and intuition never cease to amaze me and always give me something to aspire to. Working by your side has been an incredibly rewarding and fun experience.

I would not be the scientist, or the person, I am today without the friendships I have been fortunate to

build along the way. Your steady presence, from near or far, during both the hardest and best moments is the fuel that has kept me going. A special thank you goes to Abi for setting the ultimate example of what living life to the fullest and joyfully chasing one's dreams looks like.

But of course this journey would have never even started if it were not for those who were there all along, well before I even knew about the existence of quantum computing. To my parents, mamma e babbo, thank you for providing a strong launchpad for all my endeavours and a soft crashpad I know I can always land on. And to my brother, Lory, I would have never gone even half as far without your example and support, and for this I am forever grateful.

Finally, Justin, this thesis is only one of the many things in my life I could not have done without you. Thank you for inspiring me to be curious, to become a better version of myself every day, and, above all, for your love.

[...] to us, the world appears convergent, until we realize, with a jolt, that it isn't.

Brian Klaas, *Fluke*

Statement of Originality

The work presented in this thesis is the result of the research performed during my graduate studies in collaboration with my supervisors and other researchers who are adequately acknowledged. The thesis has been written solely by me and it presents my own work. Results from the literature mentioned in the thesis are appropriately referenced.

Statement of Author Contributions

Chapters 1 and 2 mostly contain introductory material from the literature to provide the necessary background and set the context for the work described in the rest of this thesis. Chapters 3–7 describe original work conducted in collaboration with other researchers at PsiQuantum. My contributions are summarised below.

Chapter 3 introduces novel linear optical multi-way fusion circuits. The constructive method to obtain new families of circuits by adding beamsplitters to the inputs of the n-GHZ analyzer was devised in collaboration with Patrick Birchall. Terry Rudolph proposed the ((1, 3), (2, 4)) 4-GHZ fusion, which inspired my analysis of depth-2 n-GHZ fusions as compositions of 4-Hadamards, as well as the 4-DFT-based 4-line fusion.

The error modelling frameworks described in Chapters 4 and 5 were developed by Chris Sparrow, Patrick Birchall and Fernando Pastawski. I used those to study Pauli errors from leading-order noise sources in multi-way fusions. The simulations and calculations to obtain fusion Pauli error models were performed by me with in-house software developed by the PsiQuantum team. The analysis of the results and investigation of Pauli error mitigation techniques were performed by me.

Chapters 6 and 7 describe the application of multi-way fusions to resource state generation and as computational primitives. This analysis was conducted in close collaboration with Mihir Pant. The software to find footprint-optimal schemes was developed by Tom Bell starting from my own implementation. I performed the calculations to compare schemes to generate specific resource states.

Copyright

Contents

List of Figures

List of Tables

Chapter 1

Fault-Tolerant Quantum Computation

This chapter introduces the core principles of fault-tolerant quantum computation, the broad context for the work described in this thesis, and gives an overview of different models which can be used to describe it. An introduction to the ZX-calculus language is also given, since this is extensively used in our analysis. The content of this chapter is a summary of relevant literature and an introduction to the notation used in this work.

In 1981, Richard Feynman gave a seminal talk [1] in which he posed the question "What kind of computer are we going to use to simulate physics?", in particular quantum physics, and answered it by proposing the use of "a machine of a different kind", a computer which is itself a quantum system. Since then, interest in quantum computing has grown very rapidly across academia, industry and government, due to the promise that such a machine would not only enable efficient simulation of physical systems, but also drastically speed up a number of classical algorithms, including database search [2] and factoring [3]. The range of problems expected to benefit from such capabilities is very wide, from the simulation of quantum chemistry processes for carbon capture [4], battery development [5] and drug discovery [6], to cryptography [7], machine learning [8], and more. Significant effort has gone into simplifying the implementation of quantum algorithms for these applications [9, 10], but it is still expected that any useful quantum computation will require a very large number

of operations. Therefore, it will be necessary to employ schemes to perform the computation in a fault-tolerant way, such that physical noise in the machine does not translate into logical errors.

This chapter describes the principles of fault-tolerant quantum computation and provides an overview of different models which can be used to understand and analyse it. In particular, the model of Fusion Based Quantum Computation (FBQC) is described, which constitutes the foundation for the work described in the rest of this thesis. An introduction to the ZX-calculus language is also provided, since this is used throughout this thesis.

1.1 The basics of quantum computation

1.1.1 Computing with quantum bits

The fundamental unit of information in classical computation is the *bit*, a system which can be in one of two states, commonly referred to as 0 and 1. Its analogue in quantum computation is the *qubit*, short for "quantum bit", a two-level quantum system whose state is usually described in terms of the two computational basis states $|0\rangle$ and $|1\rangle$. The key difference between a bit and a qubit is that the latter can be in a *superposition* state, i.e. a linear combination of the two basis states:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{1.1}$$

where α and β are complex coefficients such that $|\alpha|^2 + |\beta|^2 = 1$. Another way to express these is $\alpha = \sin\frac{\theta}{2}$ and $\beta = e^{i\phi}\cos\frac{\theta}{2}$, which allows to interpret the state of a single qubit as a vector within a sphere of radius 1, forming an angle θ with the $x-y$ plane and ϕ with the $x-z$ plane. This is known as the *Bloch sphere*. A single-qubit state corresponds to a 2D vector with rows corresponding to the

coefficients of the two basis states, i.e.,

$$|\psi\rangle = \begin{pmatrix} \alpha \\ \\ \beta \end{pmatrix}. \tag{1.2}$$

The state of n qubits is then a vector with shape $2^n \times 1$ in a 2^n-dimensional Hilbert space with entries corresponding to the coefficients of the 2^n states in the basis $\{|0\rangle, |1\rangle\}^{\otimes n}$.

A quantum computation is carried out by performing operations, *gates*, on the logical state of the input qubits. It is possible to define a *universal gate set*, a discrete set of gates sufficient for *universal quantum computation*, defined as the ability to perform any unitary transformation on the input qubits with arbitrary accuracy [11]. Multiple such gate sets exist. A common choice consists of the Hadamard, H, and the $\pi/8$, T, single-qubit gates, and the two-qubit controlled-NOT gate, CNOT [12], where

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ \\ 1 & -1 \end{pmatrix} \qquad T = \begin{pmatrix} 1 & 0 \\ \\ 0 & e^{i\pi/4} \end{pmatrix} \qquad \text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ \\ 0 & 1 & 0 & 0 \\ \\ 0 & 0 & 0 & 1 \\ \\ 0 & 0 & 1 & 0 \end{pmatrix}, \tag{1.3}$$

are the matrices which act on the single- and two-qubit state vectors.

Conjugating the target port of the CNOT with a Hadamard gate gives another common two-qubit gate, known as the controlled-phase gate, CZ:

$$\text{CZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ \\ 0 & 1 & 0 & 0 \\ \\ 0 & 0 & 1 & 0 \\ \\ 0 & 0 & 0 & -1 \end{pmatrix}. \tag{1.4}$$

This can be used to generate *entanglement*, a property unique to quantum systems whose joint state cannot be expressed as a product of individual states for each subsystem. An example of this is the following:

$$|00\rangle = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} \xrightarrow{H_1 \otimes H_2} \frac{1}{2} \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \xrightarrow{\text{CZ}} \frac{1}{2} \begin{pmatrix} 1 \\ 1 \\ 1 \\ -1 \end{pmatrix} = \frac{1}{2}\left(|00\rangle + |01\rangle + |10\rangle - |11\rangle\right) = \frac{1}{\sqrt{2}}\left(|+0\rangle + |-1\rangle\right), \tag{1.5}$$

where $|\pm\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle \pm |1\rangle\right)$. We see that this series of operations turns the input product state into a maximally-entangled state known as the two-qubit *linear cluster state*, equivalent to a Bell pair up to a Hadamard gate on one qubit. The ability to entangle qubits in the course of the computation is crucial for many quantum computing algorithms which result in exponential speedups over their classical counterparts, although the precise role that entanglement plays in enabling such resource reductions remains unclear [13, 14]. As we will see later, entangling operations are central to the framework of Fusion-Based Quantum Computation (FBQC) which underlies the work described in this thesis.

The final step in a quantum computation is the readout of the output qubit state after the necessary gates have been applied. This is done by performing measurements on the qubits in the computational basis, which returns a classical bit-string corresponding to the result of the computation.

1.1.2 Noisy quantum computation

Quantum error correction

When qubits, gates and measurements are realised in practice as physical quantum systems, they are subject to noise due to imperfections of the experimental setup and to the influence of the external environment. It is expected that any useful quantum algorithm to be executed on a quantum computer

will involve millions of qubit operations, meaning that error rates well below 10^{-6} would need to be achieved to avoid the risk of failure of the computation. Exceptional single-qubit gate fidelities have been obtained experimentally in some hardware platforms [15], but even the best demonstrations remain orders of magnitude away from the required error rates. The technique of Quantum Error Correction (QEC) provides a solution to this problem.

QEC is inspired by classical error correction, which is based on the idea of *encoding* the information by adding redundancy to it, and then performing a *recovery* procedure to read it out at the end of the protocol. However, there are two fundamental differences between quantum and classical systems which make a direct translation of classical error correction protocols to quantum computation impossible. First, the quantum state cannot be directly observed during the computation to determine whether an error has taken place, since this could alter the logical information. Moreover, the no-cloning theorem [16] states that it is not possible to perfectly copy an unknown quantum state, which means that any classical error correction protocol which relies on making multiple copies of the bit state is unsuitable for quantum systems.

In QEC, these challenges are overcome by unitarily encoding the state of a *logical qubit* into a *quantum code* consisting of multiple *physical qubits*. On this encoding, it is possible to perform measurements which do not disturb the quantum state of the logical qubit while allowing to identify the presence of errors, in a process referred to as *syndrome diagnosis*. *Decoding* is then performed to choose a correction operator to apply to try to restore the original quantum state.

In general, not all types of noise can be corrected by all quantum codes, but only those that satisfy the *Quantum Error Correction conditions* [11]:

Theorem 1. *Let P be a projector onto a codespace C, and $\mathcal{E}$ a quantum operation with operation elements $\{E_i\}$. A necessary and sufficient condition for the existence of an error-correction operation*

correcting $\mathcal{E}$ on C is:

$$PE_i^\dagger E_j P = \alpha_{ij} P \tag{1.6}$$

where α_{ij} are the complex elements of a Hermitian matrix.

A proof of this theorem is given in Ref. [11]. These conditions can be extended to any linear combination of errors E_i, meaning that the recovery operation can be used against not just a specific noise process, but against a whole class of them. This can be seen by considering two independent linear combinations of errors, $E = \sum_i \alpha_i E_i$ and $F = \sum_j \beta_j E_j$, and showing that eq. 1.6 implies that they also satisfy the quantum error correction conditions. This leads to the very powerful property of QEC that the continuous set of errors which can affect quantum states can be corrected by only correcting a discrete subset of them.

A common way to describe QEC codes is by using the notation $[[n,k,d]]$, where n is the number of *physical* qubits in the encoding, k the number of *logical* encoded qubits, and d is the *distance* of the code, i.e. the minimum number of single-qubit errors which can result in a logical error.

Stabilizer codes

Stabilizer codes are a widely-studied class of QEC codes which are particularly easy to represent and analyse within the framework of the *stabilizer formalism* [17]. In this picture, a quantum state $|\psi\rangle$ is described by the set of Pauli operators of which it is an eigenstate with eigenvalue $+1$. These operators are members of the Pauli group, $\mathcal{P}_1 = \{e^{i\frac{k\pi}{2}}I, e^{i\frac{k\pi}{2}}X, e^{i\frac{k\pi}{2}}Y, e^{i\frac{k\pi}{2}}Z\}$ with $k \in \{0,1,2,3\}$ for a single qubit, and $\mathcal{P}_n = \mathcal{P}_1^{\otimes n}$ for n qubits. Here I,X,Y,Z are the four 2×2 Pauli matrices. A set of states on n qubits can be defined as

$$V_S = \{|\psi\rangle : G_i|\psi\rangle = |\psi\rangle \ \forall G_i \in S\}, \tag{1.7}$$

where S is a subgroup of the Pauli group known as the *stabilizer group* of V_S. The elements of S are known as *stabilizer operators* and any product of them also belong to S. A more compact representation of this group can be given using its *generators*, a subset of its elements such that all the elements in S can be expressed as products of them. This is usually written as $S = \langle g_1, \ldots, g_l \rangle$, where l is at most $\log_2(|S|)$ and the number of generators is related to the dimension of the stabilized vector space. This compact description allows efficient classical simulation of any computation which consists of quantum states which are eigenstates of the Pauli matrices and their transformation through Pauli measurements and Clifford group operations, i.e., unitaries that evolve Pauli operators into Pauli operators. This was formally stated in the *Gottesman-Knill theorem* [18].

Unitary operations and measurements on quantum states can be described in this formalism through their effect on the stabilizer operators. Applying a unitary U to a state corresponds to conjugating all its stabilizer generators by it, such that

$$S \to S' = \langle U g_i U^\dagger \; \forall g_i \in S \rangle. \tag{1.8}$$

The action of a Pauli measurement g on the stabilizer group S depends on whether the operator commutes with all its generators. If it does, the state is an eigenstate of the operator, so it is not affected by the measurement, whose outcome is deterministic, and the stabilizer group is unchanged. However, if g anti-commutes with a subset of the stabilizer generators $g_1, \ldots, g_k$, it is possible to define a new set of generators $S = \langle g_1, g_1 g_2, \ldots, g_1 g_k, g_{k+1}, \ldots, g_l \rangle$, such that only the first element anticommutes with g. In this case the measurement outcome is random and the resulting stabilizer group becomes

$$S \to S' = \langle \pm g, g_1 g_2, \ldots, g_1 g_k, g_{k+1}, \ldots, g_l \rangle, \tag{1.9}$$

where the ± 1 phase is determined by the result of the measurement.

Within this formalism, a stabilizer code which encodes k logical qubits in n physical qubits can be

viewed as a subspace V_S stabilized by a group S with $n-k$ generators [19]. V_S is the *codespace*, a protected subspace of the full Hilbert space of the physical qubit state. For each encoded logical qubits, the code also has two logical operators, $\overline{X}$ and $\overline{Z}$, which act on its state. These are defined such that they anti-commute with each other and commute with all the elements of S, but they are not part of S themselves.

The way QEC is implemented with this class of codes is by performing measurements of the stabilizer operators, which do not alter the state of the code since they commute with the logical operators. For this reason, these operators are also referred to as *checks* of the code. In the ideal case, only $+1$ outcomes should be obtained, so any -1 outcome signals the presence of an error. Once the stabilizer measurement outcomes have been obtained, known as the *syndrome*, deducing what correction operation should be applied to restore the ideal state is not a straightforward task. Multiple decoding algorithms have been developed with different performance, both in terms of accuracy and computational efficiency, two properties which typically trade-off against each other [20].

The Shor code [21] is an example of a stabilizer code which can protect against any single-qubit error. The two logical qubit basis states are defined on nine physical qubits as follows:

$$|\overline{0}\rangle = \frac{1}{2}\left(|0\rangle^{\otimes 3} + |1\rangle^{\otimes 3}\right)^{\otimes 3}, \quad |\overline{1}\rangle = \frac{1}{2}\left(|0\rangle^{\otimes 3} - |1\rangle^{\otimes 3}\right)^{\otimes 3} \tag{1.10}$$

and an arbitrary logical qubit state is:

$$|\overline{\psi}\rangle = \alpha|\overline{0}\rangle + \beta|\overline{1}\rangle. \tag{1.11}$$

This is a $[[9,1,3]]$ code and therefore has $n-k=8$ stabilizer generators:

$$\begin{aligned} S_{\text{Shor}} = \langle & Z_1Z_2, Z_2Z_3, Z_4Z_5, Z_5Z_6, Z_7Z_8, Z_8Z_9, \\ & X_1X_2X_3X_4X_5X_6, X_4X_5X_6X_7X_8X_9 \rangle. \end{aligned} \tag{1.12}$$

The logical operators of this code are:

$$\begin{aligned} \overline{X}_{\text{Shor}} &= Z_1Z_2Z_3, \\ \overline{Z}_{\text{Shor}} &= X_1X_4X_7, \end{aligned} \tag{1.13}$$

although equivalent choices can be made by multiplying these by any of the elements in S_{Shor}.

If qubit i suffers an X flip, the measurement of the Z-type stabilizer generator(s) with support on it gives a -1 outcome. Similarly, if a Z error occurs, the X-type generator involving the flipped qubit is impacted. It can be seen that any single-qubit X error can be uniquely identified by the syndrome it produces. For instance, X_1 flips the Z_1Z_2 operator, X_3 the Z_2Z_3 operator, X_2 both of them, and so on. On the other hand, multiple Z errors can result in the same syndrome, such as Z_1, Z_2 and Z_3 all being consistent with only the $X_1X_2X_3X_4X_5X_6$ being flipped. This illustrates the fact that identifying the correction operation to perform is not always a trivial task and decoding algorithms are needed in general. In this case, applying any of the three Z operators returns the state to the codespace. By Theorem 1, then, the Shor code can correct any single qubit error which can be expressed as a linear combination of Pauli operators.

Single-qubit Pauli rotations are an important type of error which can occur, but not the only one. In general, the specifics of the underlying physical system are essential to determine the most suitable error model to describe it. *Erasure* is another kind of error which is common across many hardware implementations, corresponding to the occurrence of an error at a known position [22], for example as a result of a physical qubit being lost or of its state being otherwise leaked outside of the qubit

subspace. As we will see, this is in fact the dominant type of error in linear optics, the physical platform considered in this thesis. In this case, the state of the qubit is unknown and the error can be modelled by replacing the erased qubit with a fully depolarized qubit state, equivalent to a model where a random Pauli is applied to it. The fact that the location of the error is known makes QEC much more effective against erasure than other types of errors [19].

The Quantum Parity Code (QPC) [23] is a generalisation of the Shor code to $m_1 m_2$ physical qubits, obtained by replacing the two levels of $\otimes 3$ in eq. 1.10 with $\otimes m_1$ and $\otimes m_2$ respectively, i.e., its logical basis states are:

$$|\overline{0}\rangle = \frac{1}{2}\left(|0\rangle^{\otimes m_1} + |1\rangle^{\otimes m_1}\right)^{\otimes m_2}, \quad |\overline{1}\rangle = \frac{1}{2}\left(|0\rangle^{\otimes m_1} - |1\rangle^{\otimes m_1}\right)^{\otimes m_2}. \tag{1.14}$$

This code is also commonly referred to as the *generalised Shor code*. This family of codes belongs to a specific subclass of stabilizer codes known as *Calderbank-Shor-Steane* (CSS) codes [24, 25], whose construction is based on classical *linear codes*. They have the property that they can correct X and Z errors independently, and every stabilizer generator is either a product of only X or only Z single qubit operators. *Topological codes* [26] are another important group of stabilizer QEC codes. In these codes, the physical qubits can be thought of as forming a geometrically local lattice structure with two key properties: the stabilizer generators of the code consist of operators acting only on the local neighbourhood of each qubit, and the quantum information is stored in global degrees of freedom of the lattice. The *surface code* [27, 28] is the most widely-used topological code due to its simple structure as a square lattice in 2D, which makes it particularly easy to implement in a physical system, and its good error correction properties. The first experimental demonstration of logical error rate suppression with a QEC code was recently performed using this code [29].

Fault tolerance

QEC codes protect quantum information from noise while it is stored or transmitted. For fully fault-tolerant quantum computation to be possible, it is also necessary to prevent errors from accumulating during state preparation, measurement and gate application. One of the definitions of a fault-tolerant operation on an encoded quantum state is that it can be done in such a way that, given a probability p of an error occurring on one component, the probability of failure of the whole operation scales as $\mathcal{O}(p^2)$ [11]. *Transversal gates*, which act on a single physical qubit from each logical qubit, are one such type of operation, since they prevent the spread of errors within the QEC code. Unfortunately, it has been shown that not all operations in a universal gate set can be performed transversally [30]. Luckily, for gates which cannot be implemented in this way, more complex protocols have been proposed to perform them fault-tolerantly, for example with the introduction of low-error ancillary *magic states* distilled from a large number of noisy copies [31].

This means that, in principle, it is possible to perform universal quantum computation in a fault-tolerant manner. However, in order for this to be useful, it is also essential for fault tolerance not to result in an exponential resource overhead, which would negate any advantage compared to classical computation. The *threshold theorem* [32] is a fundamental result which establishes the possibility of realising a fault-tolerant quantum computer in practice:

Theorem 2 (**Threshold theorem**). *A quantum circuit containing $p(n)$ gates may be simulated with probability of error at most ε using $O(\mathrm{poly}(\log p(n)/\varepsilon)p(n))$ quantum gates on hardware whose gates fail with probability at most p_e, provided p_e is less than some constant threshold p_{th}.*

In other words, as long as the error rate on each operation is below a given threshold value, it is possible to perform a quantum computation of arbitrary size with polylogarithmic resource scaling. The threshold value is determined by the specifics of the QEC code, the error model for the underlying system, and the decoding algorithm used. Under a model where Pauli errors are independent and

identically distributed (i.i.d.), the 2D surface code has a threshold of $\sim 10.3\%$ [33]. In the case of erasure only, the threshold is 50% [34].

1.2 Models of quantum computation

A variety of physical systems can be used to implement universal quantum computation, provided that they meet a series of well-defined criteria [35, 36], among which is scalability. Several technologies have been found to satisfy these criteria and are concurrently being pursued across industry and academia for the realisation of a fault-tolerant quantum computer. A few examples are: superconducting qubits consisting of artificial atoms with two isolated energy levels, known as *transmons*, obtained from cooled circuits of capacitors and Josephson Junctions [37]; ion qubits which encode a quantum state in two of the energy levels of a trapped ion [38]; neutral atom (e.g., Rubidium) qubits whose state is encoded in two electronic levels of the atom and is controlled by lasers [39]; and spin qubits associated with single electrons in semiconductor (e.g., silicon) quantum dots [40]. As an alternative to these matter-based platforms, photons can also be turned into qubits by exploiting several of their degrees of freedom, such as polarization, spatial or time position [41]. This thesis focuses on quantum computation with single photons and linear optics, an approach which is described in more detail in the next chapter.

Each of these technologies presents advantages and disadvantages for the realisation of a large-scale fault-tolerant quantum computer, and criteria which are easy to meet for one system may be harder for others. For instance, two-qubit gates can be performed deterministically and with high fidelity on most of the aforementioned matter-based qubits, whereas they are harder to implement with linear optics, as we will see in the next chapter. On the other hand, measurement speeds for photonic qubits are significantly higher than for many types of matter-based qubit. The framework used to describe quantum computation so far is known as the *quantum circuit model*, wherein qubits are initialised,

	Quantum circuit model	Fusion-based quantum computing	Measurement-based quantum computing
	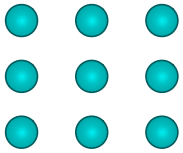	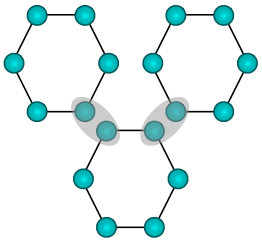	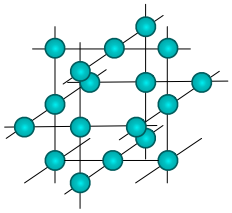
Objects	Static qubits	Constant-size entangled resource states	Large entangled cluster state
Operations	Non-destructive entangling gates and measurements	Possibly destructive entangling measurements (fusions)	Possibly destructive single-qubit measurements

Figure 1.1: Summary and comparison of the key elements of three models of quantum computation. Fusion-based quantum computing (FBQC) is the model underlying the work described in this thesis. (Figure adapted from Ref. [44])

acted on by gates, and eventually read out. While helpful to understand the fundamental principles of quantum computation, due to its resemblance to the structure of classical computation, some of the requirements of this model can pose challenges for certain physical platforms, e.g., by requiring repeated application of multi-qubit gates at runtime. Alternative paradigms for quantum computation have been developed over time which diverge from the common description of classical computation and are better tailored to the properties of different underlying physical platforms. Among these, two models have emerged as particularly advantageous frameworks for photonic quantum computation: Measurement-Based Quantum Computation (MBQC) [42] and Fusion-Based Quantum Computation (FBQC) [43], described in detail below. A summary and comparison of their key elements and of those of the quantum circuit model is given in fig. 1.1.

1.2.1 Measurement-Based Quantum Computation

MBQC was introduced by Raussendorf and Briegel in 2001 [42]. In this model, before the start of the computation, the input qubits are prepared in a highly entangled state, known as the *cluster state*. The computation is then carried out through single-qubit measurements and classical feedforward only.

Since the input resource is generated offline, its preparation can be attempted multiple times without affecting the runtime of the computation. This makes MBQC particularly relevant for physical quantum computing platforms for which single-qubit gates are substantially easier to implement than entangling multi-qubit operations, as in the case of linear optics.

The cluster state is a *graph state*, i.e. a state whose qubits can be represented as vertices connected by edges on a graph. The corresponding quantum state is the one obtained by assigning a $|+\rangle$ state to every vertex and applying CZ gates between pairs of connected vertices. The result is a state with stabilizer generators $X_i \prod_{j \in \mathcal{N}(i)} Z_j$, where $\mathcal{N}(i)$ is the neighbourhood of qubit i. For example, the two-qubit cluster state $\frac{1}{\sqrt{2}}(|+0\rangle + |-1\rangle)$ can be represented in this way as two vertices with an edge between them. For larger numbers of qubits, the cluster state can be thought of as a subgraph of a simple cubic lattice in d dimensions. The cluster state has been proven to be universal for quantum computation [45], and so have several other initial states [46, 47]. The use of MBQC for fault-tolerant quantum computation has been widely studied using several types of cluster states which have been shown to be robust to noise [48–51].

The entanglement in the cluster state is the resource consumed through single-qubit measurements during the computation. Initially the structure of the cluster state is modified through single-qubit Pauli Z measurements as required by the algorithm to be implemented. Then the computation proceeds as a sequence of adaptive single-qubit measurements, so that the results of previous stages can be used to choose the basis of the following. The final round of measurements is performed to read out the result of the computation. Throughout the computation, the entanglement in the cluster state is used to perform *teleportation* [52] of the quantum information across it while the desired gates are applied to it. In this way, MBQC does away with the need to perform entangling multi-qubit operations at runtime or to repeatedly apply a gate to a fixed physical system. However, this comes at the cost of a different set of challenges, particularly related to the generation and storage of a large, highly entangled input quantum state, whose size increases with that of the computation. Generating a

cluster state large enough to run useful algorithms and keeping it from decohering in the time needed to complete the computation is an arduous task for most physical platforms.

1.2.2 Fusion-Based Quantum Computation

FBQC, introduced in 2021 [43], avoids the complications associated with offline preparation of a large cluster state by replacing it with a *fusion network*, which specifies an entanglement structure to be implemented at runtime by performing multi-qubit entangling projective measurements, *fusions*, between small entangled *resource states*. A graphical representation of one of the fusion networks introduced in the original FBQC proposal is shown in fig. 1.2(a).

Resource states are typically chosen to be stabilizer states, and one of their key properties is that they are constant in size. A larger computation requires generating more resource states, but not larger ones, to realise a bigger fusion network. This implies that, differently from the cluster state of MBQC, the number of resources and operations needed to produce a resource state is fixed, which bounds both its cost and the amount of noise that can affect the qubits.

Fusions are measurements which project onto entangled stabilizer states. The simplest example is a Bell State Measurement (BSM) which measures the operators XX and ZZ on the two input qubits. When applied between resource states as prescribed by the fusion network, fusions create long-range entanglement across it which is used as a resource for the quantum computation. While the global structure of the fusion network is reminiscent of the cluster state, the key difference between this approach and MBQC is that here a large entangled state is never required to be present at once. Instead, the entanglement in the resource states, together with that introduced by the fusion measurements between them, generates the necessary long-range quantum correlations in the course of the computation as the qubits are consumed. In this approach, fusions can be performed sequentially and a resource state does not need to be produced until its first qubit has to be fused, thus avoiding a long idling time during which errors could occur.

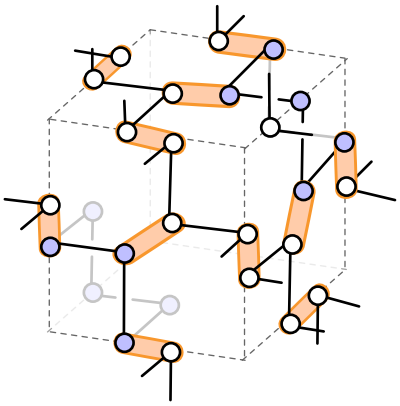

(a) Fusion network

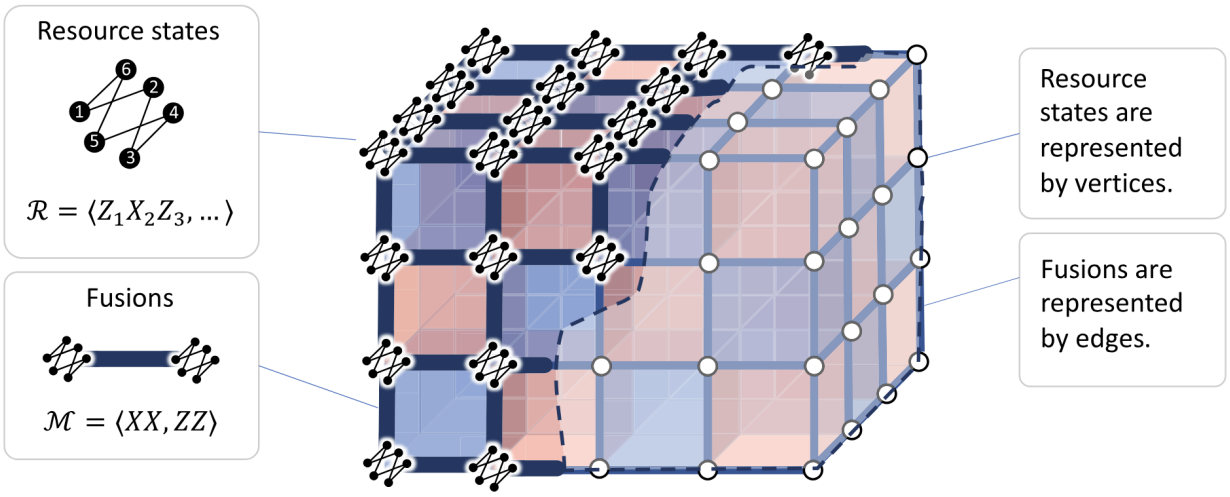

(b) Fusion graph

Figure 1.2: Descriptions of one of the FBQC schemes introduced in Ref. [43], which uses six-ring resource states and two-qubit fusions. (a) Unit cell of the fusion network of the scheme. The resource states are represented as six-qubit graphs placed at opposite corners of the cubic cell, and fusions between qubits of different resource states are shown in orange. (figure from Ref. [43], Open Access) (b) Fusion graph representation of the same scheme, which highlights its fault tolerance properties. Resource states are placed at the vertices of the graph, and edges are fusions between them. The parity checks of the scheme can be easily identified as cells of the graph. (figure from Ref. [53], Open Access)

Fault tolerance is achieved by choosing a fusion network whose structure corresponds to a QEC code, such that the outcomes of the fusion measurements can be used to evaluate its check operators. There are two stabilizer groups associated with the fusion network: a group R of all the stabilizer operators of the resource states, and a group F of fusion operators. When fusions are applied between resource states, the result is a non-trivial stabilizer check group $C := R \cap F$, referred to as the *surviving stabilizer group*. This is the subset of R whose elements commute with the elements of F, or, equivalently, the

subgroup of R which can be obtained from the fusion measurements. The fusion outcomes can be compared to the expected values of the resource state stabilizers, allowing to identify the presence of errors. A simple example is explicitly worked through in Supplementary note IVB of Ref. [43]. The fusion network is typically chosen to have a geometrically local structure to implement topological QEC and fault tolerance. Logical qubits are encoded in patches of a QEC code, *logical blocks*, and operations are applied to them by modifying their *topological features*, such as boundaries, through changes in the fusion measurement basis [53]. Schemes based on the surface code can be described as *fusion complexes* [54] and represented as *fusion graphs* with resource states at the vertices and fusions along the edges [53]. The checks of the scheme correspond to the cells of the graph. The fusion graph representation of the six-ring FBQC scheme from Ref. [43] is shown in fig. 1.2(b).

One of the main advantages of FBQC is the flexibility it provides. A given global entanglement structure can be realised with many fusion networks composed of different resource states and fusions. Moreover, the locations of resource states and fusions in space and time are also not fixed by the fusion network. The result is a variety of possible schemes to implement fault-tolerant quantum computation, which can be optimised according to the primitives and constraints of the underlying physical system. This framework is particularly well-suited for hardware platforms with native multi-qubit projective measurements and resource state generation, among which is photonics. As we will see in the next chapter, the use of FBQC has enabled significant improvements in the performance of linear optical quantum computation.

1.3 The ZX-calculus

The description of different models of fault-tolerant quantum computation may appear to suggest that these are fundamentally distinct approaches. While they certainly present significant differences, from the primitives they are based on to the way these are used for computation, it is nevertheless possible

to represent and analyse all these paradigms using a common framework based on the diagrammatic language of the ZX-calculus [55]. Introduced in Refs. [56, 57], the ZX-calculus has emerged as a powerful tool in various fields of quantum computing research, ranging from the analysis of computational models to quantum circuit optimisation, to the exploration of quantum error correction and fault tolerance [55, 58–60]. As will be clear in later chapters, this formalism is particularly useful for the investigation of various aspects of linear optical FBQC architectures.

The ZX-calculus is a graphical language consisting of a set of *diagrams*, which represent quantum systems, accompanied by *rules* used to describe their properties. The original proposal showed that this language is *universal*, meaning that it can capture any linear map between qubits, and *sound*, such that any equality derived in it can also be shown within other representations. While it is known that the ZX-calculus in its original form is not *complete*, new axiomisations were proposed to make the unrestricted language complete [61, 62]. In this thesis, we are principally interested in *stabilizer ZX-calculus*, a fragment of the language whose completeness was shown in [63]. Here we give a brief overview of the basic elements and rules of the language which are most relevant to the work described in the rest of this thesis, related to the design and analysis of linear optical FBQC architectures. A more comprehensive description of the ZX-calculus can be found in existing reviews such as [64–66].

1.3.1 The building blocks

The fundamental units of ZX-calculus diagrams are so-called *spiders*, round nodes with a phase α and with legs (or ports) representing input and output qubits of a linear map. There are two types of spiders differentiated by the colour of the nodes. Green (or white) spiders are also referred to as Z-spiders because they represent linear maps in the computational basis. A green spider with n input and m output ports corresponds to the map $|0\rangle^{\otimes m}\langle 0|^{\otimes n} + e^{i\alpha}|1\rangle^{\otimes m}\langle 1|^{\otimes n}$. Similarly, red (or grey) spiders are X-spiders as they represent linear maps in the Pauli X eigenbasis. A red spider with n input and m output ports corresponds to the map $|+\rangle^{\otimes m}\langle +|^{\otimes n} + e^{i\alpha}|-\rangle^{\otimes m}\langle -|^{\otimes n}$. These spiders can

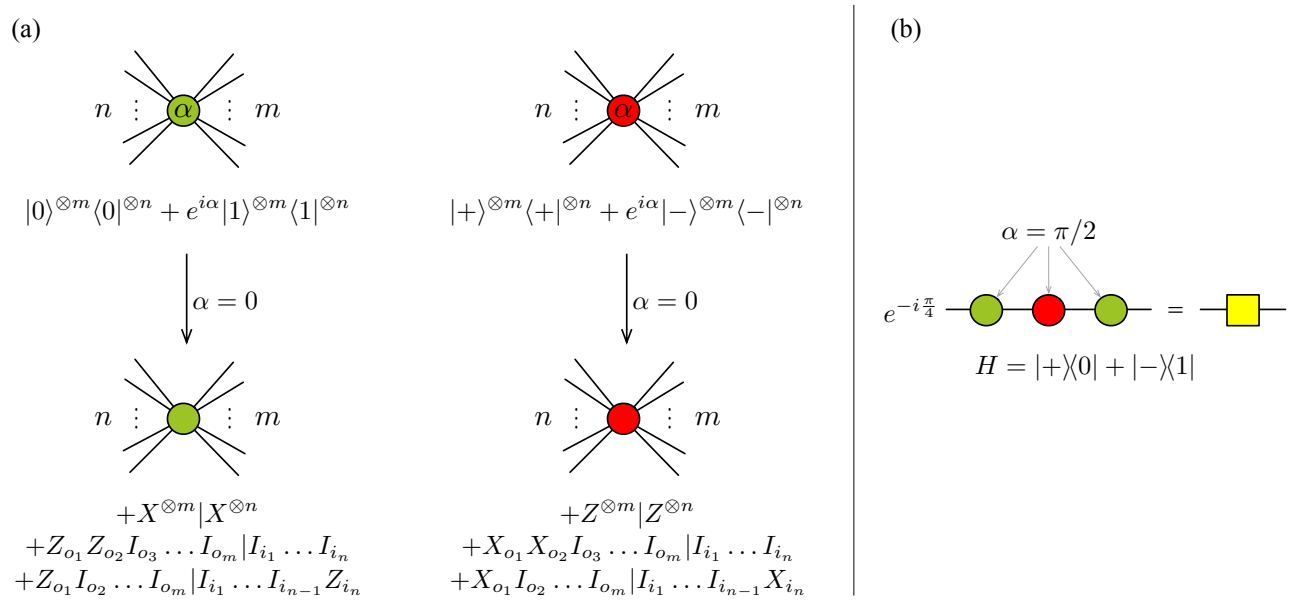

Figure 1.3: The fundamental building blocks of the ZX-calculus. (a) Green and red *spiders*, nodes with input (left) and output (right) legs and a phase parameter α, represent linear maps in the Z and X basis respectively. When $\alpha = 0$, we omit it from the diagrams, and the green (red) spiders have associated stabilizer generators consisting of X (Z) operators on all legs and pairwise Z (X). The stabilizer operators are expressed as $+P_{o_1}\ldots P_{o_m}|P_{i_1}\ldots P_{i_n}$, where P_j is a Pauli operator on qubit j. Operators to the left of the bar act on output qubits, and those on the right on the inputs. (b) A Hadamard unitary can be constructed as a chain of green and red spiders with $\alpha = \pi/2$, and it is represented as a yellow square node.

be used to represent multiple types of quantum objects and operations. A spider with only output legs is a pure quantum state, while one with only input legs is a measurement operator. Any single qubit unitary can be constructed as a chain of green and red spiders (Z and X rotations). Among these, the Hadamard gate has its own special notation as a yellow (or white) square with one input and one output. Refs. [55, 65, 66] also introduced ways to use spiders to represent *quantum instruments* which include classical information. Fig. 1.3 provides a graphical summary of the main building blocks.

In general, the phases associated to the spiders are such that $0 \leq \alpha < 2\pi$. In the special case that only multiples of $\pi/2$ are used, each spider can be associated with a set of stabilizer generators: a phaseless green (red) spider with n input and m output legs has one nm-body stabilizer generator $+X^{\otimes nm}$ ($+Z^{\otimes nm}$) and $(nm-1)$ two-body $+ZZ$ ($+XX$) on pairs of legs. The same holds for spiders with $\alpha = \pi$, except that the sign of the nm-body stabilizer generator is flipped. This restriction to stabilizer states is sufficient for our purposes of using the ZX-calculus for the analysis of FBQC [55],

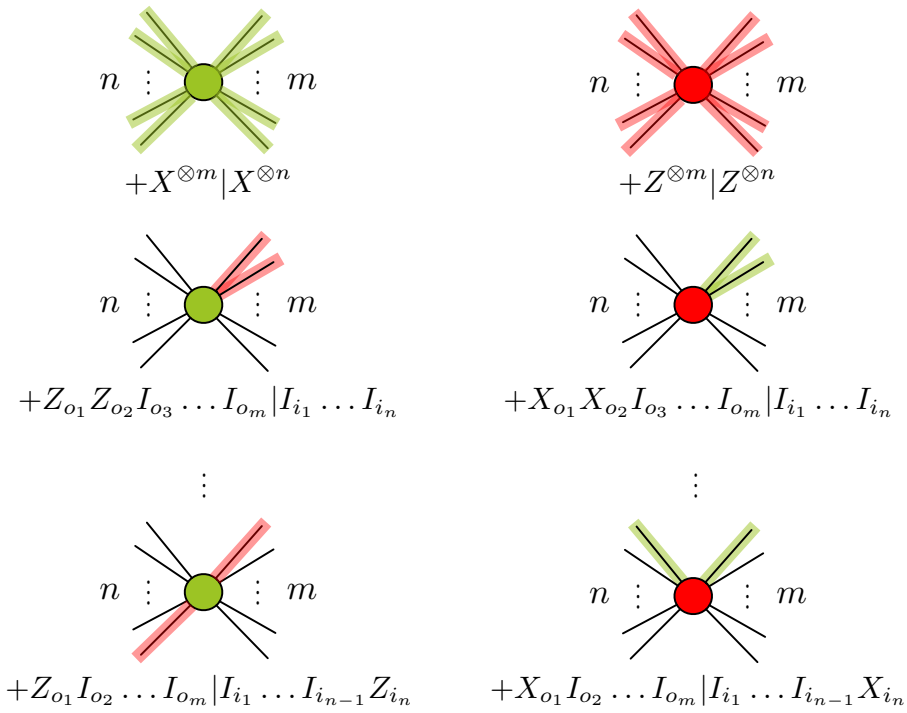

Figure 1.4: Representation of the stabilizer generators of green and red spiders as Pauli webs. Green (red) highlighting of an edge represents an X (Z) operator, and the total operator is the product of all the highlights.

and we usually assume $\alpha = 0$ unless otherwise specified. This is due to the property of *Pauli frame symmetry* of ZX-diagrams, thanks to which $\alpha = \pi$ phases on the spiders do not alter the structure of a network but only change the sign of the corresponding operators. Therefore, when analyzing a ZX-diagram all such phases can be ignored and assumed to be tracked separately. In some cases it is useful to keep track of which Pauli operators act on input or output ports, so we use the notation $\pm P_{o_1}\ldots P_{o_m}|P_{i_1}\ldots P_{i_n}$ for the stabilizer operators, where the P_{o_i} on the left of the bar act on each of the output legs, and the P_{i_i} on the right on the inputs. The bar can be omitted when the orientation of the ports corresponding to each operator is clear from the context.

The stabilizer operators associated with a spider can be represented using so-called *Pauli webs* [55, 67]. In this notation, a stabilizer operator is represented by highlighting the ports in its support in green or red, corresponding to X or Z single-qubit operators. The Pauli webs for some of the stabilizer generators of green and red phaseless spiders are shown in fig. 1.4. This graphical representation is useful to identify the stabilizer operators of larger ZX-diagrams, obtained by connecting spiders to

$$\text{CNOT} = |00\rangle\langle 00| + |01\rangle\langle 01| + |11\rangle\langle 10| + |10\rangle\langle 11|$$

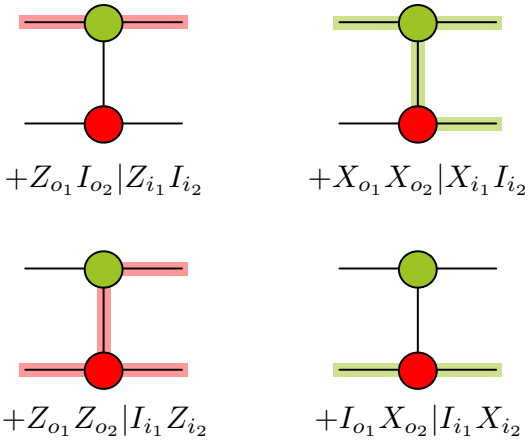

Figure 1.5: ZX-diagram and Pauli webs of the CNOT gate.

one another through their legs. For instance, connecting one green and one red three-legged spider through one port gives the two-qubit CNOT gate. This can be verified by considering the Pauli webs for the diagram, given by the subset of Pauli webs from the two spiders which are compatible, i.e., which do not have different operators on a shared leg. The resulting generating set of stabilizer operators for this two-spider diagram is shown in fig. 1.5.

1.3.2 Key rules

The ZX-calculus provides a set of transformation and equivalence rules which can be used to perform calculations on the diagrams. Here we list the subset of these rules which are most often used in the application of the ZX-calculus to linear optical FBQC. These are summarized graphically in fig.1.6.

1. *Spider fusion*: Two spiders of the same colour connected through one or more ports can be combined into a single spider with a number of legs equal to the total of the two minus the connecting ports. If the spiders have phases on them, the resulting spider has a phase which is the sum of the two.

2. *Identity*: A green or red spider with two legs and $\alpha = 0$ is equivalent to a connector without a node. This also means that red and green spiders of this kind are equivalent to each other.

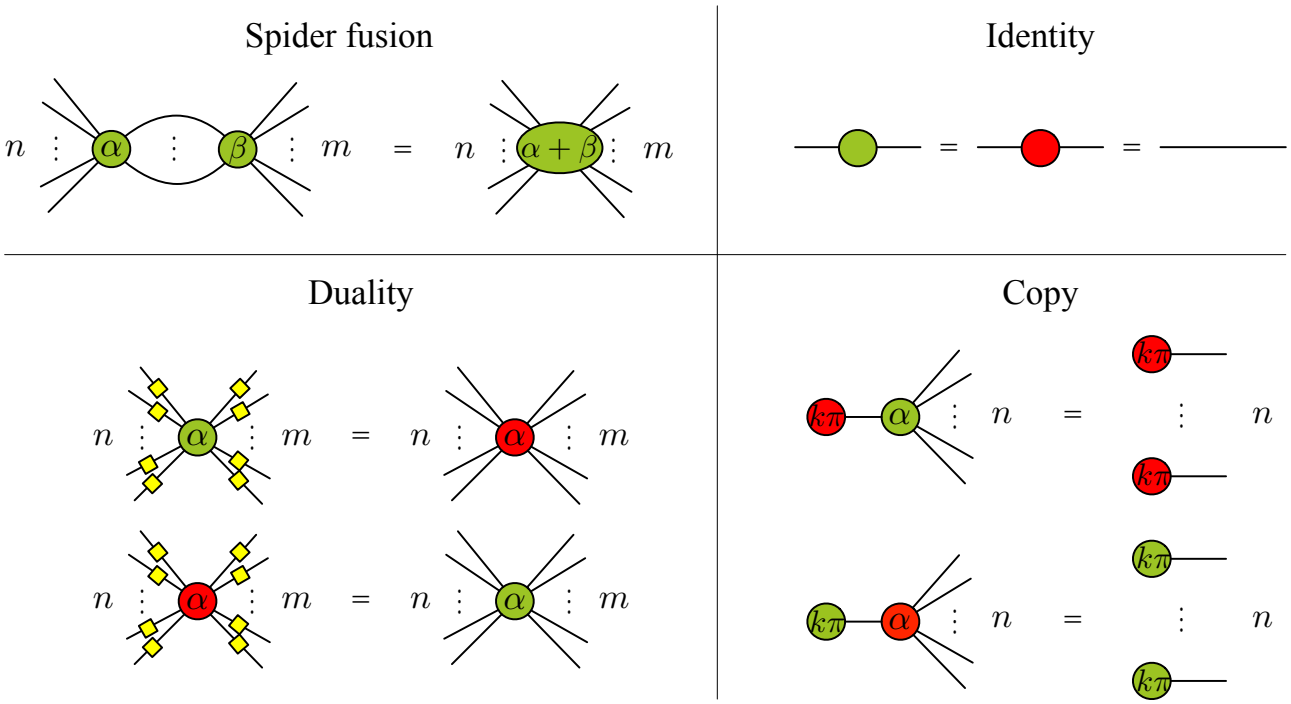

Figure 1.6: Key rules of the ZX-calculus used in the rest of this work.

3. *Duality*: A red spider can be turned into a green one by applying Hadamards to all its legs, and viceversa. This implies an important property of Hadamard nodes, namely that two of them cancel out when connected, resulting in a connection without a node.

4. *Copy*: A single-leg spider with phase $\alpha = k\pi$ connected to one leg of a spider of the opposite colour with any phase is "copied" onto the remaining legs, resulting in a diagram with $(n-1)$ spiders with phase $k\pi$.

Additional rules of the complete ZX-calculus include the *bi-algebra* and π-*commutation* rules.

The ZX-calculus spiders described above can be directly used to represent both primitives of FBQC, fusion measurements and resource states, and the rules provide ways of constructing and manipulating schemes. The next chapter will provide more details on this correspondence specifically in the context of linear optical FBQC.

1.4 Conclusion

Useful quantum algorithms will need a large-scale quantum computer to execute millions of operations on many qubits, requiring unrealistically low physical noise levels to prevent errors in the computation. The solution to this problem is provided by quantum error correction and fault tolerance, thanks to which many imperfect physical qubits can be used to encode and manipulate the state of a less noisy single logical qubit, provided that the physical error rates are below a certain threshold value.

Different models of quantum computation exist. Fusion-based quantum computation (FBQC) has emerged as a useful framework to describe and analyse fault-tolerant quantum computation, and particularly its implementation using physical systems for which multi-qubit entangling projective measurements, fusions, are native operations. In the next chapter we focus on the realisation of FBQC in one such platform: linear optics.

Chapter 2

Linear Optical Fusion-Based Quantum Computation

This chapter describes the implementation of fault-tolerant quantum computing in a specific physical platform: linear optics. An overview of the key elements of Linear Optical Quantum Computing architectures is provided, with a brief summary of the history of their development. The realisation of Fusion-Based Quantum Computation within this system is described in detail, and the evaluation of the performance of an architecture of this kind is discussed. The majority of the content of this chapter is a summary of relevant literature, except for Section 2.3, which introduces the framework and intuition which motivates the work presented in the rest of this thesis.

After reviewing the general principles and a few of the existing models of fault-tolerant quantum computation in a platform-agnostic way, we now turn our focus to its implementation using photonic qubits, the physical system considered in the rest of this thesis. In particular, the work in this thesis fits within the framework of Linear Optical Quantum Computing (LOQC), based on the interference and measurement of single photons, and more specifically to its application to the FBQC model, referred to as Linear Optical FBQC (LO-FBQC).

This section provides an overview of the key elements and methods of LOQC, focusing in particular on its implementation in integrated photonics, widely regarded as the most promising platform for the realisation of large-scale photonic systems [68–71]. A brief history of the evolution of LOQC

schemes is given, before focusing on the most recent proposal based on FBQC. The linear optical version of the primitives of this model is described, and an outline of the structure of a full LO-FBQC architecture and its properties is provided.

2.1 Computing with photons

Photons have a number of properties which make them appealing carriers of quantum information and fundamental units of a large-scale fault-tolerant quantum computer [70, 72]. First, they do not interact with the thermal environment or with each other, which limits decoherence of the quantum state due to noise and quantum cross-talk, while also relaxing cooling requirements. Moreover, photonic components can operate at fast clock rates, which reduces idling times and therefore the total error rate, and also enables fast operations. Single qubit gates can be implemented with high fidelity using simple optical components, and large numbers of photonic qubits can be manipulated and moved between different spatial locations with optical fibre, greatly simplifying the construction of large-scale quantum computing architectures.

2.1.1 Single photons as qubits

A common definition of a photonic qubit is based on the so-called *dual-rail encoding*, in which the two computational states correspond to the photon occupying one of two modes (or rails):

$$\begin{aligned} |0\rangle &= |10\rangle\rangle \\ |1\rangle &= |01\rangle\rangle, \end{aligned} \tag{2.1}$$

where the $|\cdot\rangle\rangle$ notation on the RHS represents the Fock state of the two qubit rails, while the $|\cdot\rangle$ notation is used for qubit states. Common choices of modes to implement this are polarisation, time-bins, or spatial rails (e.g., waveguides on a chip). In this thesis we primarily use the latter, but the results

presented can be translated to other encodings, for example using a simple passive transformation in the case of polarisation [41], or with the addition of an active element for time-bin.

Alternative photonic qubit encodings are possible, e.g., a single rail with vacuum or 1 photon in it, but the dual-rail encoding is the most widely studied due to its favourable properties. Because the photon number is fixed, it allows detection of errors which add or subtract photons to the rails, among which is photon loss, the dominant source of error in photonic systems [11]. Moreover, this encoding allows the relaxation of phase stability requirements. Consider a state of two dual-rail qubits:

$$|\psi\rangle = c_1|10\rangle\rangle_1|10\rangle\rangle_2 + c_2|10\rangle\rangle_1|01\rangle\rangle_2 + c_3|01\rangle\rangle_1|10\rangle\rangle_2 + c_4|01\rangle\rangle_1|01\rangle\rangle_2, \tag{2.2}$$

where the subscripts identify each qubit and the c_i are complex coefficients such that $\sum_i |c_i|^2 = 1$. If both rails of qubit 1 are subject to an arbitrary phase-shift ϕ, the resulting state is:

$$|\psi\rangle \rightarrow c_1 e^{i\phi}|10\rangle\rangle_1|10\rangle\rangle_2 + e^{i\phi}c_2|10\rangle\rangle_1|01\rangle\rangle_2 + e^{i\phi}c_3|01\rangle\rangle_1|10\rangle\rangle_2 + e^{i\phi}c_4|01\rangle\rangle_1|01\rangle\rangle_2 = e^{i\phi}|\psi\rangle, \tag{2.3}$$

which is equivalent to the original quantum state up to a global phase. This removes the need to control phases between rails of different qubits, greatly simplifying hardware requirements. Finally, arbitrary single qubit gates can be performed easily on dual-rail encoded qubits, since any unitary transformation on two rails can be implemented with a combination of linear optical devices, as described in more detail below.

2.1.2 Operations on photonic qubits

Among the possible implementations of large-scale photonic systems, photonic integrated circuits are commonly considered to be the leading approach [68–71]. Most of the components needed for photonic quantum computation can be manufactured in a compact manner in the same fabrication facilities that mass-produce chips for Complementary Metal–Oxide–Semiconductor (CMOS) tech-

nologies, which is a significant advantage for the reliable realisation of a large-scale fault-tolerant quantum computer. Over the last decade, circuits of this kind have been used for progressively more complex quantum information processing experiments [73–76], and the manufacturing of a complete set of hardware components for linear optical fault-tolerant quantum computing at a semiconductor fabrication facility has been demonstrated [77]. Therefore, this is the platform we consider for the implementation of the LO-FBQC architectures described in the rest of this thesis.

In silicon, the generation of the single photons relies on a probabilistic process known as Spontaneous Four-Wave Mixing (SFWM), which exploits the third-order non-linearity of the material [78]. This is analogous to the Spontaneous Parametric Down-Conversion (SPDC) process which takes place in materials with second-order non-linearity. When a pump laser is applied to these materials, the effect is described by the two-mode squeezing Hamiltonian, which generates the state:

$$\frac{1}{\cosh(r)} \sum_{n=0}^{\infty} (\tanh(r) e^{i\phi})^n |n\rangle\rangle_S |n\rangle\rangle_I, \tag{2.4}$$

where r is a parameter related to the intensity of the laser and $|n\rangle\rangle$ is the n-photon Fock state. The subscripts stand for *signal* and *idler*. When r is small, the dominant terms in the state are:

$$\frac{1}{\cosh(r)} |0\rangle\rangle_S |0\rangle\rangle_I + \frac{\tanh(r)}{\cosh(r)} e^{i\phi} |1\rangle\rangle_S |1\rangle\rangle_I, \tag{2.5}$$

which is a superposition of vacuum and a pair of single photons across the two modes. By sending one of the two modes to a detector, it is possible to herald the presence of a single photon in the other mode when a click is obtained. This setup is referred to as a Heralded Single Photon Source (HSPS), and it succeeds at generating a single photon with a probability which scales with the pump power up to a maximum of 25% [79]. While the photon generation probability is not very high, these sources can be fired at a high repetition rate, so the technique of *multiplexing* [80], or *muxing*, can be used to increase the success probability via repeated attempts, as described in more detail below.

A state of n photons in a mode can be written as

$$|n\rangle\rangle = \frac{(a^\dagger)^n}{\sqrt{n!}}|0\rangle\rangle, \tag{2.6}$$

where $a^\dagger$ is referred to as the bosonic *creation operator*, since its action on an n-photon state is

$$a^\dagger|n\rangle\rangle = \sqrt{n+1}|n+1\rangle\rangle. \tag{2.7}$$

The action of linear optical components on the photonic state can be expressed as a unitary *transfer matrix*, T, which transforms its creation operators as [41]:

$$a_i^\dagger \to \sum_j T_{ji} a_j^\dagger. \tag{2.8}$$

Once generated on chip, the dual-rail-encoded qubits are routed using waveguides, and they are transported between chips in optical fibre. Qubit manipulation is done on-chip with networks of beamsplitters and passive phase shifters. The former are realised on chip as *directional couplers*, structures which bring two waveguides close to each other so that their modes overlap. The transfer matrix of a beamsplitter acting on two rails is:

$$T_{\mathrm{BS}}(\theta) = \begin{pmatrix} \cos\theta & -i\sin\theta \\ -i\sin\theta & \cos\theta \end{pmatrix}, \tag{2.9}$$

where the parameter θ determines the *transmissivity*, $\cos^2\theta$, and *reflectivity*, $\sin^2\theta$, of the beamsplitter. $\theta = \pi/4$ corresponds to a 50:50 beamsplitter, which transmits and reflects equal amounts of light. When this kind of beamsplitter is applied to two modes containing one photon each, the resulting

state is:

$$|11\rangle\rangle = a_1^\dagger a_2^\dagger |00\rangle\rangle \tag{2.10}$$

$$\rightarrow \frac{1}{2}(a_1^\dagger + ia_2^\dagger)(a_1^\dagger - ia_2^\dagger)|00\rangle\rangle \tag{2.11}$$

$$= \frac{1}{2}\left[(a_1^\dagger)^2 + \cancel{a_1^\dagger a_2^\dagger} - \cancel{a_1^\dagger a_2^\dagger} + (a_2^\dagger)^2\right]|00\rangle\rangle \tag{2.12}$$

$$= \frac{1}{\sqrt{2}}(|20\rangle\rangle + |02\rangle\rangle). \tag{2.13}$$

The two terms corresponding to the two photons exiting the beamsplitters in different modes cancel out and *bunching* occurs, whereby both photons are found in the same output mode. This only occurs if the two input photons are perfectly indistinguishable in all their internal degrees of freedom. As will be explained in more detail later, this phenomenon, known as the Hong-Ou-Mandel (HOM) effect [81], is central to the generation of entanglement with linear optics, and therefore to LO-FBQC.

Passive phase shifts are imparted by varying the length of the waveguide on the photon path. A phase shift by an angle ϕ applied to input mode i is described by a diagonal transfer matrix where all elements are 1 except the i^{th} one:

$$T_{\text{P},i}(\phi) = \begin{pmatrix} 1 & \cdots & \cdots & \cdots & 0 \\ \vdots & \ddots & & & \vdots \\ \vdots & & e^{i\phi} & & \vdots \\ \vdots & & & \ddots & \vdots \\ 0 & \cdots & \cdots & \cdots & 1 \end{pmatrix}. \tag{2.14}$$

Any two-mode unitary can be implemented with beamsplitters and phase shifters [82], meaning that arbitrary single-qubit gates can be realised with these two elements. As an example, the Hadamard single-qubit gate, H, can be realised by applying a 50:50 beamsplitter with a $\pi/2$ phase shift on the

second input and output to the two rails of a qubit:

$$\begin{aligned} H &= T_{\mathrm{P},2}(\pi/2)T_{\mathrm{BS}}(\pi/4)T_{\mathrm{P},2}(\pi/2) \\ &= \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix} \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix} \\ &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \end{aligned} \tag{2.15}$$

A beamsplitter with variable transmissivity can be implemented as a Mach-Zehnder interferometer (MZI), where a variable phase shift is applied to one port between two 50:50 beamsplitter. Arbitrary unitaries on any number of modes can be built by connecting multiple of these [82, 83]. For this purpose, MZIs require the use of active phase shifters. Active phase shifters, whose phase can be varied as needed, are also typically necessary to implement adaptive measurement choices and multiplexing[1]. Currently, the favoured approach to implement active phase shifters requires the integration of materials with electro-optic effects on chip [85]. Recently, active phase shifters based on the Pockels effect in BTO have been shown to have low loss and to be compatible with the integrated photonics platform [77].

Multiplexing is a technique used to increase the probability of non-deterministic, heralded processes such as the generation of single photons [80]. The probabilistic operation is attempted multiple times and the outputs are sent into a switch network consisting of passive linear optical elements as well as active phase shifters. The heralding information is used to identify one or more successful inputs to route to the output, and the active phases are set accordingly so that the transfer matrix of the network implements the desired routing. Multiplexed single photon source designs have been proposed and experimentally realised [86–92]. In its simplest form, multiplexing allows the use of N copies of a

[1] It has been shown that in principle it is possible to perform photonic quantum computing without coherent switching, which preserves the quantum state [84]. However, the resulting architectures have very large resource overheads which make them unrealistic to realise in practice.

state generated with probability p to obtain a single copy of it with probability $p_{\mathrm{MUX}} = 1-(1-p)^N$. This strategy, known as N-to-1 multiplexing, is inherently inefficient, since on average Np copies of the state are generated but only one of them is used. In general, inefficiency can be quantified by a parameter $\gamma = N_{\mathrm{in}}/N_{\mathrm{out}}$, where N_{in} is the average number of available inputs and N_{out} the average number of outputs which can be used ($N_{\mathrm{out}} = 1$ in the case of N-to-1 multiplexing). This corresponds to the number of input resources needed per useful output, so $\gamma \geq 1$, where $\gamma = 1$ corresponds to perfectly efficient multiplexing. Strategies and switch network designs with improved efficiencies and hardware requirements have been proposed, most of which rely on the selection and routing of multiple input copies in what is known as N-to-M multiplexing [93, 94].

The worldline of every photon in an LOQC architecture ends with a detector, a component which absorbs the photon and turns its energy into an electrical signal [95]. Detectors are used to herald the outcome of operations throughout the computation, from single photon generation, qubit entanglement, and the readout of the outcome of the computation. In most cases, it is necessary to be able to resolve the number of photons detected, i.e., to apply a Fock state projector $|n\rangle\langle n|$. Detectors with this capability are referred to as Photon Number Resolving Detectors (PNRDs) and they are usually constructed by cascading multiple individual detectors. Superconducting Nanowire Single Photon Detectors (SNSPDs) are the leading detector technology for scalable implementation, thanks to their compatibility with integrated photonic materials and high operation speed [77, 96]. As the name suggests, these devices are made of a material which is superconducting below a critical temperature. The absorption of a photon causes the material to heat up and become resistive, resulting in a voltage which can be read out. These are the only components in a photonic quantum computer which need to operate at low temperatures, on the order of a few Kelvin.

In the rest of this thesis, we use the graphical representation in fig. 2.1 for the components of linear optical circuits acting on dual-rail qubits. We follow Ref. [97], in using two black nodes connected by a vertical line as shorthand for a Hadamard interferometer built out of two $\pi/2$ passive phase

(a) **(b)** **(c)** **(d)**

Figure 2.1: Graphical representation of the main components used in linear optical fusion circuits. (a) A Hadamard interferometer acting on two photon modes. (b) A crossing of photon modes. (c) A photon number resolving detector (PNRD) capable of discriminating the number of photons detected. (d) A passive phase shifter imparting a phase $e^{i\alpha}$. In this notation, lines represent individual rails of a linear optical circuit, not qubits as in the representation of standard quantum circuit diagrams.

shifters and a 50:50 beamsplitter, as seen in eq. 2.15. Often diagrams will include these Hadamards acting on non-adjacent rails. In practice, modes have to be brought together to perform the coupling, so the diagrams implicitly assume the presence of rail crossings. When represented explicitly, these are drawn as two intersecting mode lines as in fig. 2.1b. PNRDs are represented by the symbol in fig. 2.1c. Active or passive phase shifters are represented by white triangles with a phase value α, as in fig. 2.1d.

Single-qubit operations are straightforward to implement on dual-rail encoded qubits, since they correspond to two-mode unitary operations. As described above, these can be realised as simple networks of beamsplitters and phase shifters, and very high single-qubit gate fidelities have been demonstrated experimentally [72]. However, single-qubit gates alone are not universal for quantum computation. Because photons do not interact, implementing two-qubit entangling gates is a major challenge for the linear optical approach to quantum computing. In fact, it is impossible to maximally entangle polarisation-encoded qubits, and therefore dual rail qubits, using only linear optical operations, since a maximally entangled state cannot be factorised into linear combinations of creation operators [98]. This was thought to imply the impossibility of implementing universal quantum computation with linear optics, since non-linearity would be necessary to implement a CZ or CNOT operation needed to achieve a universal gate set. That was until it was shown that this non-linearity could be obtained by supplementing linear optics with measurement, launching the field of LOQC. In the next section, we review this landmark result and subsequent developments in the field of LOQC to date.

2.1.3 A brief history of Linear Optical Quantum Computation

The proposal of the Knill, Laflamme & Milburn (KLM) scheme in 2001 [99] was groundbreaking for the field of LOQC. It showed that scalable universal quantum computation is possible using only linear optical elements: single photon sources, beamsplitters, phase shifters and single photon detectors with classical feedforward. The key realisation was that *measurement-induced non-linearity* [100] can be used to implement a two-qubit CZ gate out of linear optical elements followed by detectors, completing a universal gate set. The operation succeeds probabilistically upon measurement of a specific detection pattern, but, crucially, it can be made to approach unit probability with a polynomial resource overhead. In practice, the use of entangled ancilla states and QEC codes are required, resulting in extremely high resource counts which make it a proof of principle more than a practical scheme [41, 101].

A number of proposals were subsequently put forward to reduce the resource overhead towards a feasible quantum computing architecture. Some improvements were initially obtained by considering variations of the KLM scheme [102–104], but the most signficant gains came from the introduction of protocols based on MBQC. In fact, this model is ideal for LOQC, since it allows to shift the impact of the probabilistic nature of entangling measurements to the offline generation of the cluster state, while the computation only involves deterministic single-qubit measurements.

Nielsen [105] proposed the first scheme of this kind, where KLM CZ gates were used to probabilistically add qubits to the cluster state. Browne & Rudolph [106] then introduced a different way to grow the cluster state using *fusion gates*, simple interferometers followed by detection which take two qubits as inputs and perform a parity check operation on them. There are two types of fusions, used for different purposes in the construction of the cluster state: the Type-I (T1) fusion outputs one qubit and is used to grow linear clusters, while the Type-II (T2) fusion consumes both inputs and is used to form two-dimensional cluster states. Both operations succeed with 50% probability

and a failure removes vertices and edges from the cluster state. A more detailed description of these operations is given in the next section. The simplicity of the fusion interferometers, compared to the CZ gate construction, and the efficiency of the scheme to grow the cluster state make the Browne & Rudolph scheme substantially cheaper than previous proposals in terms of both number of operations and resources consumed [101]. It was shown that scalable quantum computation is possible using this approach, but with threshold values of 3×10^{-3} for photon loss and 10^{-4} for depolarising noise, both of which are significantly lower than the expected error rates in realistic systems [107].

Another crucial advance towards the practical realisation of LOQC was the introduction of completely *ballistic* schemes, where every operation is only attempted once. Previous repeat-until-success strategies suffered from the significant drawback of requiring large amounts of fast active switching and photon routing, which would be challenging to implement and prone to high error rates. The first fully ballistic proposal came from Kieling, Rudolph & Eisert [108]. Their scheme accepts the probabilistic nature of fusions and builds an imperfect cluster state with missing vertices and edges. Using concepts from percolation theory, a renormalisation procedure is used to coarse-grain the cluster state and obtain a resource for MBQC. The cost of avoiding large amounts of switching and routing is that the resource overhead is higher compared to other repeat-until-success schemes, but later proposals showed that savings are possible within this paradigm while also improving the error tolerance [109–111].

The most recent development in LOQC is the design of architectures based on the FBQC model of quantum computation. In the next section we describe the key elements of LO-FBQC, including the implementation of the primitives of this model in linear optics and the main features of a physical architecture constructed with them.

2.2 FBQC with photons

Linear optics is a platform particularly well-suited for the implementation of FBQC, since it natively has access to multi-qubit projective entangling measurements: the fusions introduced by Browne & Rudolph. In fact, the development of this model for quantum computation was inspired by these operations, which it is named after. The first proposal of an LO-FBQC architecture outlined in Ref. [43] demonstrated the potential for this framework to lead to significant design simplifications and performance improvements over previous schemes. Avoiding the need to generate and store a growing entangled cluster state streamlines the physical architecture, reducing component depth and classical processing requirements, which, together with the flexibility to tailor the design of the fusion network to the underlying fault-tolerant scheme, results in high thresholds particularly against photon loss, the main source of noise in linear optics. This advantage has been demonstrated by following proposals for FBQC architectures with further performance improvements [112–117].

In this section, we describe the way in which the two fundamental primitives of FBQC, fusions and resource states, are realised with linear optics, and give an overview of the key elements of a physical architecture for an LO-FBQC architecture.

2.2.1 Two-qubit linear optical fusions

The linear optical two-qubit T1 and T2 fusions introduced by Browne & Rudolph [106] can be applied to spatially-encoded dual-rail qubits via simple interferometers of beamsplitters followed by detectors. Their operation relies on the HOM effect taking place when the input photons meet at a beamsplitter, which causes different detection patterns to be consistent with only a subset of the possible input qubit states. In some cases, this results in an entangling projection on the inputs, or a *success* outcome, and in others in a projection onto a product state, or a *failure* outcome. Therefore fusions are inherently probabilistic entangling operations. The type of measurement performed is heralded

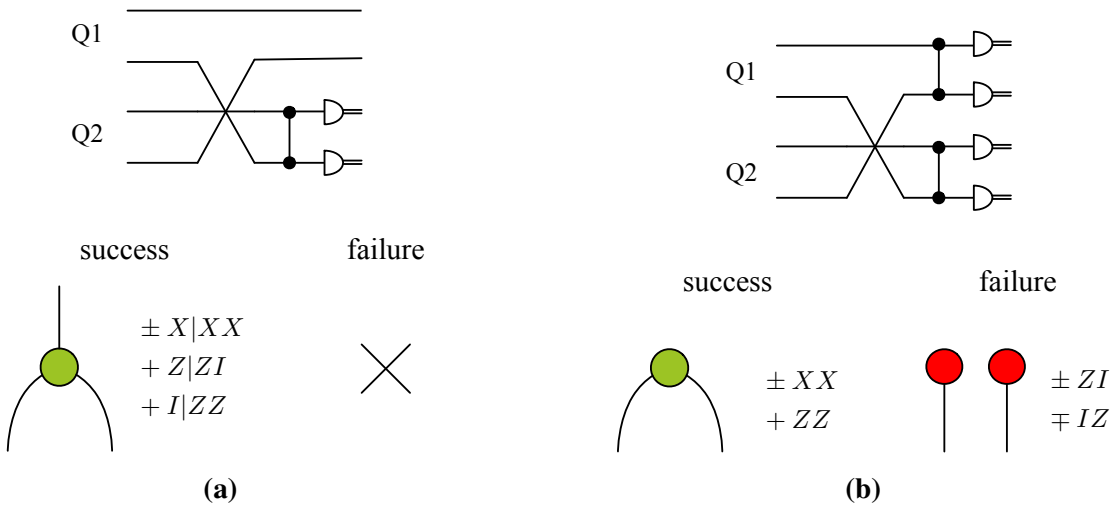

Figure 2.2: Circuit and ZX-diagram representations of the outcomes of linear optical two-way fusions. (a) T1 fusion, performed by interfering one mode from each input qubit at a beamsplitter and detecting the output. In case of success, a qubit is defined on the two output rails and the ZX-diagram is a three-legged green spider. When the fusion fails, the two input qubits are measured in Z, but the state of the output rails is outside of the qubit subspace, which cannot be represented in ZX-notation. (b) T2 fusion. The circuit is the same as the T1 fusion, but with an extra beamsplitter and detectors on the other pair of rails. In case of success a Bell state measurement is performed, represented by a two-legged green spider. When the fusion fails, the two inputs are measured in Z.

by the pattern observed at the detectors which follow the interferometer, called the *detection pattern* or *heralding pattern*.

The T1 fusion circuit interferes two rails from the input qubits at a beamsplitter and then detects them, as shown in fig. 2.2a. Consider the evolution through the circuit of the four basis states for the two input qubits, expressed in terms of the creation operators of the four input modes, labelled $1,\ldots,4$:

$$\begin{aligned}
a_1^\dagger a_3^\dagger &\to \frac{1}{\sqrt{2}} a_1^\dagger \left(a_3^\dagger + a_4^\dagger\right) \\
a_1^\dagger a_4^\dagger &\to a_1^\dagger a_2^\dagger \\
a_2^\dagger a_3^\dagger &\to \frac{1}{\sqrt{2}} \left[\left(a_3^\dagger\right)^2 - \left(a_4^\dagger\right)^2\right] \\
a_2^\dagger a_4^\dagger &\to \frac{1}{\sqrt{2}} a_2^\dagger \left(a_3^\dagger - a_4^\dagger\right),
\end{aligned} \tag{2.16}$$

where the indices on the RHS refer to the output modes. Rails 3 and 4 are detected in the circuit after this evolution. From this we see that the detections of 0 or 2 photons (second and third row in eq. 2.16) are only consistent with one input basis state, $|10,01\rangle\rangle = |01\rangle$ and $|01,10\rangle\rangle = |10\rangle$ respectively. So in

this case the input qubits are projected onto one of these states, and the output rails are not in a qubit state since they contain either 2 or 0 photons. This operation is described by the Kraus operators:

$$F_{\mathrm{T1}}^{\mathrm{fail}} \in \{|11\rangle\rangle\langle 01|, |00\rangle\rangle\langle 10|\}, \tag{2.17}$$

where the bras are the projections applied to the two input qubits and the kets the state of the two output rails.

On the other hand, if one photon is detected in output rail 3 or 4 (first and fourth row in eq. 2.16), the input state could have been either $|10,10\rangle\rangle = |00\rangle$ or $|01,01\rangle\rangle = |11\rangle$, so the qubits are projected onto an even superposition of these two terms. In the two cases the output rails contain a qubit in the $|10\rangle\rangle = |0\rangle$ or $|01\rangle\rangle = |1\rangle$ state respectively. The corresponding Kraus operators are:

$$F_{\mathrm{T1}}^{\mathrm{succ}} = \frac{1}{\sqrt{2}} \left(|0\rangle\langle 00| \pm |1\rangle\langle 11|\right), \tag{2.18}$$

where the sign depends on the specific detection pattern obtained. Since the transformation of both $a_1^\dagger a_3^\dagger$ and $a_2^\dagger a_4^\dagger$ leads to $a_3^\dagger$ with the same sign in eq. 2.16, the detection of a single photon in mode 3 is associated with a $+$ sign in eq. 2.18, whereas if the photon is detected in mode 4 the sign is $-$. Another way to see this would be to evolve the detection patterns across the T1 fusion circuit in reverse and project the resulting state onto the relevant two-qubit subspace.

The operation of this circuit can be understood as a ZZ parity check on the two input qubits. The fusion succeeds when the two inputs have parity $ZZ = +1$, and it fails if $ZZ = -1$. Assuming that the input qubits are maximally mixed, or equivalently, maximally entangled with separate systems not being measured, and that no noise is present, the two outcomes are each obtained with probability $1/2$. The measurement performed upon success can be described by a set of stabilizer generators:

$$S_{\mathrm{T1}}^{\mathrm{succ}} = \langle \pm X|XX, +Z|ZI, +Z|IZ\rangle, \tag{2.19}$$

which can be represented in ZX-calculus notation as a green spider with two inputs and one output leg. In the case of failure, each input qubit is measured in Z, giving operators $\{\pm ZI, \mp IZ\}$ on the two inputs, which correspond to two red spiders with a single input port, but the output rails do not contain a qubit, so the outcome cannot be fully described in the stabilizer ZX-calculus notation.

The T2 fusion is a fully destructive two-qubit measurement implemented by the circuit in fig. 2.2b. It is similar to T1 fusion, with the addition of one beamsplitter and one pair of detectors on the top modes. In this case, the evolution of the input basis states is the following:

$$\begin{aligned}
a_1^\dagger a_3^\dagger &\to \frac{1}{\sqrt{2}}\left(a_1^\dagger + a_2^\dagger\right)\left(a_3^\dagger + a_4^\dagger\right) \\
a_1^\dagger a_4^\dagger &\to \frac{1}{\sqrt{2}}\left[\left(a_1^\dagger\right)^2 - \left(a_2^\dagger\right)^2\right] \\
a_2^\dagger a_3^\dagger &\to \frac{1}{\sqrt{2}}\left[\left(a_3^\dagger\right)^2 - \left(a_4^\dagger\right)^2\right] \\
a_2^\dagger a_4^\dagger &\to \frac{1}{\sqrt{2}}\left(a_1^\dagger - a_2^\dagger\right)\left(a_3^\dagger - a_4^\dagger\right).
\end{aligned} \tag{2.20}$$

From this we see that there are two types of detection patterns. When two photons reach the same detector (second and third row in eq. 2.20), the input state could only have been $|10,01\rangle\rangle = |10\rangle$ or $|01,10\rangle\rangle = |01\rangle$ depending on which rail the photons are found in. Therefore, two single-qubit Z measurements are applied, described by the Kraus operators:

$$F_{\mathrm{T2}}^{\mathrm{fail}} \in \{|01\rangle\langle 01|, |10\rangle\langle 10|\}. \tag{2.21}$$

Detection patterns where the two qubits reach different detectors, however, are consistent with both input states $|10,10\rangle\rangle = |00\rangle$ and $|01,01\rangle\rangle = |11\rangle$, so the two qubits are projected onto a superposition of the two. The corresponding Kraus operators are those of a BSM:

$$F_{\mathrm{T2}}^{\mathrm{succ}} = \frac{1}{2}(|00\rangle \pm |11\rangle)(\langle 00| \pm \langle 11|), \tag{2.22}$$

where the sign is determined by the specific detection pattern as in the case of the T1 fusion.

Similarly to the T1 fusion, this operation can also be seen as a ZZ parity check on the inputs, with success and failure associated with the $+1$ and -1 check values respectively. When the two qubits are maximally mixed, the two outcomes occur with probability $1/2$ each. In this case both types of measurements have an associated set of stabilizer generators, since they project the inputs onto pure stabilizer states. In the case of success this is:

$$S_{\mathrm{T2}}^{\mathrm{succ}} = \langle \pm XX, +ZZ \rangle, \tag{2.23}$$

corresponding to a green ZX-spider with two input legs. For the failure outcome:

$$S_{\mathrm{T2}}^{\mathrm{fail}} = \langle \pm ZI, \mp IZ \rangle, \tag{2.24}$$

represented by two single-input red ZX-spiders.

In both fusions, when a failure outcome is obtained the value of the ZZ operator on the input qubits can be reconstructed by multiplying the single-qubit measurement values. In a T2 fusion, this means that half of the desired information can be obtained, and only the XX stabilizer value is lost. This gives a way of understanding fusion failure as a form of partial erasure. It is possible to change the *failure basis* of the fusion, i.e. the type of single-qubit measurement applied in case of failure, by performing single-qubit rotations on the inputs of the fusion [101]. For example, performing a Hadamard on each qubit results in a fusion whose success outcome is described by the same stabilizer operators, up to differences in the signs, but which measures the qubits in X when it fails. This operation corresponds to a red ZX-spider with two input ports. We refer to the original T2 fusion described above as a Z-type fusion and to this rotated version as an X-type fusion. The ability to modify the failure basis allows to choose which of the success stabilizer operators is erased in case of failure. Staggering the fusions

in the fusion network in time and introducing this bias adaptively based on the outcomes of previous fusions has been found to lead to large gains in the performance of LO-FBQC architectures [115].

From eqs. 2.16 and 2.20 we see that neither fusion admits detection patterns with one photon at each of the two detectors which follow a beamsplitter. This is a consequence of the HOM effect which causes the two photons to bunch. However, in the presence of errors, particularly photon distinguishability, it is possible for such anti-bunched patterns to arise. This is an example of a *nonsense pattern*, i.e., an unexpected detection pattern which signals that an error occurred. Another instance of this kind is the detection of more or fewer photons than expected. For example, finding only one photon in a T2 fusions heralds the occurrence of photon loss. When this type of detection pattern is obtained, it is impossible to reconstruct either of the success measurement operators, so a full erasure occurs [43].

As a fully destructive measurement, the T2 fusion naturally offers protection against photon loss and any other errors which alter the number of photons in the circuit, such as multiphoton contamination, since their occurrence is evident in the detection pattern. The same does not apply to the T1 fusion, whose ideal detection patterns may contain 0, 1 or 2 photons. In a realistic photonic architecture, it is essential to mitigate the impact of errors, and in particular photon loss. Therefore, the work described in this thesis focuses on the T2 fusion and other fully destructive fusion circuits.

The operators in eq. 2.22 are those of a Bell state measurement (BSM) which projects the input qubits onto one of the two Bell states $|\phi^{\pm}\rangle$. This type of measurement has been widely studied, due to the key role it plays in a number of quantum information protocols. Its implementation using linear optics was first proposed and experimentally demonstrated in the '90s [118–120]. The scheme allows to discriminate at most two out of the four Bell states, which gives the success probability $1/2$. It has been shown that it is impossible to perform a complete Bell state measurement, or, equivalently, a deterministic T2 fusion, with only linear optics and vacuum ancilla modes [121, 122], and in fact, $1/2$ has been proven to be the maximum achievable success probability with these elements [123]. Methods to break this limit have been proposed based on interactions with atoms [124], hyperentanglement [125],

squeezing [126], or qubit ancilla states. Among the latter is Grice's proposal [127], which boosts the success probability to $1 - \frac{1}{2^{n+1}}$ by using ancillary entangled states of 2^n qubits. The simplest version of this scheme uses a Bell pair to perform a BSM with probability $3/4$. Ewert & van Loock showed that the same probability can be achieved using four unentangled single-photon ancillae, and $5/8$ with two [128]. Both these proposals were found to achieve the maximum level of boosting possible with the ancilla states they use [129, 130]. While the original T2 fusion is a probabilistic BSM, fusion operations can more generally perform projections onto any set of maximally-entangled states which do not necessarily need to be orthogonal [97]. Extending the definition in this way allows better boosting with single-photon ancillae: two single photons can be used to obtain a T2 fusion with a $2/3$ success probability [97], and it has been conjectured that larger numbers of photons can be used to surpass the success rate of the best-performing known single-photon boosted BSM scheme [131].

Boosted fusion circuits are more prone to errors since they require more photons, each of which can introduce noise. In Ref. [43], the impact of photon loss on architectures with T2 fusions boosted with Grice's scheme is analysed. The conclusion is that one level of boosting, corresponding to a fusion success probability of $3/4$, is optimal for their best-performing architecture. In addition to increasing the rate of erasure due to loss, considered in this analysis, the use of more photons potentially affected by physical noise can cause higher rates of Pauli errors. The complexity of the system is also increased by the need to prepare and consume entangled ancilla states, and the noise requirements on the hardware are made more stringent. These trade-offs need to be accounted for when deciding whether to make use of boosted fusions. In the rest of this thesis, we restrict our analysis to standard, un-boosted fusion circuits.

2.2.2 Resource states

Resource states are typically entangled stabilizer states, and as such they are equivalent to *graph states* up to local Clifford operations on the qubits [132]. An n-qubit graph state is a quantum state which

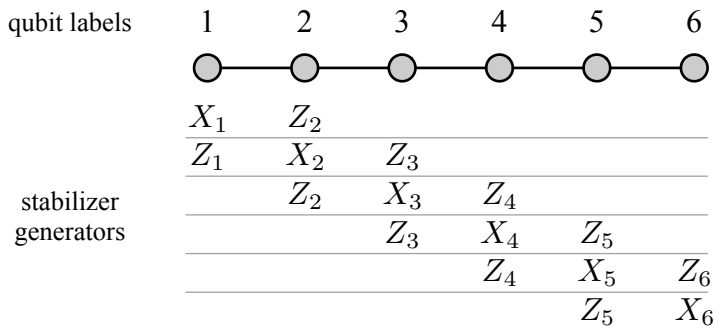

Figure 2.3: Graph state representation of a six-line linear cluster states. Vertices correspond to qubits and edges are placed between them such that the stabilizer generators of the state can be read off as products of X_i on each qubit and Z_j on all its neighbours.

can be represented as a graph G with one vertex per qubit and edges connecting the vertices in such a way that the stabilizer generators of the state can be read off as $X_i \prod_{j \in \mathcal{N}(i)} Z_j$, where $1 \leq i \leq n$ are the qubit labels and $\mathcal{N}(i)$ is the neighbourhood of vertex i in G. The simple example of a six-qubit linear cluster state is shown in fig. 2.3. In LO-FBQC, entangled resource states are typically grown by fusing so-called *seed states* in a process referred to as Resource State Generation (RSG) [43]. Seed states are smaller stabilizer states of a few qubits, produced directly from single photons using linear optical circuits followed by partial detection. This approach to RSG has the advantage of fitting naturally within the linear optical platform and allows to make use of known schemes and techniques for the generation of entanglement [97, 133, 134]. The main drawback is the resource overhead incurred due to the probabilistic nature of linear optical entangling operations. The use of matter-based emitters to produce deterministic photonic resource states is an active field of reserach [135–139], but challenges remain to be solved before such systems can be used in an LOQC architecture, including increasing the speed of the process, reliably producing states with indistinguishable photons and integrating the emitters with a photonic platform. Therefore, we focus on the all-linear-optical approach in the remainder of this thesis and describe it in more detail here. The use of quantum emitters to generate single photons for LO-FBQC is briefly considered again in Appendix C.

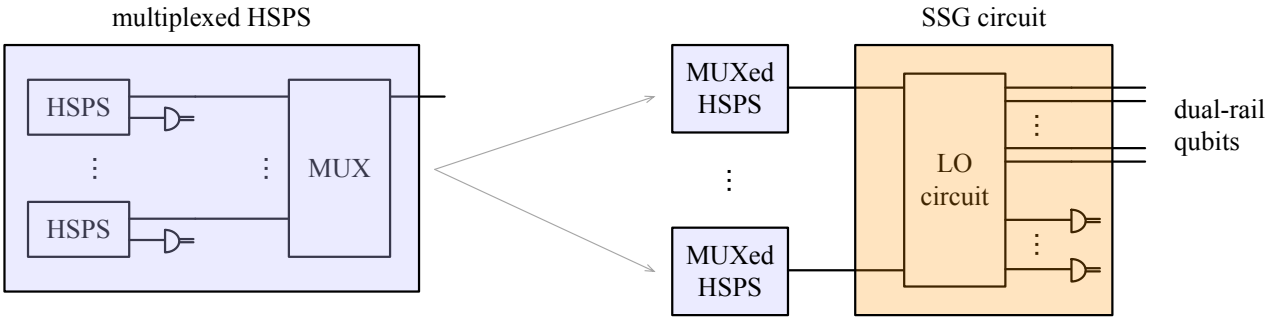

Figure 2.4: Seed state generation (SSG) from single photons. Heralded single photon sources (HSPSs) produce probabilistic single photons which are multiplexed through an active switch network (MUX) to obtain single photons with high probability. The output of these multiplexed HSPSs is then fed into an SSG circuit, comprised of a linear optical network and single photon detectors. Some of the photons are detected within the circuit and probabilistically herald the generation of an entangled state of dual-rail qubits on the output modes.

Resource State Generation

The starting point of linear optical RSG is Seed State Generation (SSG), i.e., the production of Bell pairs or small GHZ states from single photons. The latter states are named after Greenberger-Horne-Zeilinger (GHZ) [140], who introduced this new class of highly entangled states of three or more particles which have the form:

$$|n\text{-GHZ}\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle^{\otimes n} + |1\rangle^{\otimes n}\right). \tag{2.25}$$

The full basis consists of 2^n states obtained by applying local rotations to the qubits eq. 2.25. The steps of SSG are summarized in fig. 2.4. High-probability single photons are obtained from probabilistic HSPSs via multiplexing. These are then turned into seed states through the application of linear optical circuits in which some of them are detected and the rest are output. With some probability, the detectors register a pattern which heralds the preparation of the desired entangled state encoded in the output rails. This process is based on the same measurement-induced non-linearity used in linear optical fusion circuits. Many schemes to turn single photons into seed states have been proposed, starting with circuits to generate Bell pairs from four single photons with probability $3/16$ [141] and n-GHZ states from $2n$ single photons with probability $1/2^{2n-1}$ [101, 142], and including newer

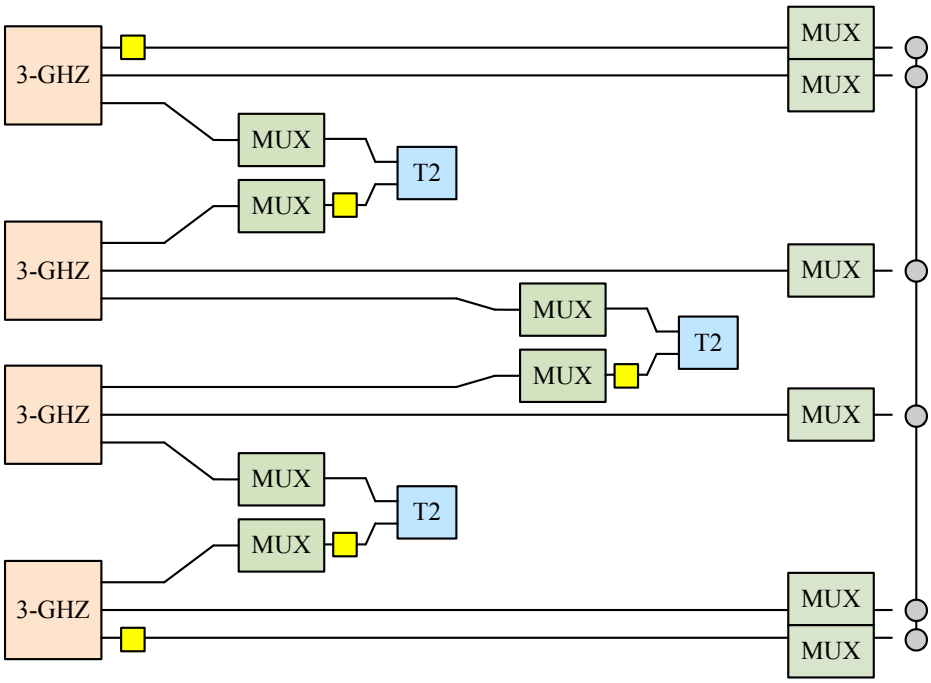

Figure 2.5: Example generation scheme for the six-qubit linear cluster state using 3-GHZ seed states (orange) and T2 fusions (blue). Lines represent dual-rail qubits and yellow boxes are single-qubit Hadamard gates, as in ZX-calculus notation. Three fusions are performed in two stages: first two disjoint pairs of 3-GHZs are fused in parallel and then a final T2 fusion is performed between their outputs. Qubits are multiplexed right before being measured or becoming part of the final resource state.

schemes with improved success rates and hardware requirements [97, 134]. The number of input photons and the success probability of SSG circuits decreases with the size of the output state, which motivates the use of small seed states as the starting point for RSG.

Once seed states are available, any resource state can be generated by performing fusions between them according to the structure of the target state, and possibly applying local Cliffords to the remaining qubits [43]. An example RSG scheme which uses 3-GHZ seed states and T2 fusions to produce the six-qubit linear cluster state above is shown in fig. 2.5. The starting seed state qubits can be divided into two groups: the *inner qubits* which are consumed by the RSG fusions, and the *outer qubits* which become part of the final resource state and go on to be detected in the fusion network. The protocol consists of a series of probabilistic fusion stages, each of which can contain multiple fusions acting on qubits from disjoint states but not on the same state. Since fusions are probabilistic, multiplexing is applied after SSG and each fusion on inner qubits to increase the output state probability.

We refer to these multiplexed fusions used to generate the resource state as RSG fusions, to distinguish them from the ballistic fusions between resource states in the fusion network. Importantly, after the SSG stage, all qubits are stored in delay lines and only sent through a multiplexing network right before they need to be fused. This minimises the number of lossy active switch networks encountered by the qubits. Compared to multiplexing all qubits at each stage, the cost is an increase in the size of the switch networks, which typically translates into a higher passive component depth. Since active phase shifters introduce significantly more loss than passive elements, this technique leads to an overall reduction in the loss rate on each qubit in the RSG.

Since the RSG scheme consists of a series of probabilistic operations, it is important to consider its resource overhead, which ultimately contributes directly to that of the full architecture. The cost of an operation which takes input resources with costs $\{C_i\}$, succeeds with probability p, and is multiplexed using a scheme with inefficiency γ can be calculated as:

$$C = \frac{\gamma}{p}\sum_i C_i, \tag{2.26}$$

where $\frac{\gamma}{p}$ is the average multiplexing amount necessary to obtain the desired output state and the sum is over all the input resources. This expression can be applied recursively to each RSG stage to calculate the total resource overhead. We calculate this cost in terms of number of multiplexed single photons, starting from the SSG stage, i.e., $C = 1$ for a muxed single photon. As an example, assuming the best case scenario of no errors, perfect multiplexing ($\gamma = 1$), an SSG scheme which uses 6 single photons to make a 3-GHZ state with probability of $1/32$ [142], and the standard T2 fusion success probability of $1/2$, we obtain a cost of 3,072 multiplexed single photons for the scheme in fig. 2.5. The result of this calculation is highly dependent on the ordering of the fusions. Consider an alternative scheme with the same seed states and fusions, but in which the resource state is built by adding one 3-GHZ at a time to the line. In that case, under the same set of assumptions, the total cost of the scheme

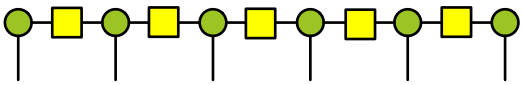

Figure 2.6: ZX-diagram representation of the six-qubit linear cluster state in fig. 2.3. This is obtained by replacing each vertex in the graph state green spider with an extra leg and placing Hadamard spiders (yellow) on all connections. In this case, legs which do not terminate at a spider on both ends are qubits.

turns out to be 4,224 muxed single photons, i.e., $\sim 40\%$ higher than for the previous ordering. As we will see in Chapter 6, optimization routines can be used to minimise the cost of generating a given resource state.

RSG with ZX-diagrams

The ZX-calculus is a particularly useful language for the representation and study of RSG schemes. The graph state corresponding to the target resource state has a straightforward mapping to the ZX-diagram picture [64]. Each vertex in the graph is replaced by a green spider with an output leg corresponding the qubit, and the edges in the graph become connections between the spiders through Hadamard nodes. This translation is shown in fig. 2.6 for the six-qubit linear cluster from fig. 2.3. The local Clifford gates necessary to capture any stabilizer states in this way can be added to the qubit legs as compositions of green and red spiders with multiple-of-$\pi/2$ phases plus the Hadamard node.

The generation of a resource state can also be represented naturally in this language. As discussed above, the success outcomes of two-qubit fusions correspond to spiders with two input legs. Similarly, seed states can be represented as spiders with only output legs representing the qubits. A 3-GHZ state has stabilizer generators $\langle +XXX, +ZZI, +ZIZ\rangle$, so it corresponds to a green spider with three output ports; a Bell pair with stabilizers $\langle +XX, +ZZ\rangle$ is a spider of either colour with two output legs, i.e., the inverse of a T2 fusion. All these spiders are usually specified up to the Pauli frame, which is dependent on the detection pattern obtained in the circuit but does not affect the structure of the diagram. For example, a different kind of Bell pair with stabilizers $\langle +XX, -ZZ\rangle$ is represented by a two-legged spider with an X rotation on one of the inputs, but this is usually omitted since it would

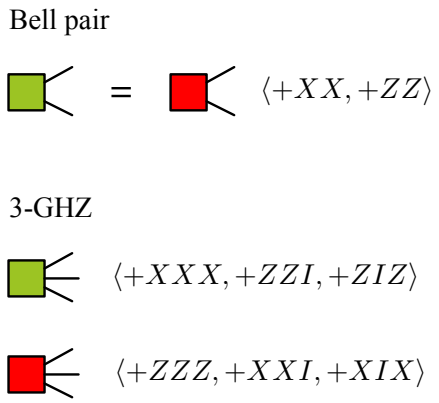

Figure 2.7: ZX-diagram representation of common seed states, namely Bell pairs and 3-GHZ states. These are green or red spiders with only output legs. In order to distinguish seed states from fusions in ZX-diagrams involving both, square spiders are used for the former.

not be implemented physically but rather tracked and accounted for classicaly at higher levels of the computing stack. In the context of RSG design and analysis, ZX-diagrams which combine seed states and fusions are commonly used, so to avoid confusion we turn to a non-standard notation whereby fusions are represented as round spiders and seed states as square spiders. The Bell pair and 3-GHZ seed state examples are shown in fig. 2.7.

The ZX-diagram representation of a resource state provides a starting point to obtain a recipe for its generation. A valid RSG scheme can be obtained from any decomposition of the ZX-diagram such that all the legs of fusion spiders are connected to seed state spiders, and in turn these have their legs either connected to a fusion, the inner qubits, or not connected to anything, the outer qubits. Using the identity rule, any connection can be replaced by a two-leg fusion or seed state spider to obtain a valid scheme, and larger spiders can be split into smaller ones by applying the reverse of the spider fusion rule, also known as the *split rule*. This introduces significant freedom in the choice of which seed states and fusions to use as building blocks of the scheme, in addition to the way in which they are connected to each other. One of the possible decompositions of the six-qubit linear cluster state into seed states and fusions obtained from in fig. 2.6 is shown in fig. 2.8.

To obtain an RSG scheme from this decomposition, an order needs to be chosen for the fusions and multiplexing stages. To capture this, we assign numerical labels to the fusions which indicate their

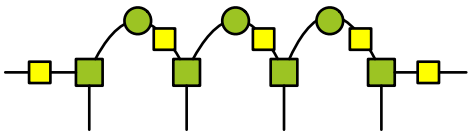

Figure 2.8: One of the possible decompositions of the ZX-diagram of the six-qubit linear cluster state into 3-GHZ seed states (green square spiders) and T2 fusions (green round spiders).

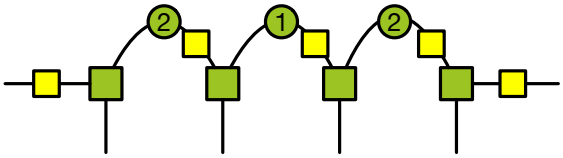

Figure 2.9: ZX-diagram representation of the RSG scheme for the six-qubit linear cluster state in fig. 2.5. Each fusion has a numerical label assigned to it, and fusion stages proceed from the highest to the lowest number, with fusions with the same label occurring in parallel.

time ordering. Fusions proceed in order from the highest to the lowest number label, and fusions with the same number occur at the same time. Only fusions acting on different connected components of the graph can be done in parallel, where a connected component is any seed state or set of seed states which have been joined by previous fusions. Fig. 2.9 shows this representation of the six-qubit cluster state RSG scheme described in the previous section.

As we have seen, the design of linear optical RSG schemes has many degrees of freedom, from the selection of the constituent seed states and fusions, to the choice of how to connect them to each other and in what order. FBQC schemes with resource states of varying sizes have been proposed, from two-qubit [116] to 24-qubit states [43], for which the design space becomes very large. Each of the possible choices impacts the performance of the scheme in terms of resource overhead and error rate. For example, large variance in the footprint requirements of different RSG schemes obtained by manipulating ZX-diagrams as described here were found in Ref. [143]. The focus of part of the work described in this thesis is on the optimization of the RSG design process and the development of schemes with improved performance on both the resource overhead and error rate fronts. This is the topic of Chapter 6.

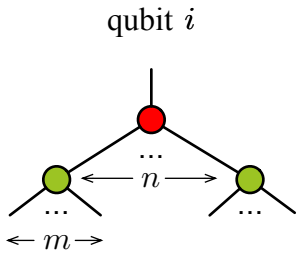

Figure 2.10: ZX-diagram of a qubit encoded using an (n,m) Shore code. The encoding uses nm physical qubits to create an n-repetition in X and an m-repetition in Z.

Encoded resource states

Large error tolerance gains have been obtained by FBQC schemes using *encoded resource states*, whose qubits are replaced by small QEC codes, referred to as *local encodings*, to protect against erasure due to fusion failure or loss [43, 54, 112, 115–117]. Quantum parity codes, described in Chapter 1 and also referred to as generalized Shor codes, are the most commonly used local qubit encodings, which have been shown to offer excellent tolerance to erasure due to photon loss [142, 144]. Following the terminology in [43], an (n,m) Shor-encoded qubit i consists of nm physical qubits and has logical operators

$$\begin{aligned} \overline{X_i} &= X_{i,1}X_{i,2}\dots X_{i,m}, \\ \overline{Z_i} &= Z_{i,1}Z_{i,m+1}\dots Z_{i,(n-1)m+1}, \end{aligned} \tag{2.27}$$

so it has an n-repetition in X and an m-repetition in Z. The $(nm-1)$ stabilizer generators are

$$\begin{aligned} S = \langle & X_{i,1}\dots X_{i,m}X_{i,m+1}\dots X_{2m}, \dots, X_{i,1}\dots X_{i,m}X_{i,(n-1)m+1}\dots X_{nm}, \\ & Z_{i,1}Z_{i,2}, \dots, Z_{i,1}Z_{i,m}, \dots, Z_{i,(n-1)m+1}Z_{i,(n-1)m+2}, \dots, Z_{(n-1)m+1,1}Z_{i,nm}\rangle. \end{aligned} \tag{2.28}$$

This structure can be represented in ZX-calculus language as in fig. 2.10. Examples of the Pauli webs corresponding to some of the logical operators and stabilizer generators of the (n,m) Shor code are shown in fig. 2.11. The constuction of Pauli webs makes it clear that there are multiple possible definitions of the logical operators of this local encoding, since different choices of ZZ (XX)

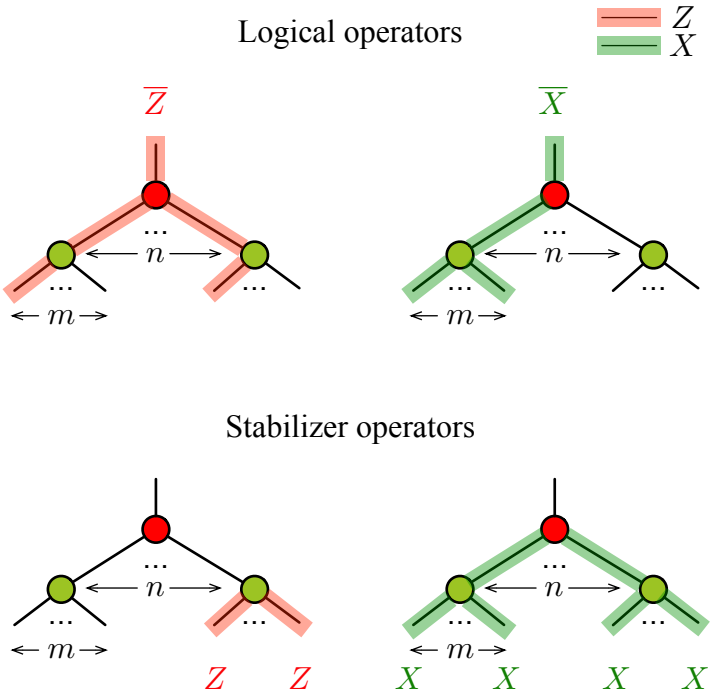

Figure 2.11: Pauli webs [55] representing two of the logical operator representatives and two of the stabilizer operators of the (n,m) Shor code.

stabilizers can be made at a green (red) spider. This redundancy protects the logical information from physical erasure.

In the fusion network, *encoded fusions* are performed between qubits encoded in this way. These consist of sequential T2 fusions between corresponding physical qubits (i,j), whose outcomes are used to reconstruct the logical Bell state measurement stabilizer generators

$$\begin{aligned}
\overline{X_iX_j} &= (X_{i,1}X_{j,1})\dots(X_{i,m}X_{j,m}) = \dots = (X_{i,(n-1)m+1}X_{j,(n-1)m+1})\dots(X_{i,nm}X_{j,nm}),\\
\overline{Z_iZ_j} &= (Z_{i,1}Z_{j,1})(Z_{i,m+1}Z_{j,m+1})\dots(Z_{i,(n-1)m+1}Z_{j,(n-1)m+1}) =\\
&= (Z_{i,2}Z_{j,2})(Z_{i,m+1}Z_{j,m+1})\dots(Z_{i,(n-1)m+1}Z_{j,(n-1)m+1}) =\\
&\dots\\
&= (Z_{i,m}Z_{j,m})(Z_{i,2m}Z_{j,2m})\dots(Z_{i,nm}Z_{j,nm}).
\end{aligned} \tag{2.29}$$

This degeneracy in the definition of the encoded measurement operators means that if some of the physical fusions fail to produce the XX or ZZ information due to failure or errors, the overall fusion

stabilizer generators can be reconstructed using the outcomes of other fusions. Moreover, if classical feedforward is allowed, the outcomes of earlier fusions can be used to set the basis of later ones so as to maximise the probability of obtaining all the information corresponding to a logical fusion success. The example of $n = m = 2$ is described explicitly in the Supplementary Information of Ref. [43]. Encoding resource states in this way has been shown to significantly improve the tolerance of FBQC schemes to both fusion failure and photon loss [115, 144–146], and in fact it has been proven that loss rates of up to 50% can in principle be tolerated using this technique [144]. Other types of graph local encodings have also been shown to offer good protection against erasure [112].

2.2.3 Physical architecture

The linear optical primitives of FBQC described in the previous two sections are used to construct a physical architecture which implements the fusion network required for the computation. The fusion network provides a prescription for which qubits from different resource states need to be fused, but it does not dictate a specific ordering of operations in space or time. Therefore, there is significant freedom in the design of its physical realisation. As we will see, the photonic platform is particularly well-suited to exploit this flexibility and give access to a wide range of possible architectures [147].

The central piece of an FBQC architecture is a *resource state generator*, or RSG module, which, as the name suggests, produces resource states. This consists of a collection of photonic chips containing all the components and linear optical circuits used in RSG, including HSPSs, SSG circuits, fusion circuits, and multiplexing networks. Photons are delayed and routed between chips using optical fibre. Control electonics is also required to choose the multiplexing settings based on the heralding signals from each probabilistic stage, and also to keep track of the Pauli frame due to the random stabilizer values associated with success outcomes of the entangling operations.

In this platform, fusion measurements are destructive, meaning that after the fusions prescribed by the fusion network are performed, only classical fusion outcomes are available, and the qubits cannot be

used again. This requires the RSG module to output resource states on a fixed clock to be consumed in rounds of fusions. The outer qubits are routed to their fusion locations through optical fibres, which also implement the delays needed to build the desired fusion network structure. All the outer qubits of a resource state are consumed by a fusion, and fusion outcomes are used to read out the state of the computation and inform subsequent steps. It is necessary for fusion circuits to be reconfigurable to implement feedfoward and the measurements necessary to perform logical operations [53]. However, the timescale of the classical processing and communication required to modify the fusion settings is decoupled from that of resource state generation, and therefore does not affect the amount of delay applied to individual photons.

A 2D grid of RSG modules can be used to generate a 3D fusion network by using time as the third dimension. This is done by sending each qubit through a different delay, i.e., an optical fibre of different length, before it is fused. For example, six-qubit RSG modules laid out in a grid of size $n_x \times n_y$ on the $x-y$ plane can be used to make a cubic fusion network. Four resource state qubits are fused in-plane with their neighbours emitted at the same time, i.e., with no delay. The other two qubits are fused in the z direction with qubits from resource states produced in the previous or following time-step by the same RSG module. To make this possible, the qubit to be fused in the future needs to go through a delay of one time step to wait for the next resource state to be generated and ready to be fused with it. Running the RSG modules for T time steps results in an $n_x \times n_y \times T$ cubic fusion network, as shown in fig. 2.12. Note that in order to allow the introduction of the topological features needed to perform logical operations, some of the fusions need to be reconfigurable, meaning that their basis can be changed or single-qubit measurement can be performed in their place. One way to enable this is by replacing a single fusion with multiple measurement options and using a switch to route between them.

The ability to route and store photonic qubits with optical fibre allows to make even more use of time in the construction of the fusion network. For example, a single resource state generator can

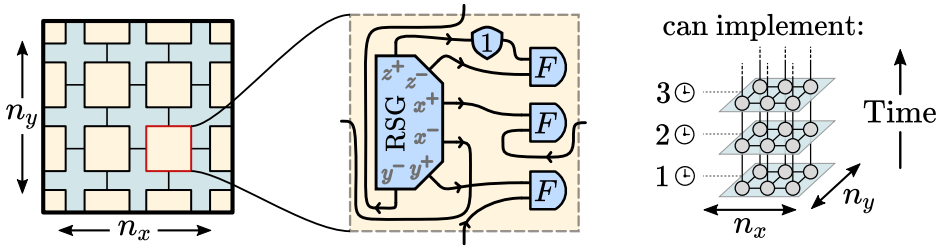

Figure 2.12: Implementation of a 3D LO-FBQC fusion network using a 2D grid of RSG modules and delays (figure from Ref. [147], Open Access).

be used to tile the full 3D fusion network, with the addition of just a few longer additional delay options and switching to choose between them. This is the most extreme incarnation of *interleaving*, which in general allows to implement a spectrum of architectures which use one RSG module to add multiple resource states to each temporal slice of the fusion network [147]. While this technique does not exclusively apply to the photonic platform, it is particularly powerful in this case since physical qubits can directly be moved across space and time with minimal loss using optical fibre, without the need for transduction.

Interleaving relies on a key property of FBQC architectures: modularity. Since resource states are constant in size and the fusion network has a regular, geometrically local structure, identical RSG modules can be used to implement the computation, with just the addition of fusion reconfigurability and classical feedforward. This is a particularly appealing property for the manufacturability of a large-scale quantum computing architecture with millions of physical qubits. The fixed resource state size also means that the architecture has constant optical depth, and the construction of both the RSG and the fusion network is such that switching, a known large contributor to the loss of linear optical architectures [148], is minimised along the path of each photon. Finally, in this framework, the RSG clock rate can be decoupled from the time scale necessary for feedforward at the level of the fusion network, since RSG and fusion cycles can continue to be performed to idle while the logical information is processed. This removes the need for photons to be stored in long delays while the higher level logic runs, further reducing the amount of error they are subject to. All these advantages of physical LO-FBQC architectures make them significantly simpler to realise and more tolerant to

noise than other schemes for photonic quantum computation.

2.3 Performance of LO-FBQC architectures

Any realistic universal, fault-tolerant quantum computing architecture is subject to requirements in terms of both its ability to tolerate noise and the resource overhead it incurs. Here we describe how these two aspects can be quantified and how improvements with respect to both are assessed in this thesis.

2.3.1 Threshold

The amount of noise a fault-tolerant quantum computing architecture can tolerate is captured by its *threshold*, i.e., the maximum error rate under which a quantum computation of any length and arbitrary accuracy can be performed efficiently. In an FBQC architecture, where all qubits are measured within a fusion, the threshold is expressed in terms of the error rate affecting fusion outcomes, which are used to evaluate the check operators of the QEC code implemented by the fusion network. The threshold value depends on the choice of error model, and when considering a realistic system it is most useful for this to capture the details of the physical noise it is subject to, such that the threshold expresses the tolerance of the scheme to real-world errors. An FBQC architecture with a higher threshold is one in which physical noise causes errors on fusion outcomes at a lower rate. The rate at which this translation happens depends in part on the choice of fusion network, but also on the design of the RSG scheme and the selection of primitive operations, which can be more or less resilient to physical noise.

In the linear optical platform, among others, two distinct types of errors exist: *heralded* errors whose presence is flagged by the value of a classical signal, e.g., an unexpected detection pattern, and *unheralded* errors which do not produce any detectable signal. The former correspond to the *erasure* of

outcome information at a fusion, and the latter to *Pauli errors* on the outcomes. The physical mechanisms behind the two types of error in an LO-FBQC architecture are described in detail in Chapter 4. For now, suffice it to say that heralded erasures are the dominant type of error in linear optics, since they can result from both fusion failure and photon loss, which are comparatively high probability events. Unheralded errors are mostly due to photon distinguishability or combinations of multiphoton contamination and loss, whose rates are significantly lower. QEC codes have higher threshold values for erasure compared to Pauli errors [107, 149], but the presence of the former has been found to lead to a decrease in the tolerance to the latter [150]. Therefore, it is important to consider both types of errors, and reducing the rate at which physical errors are converted into either is crucial to obtain a useful LO-FBQC architecture.

Although erasure occurs at much higher rates than Pauli errors at fusions, it has the advantage of being heralded, which allows to use multiplexing to discard affected attempts. Therefore, this type of error is most detrimental for ballistic fusions between resource states, whereas inside RSG it results in a larger resource overhead but not necessarily an increase in the error rate. On the other hand, unheralded errors are a cause for concern in both types of fusions, since they cannot be removed via multiplexing. Pauli errors in RSG fusions directly translate into errors on the qubits of the resource state, and they can result in much higher error rates than the Paulis occurring at individual fusions in the fusion network since errors from multiple fusions accumulate throughout the RSG scheme. Therefore, in order to increase the tolerance of the architecture to physical noise, it is most effective to focus on reducing the amount of erasure caused at fusions between resource states and the amount of Pauli errors caused at RSG fusions. This is the approach we take in this thesis. We do not explicitly calculate threshold values, but rather aim to compare the erasure and Pauli error rates in fusion networks and RSG schemes which use multi-way fusions to those of baseline architectures based on two-way fusions. Given a fusion network with a fixed set of resource states and fusions, lower RSG and fusion error rates translate into a higher physical noise threshold.

2.3.2 Footprint

Due to the probabilistic nature of linear optical entangling operations, the resource overhead of an LO-FBQC architecture is another crucial aspect which determines the feasibility of its implementation. The footprint is determined by two factors: the cost of generating the resource state and the number of resource states used in the fusion network.

The size of the fusion network is determined by its structure and by the target logical error rate: provided that the physical error rates are lower than the threshold value, the logical error rate is suppressed by increasing the size of the fusion network. The size required for a fixed target logical error rate depends on how far below the threshold the system operates - the closer to it, the larger the fusion network needs to be. So, reducing the fusion erasure and Pauli errors as discussed above also contributes to decreasing the resource overhead of the architecture. On the other hand, once a specific fusion network has been chosen, the types of fusions to be performed between resource states are in large part fixed.

The design of the RSG scheme is extremely flexible and can be easily modified to use different seed states and fusions, which allows more gains to come from selecting different building blocks. Since the focus of this thesis is the impact of changing the types of fusions used in the architecture, we focus on the RSG footprint as a proxy for the overhead of the entire architecture and compare the costs of different ways of constructing a fixed target resource state, assuming a fixed fusion network to assess the possible gains. The cost of a resource state is quantified by the number of multiplexed single photons, as described in Sec. 2.2.2.

2.3.3 Why multi-way fusions?

Fusions are generally defined as multi-qubit entangling projective measurements, but so far we have mostly focused on two-qubit operations, which are also the most commonly used in FBQC schemes in

Figure 2.13: Generation of a 6-qubit resource state from 3-qubit seed states, using fusions of different sizes. (a) When two-way fusions are used, four seed states and three fusions are needed. (b) A single three-way fusion applied to three seed states allows to generate the same state.

the literature. It has been proposed that it may be advantageous to consider *multi-way fusions*, which entangle more than two qubits, particularly in the context of RSG design, where they have been shown to potentially lead to significant reductions in resource overhead [143]. One way to understand this intuitively is based on the fact that, given a resource state with a fixed number of qubits, the number of seed states and fusion stages needed to produce it decreases as the size of the fusions increases. This is illustrated in the simple example of the generation of a 6-qubit state from 3-qubit seed states in fig. 2.13. On the other hand, as we will see in Chapter 3, the success probability of a linear optical fusion decreases with its number of inputs, introducing a trade-off between these two effects.

The same reasoning can also be applied to the scaling of Pauli errors occurring in the generation process. As described in more detail in Chapter 4, the two main sources of unheralded errors in RSG fusions are distinguishability and combinations of multiphoton contamination and loss. The former is a property of the fused qubits, so its total rate scales linearly with the number of fused qubits in the RSG. As seen in the simple example above, this decreases as larger fusions are used, so the rate of Pauli errors due to distinguishability is expected to be reduced accordingly. On the other hand, the probability of a multiphoton contamination event combining with photon loss at a fusion scales combinatorially with the number of fusion inputs. Therefore, using bigger multi-way fusions should cause an increase in the amount of Pauli errors due to this mechanism. When considering the joint impact of both these effects, we notice that to leading order, a single distinguishable photon can result

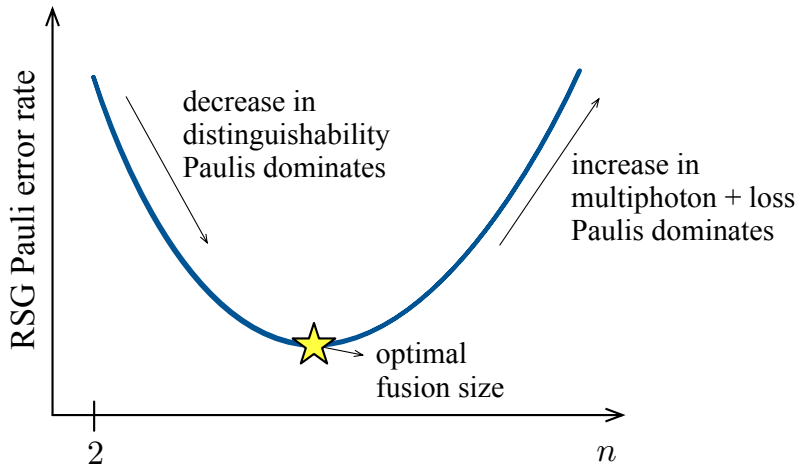

Figure 2.14: Sketch of the expected trend of Pauli errors in the generation of a fixed-size resource state as the size of the fusions, n, increases. This is not data. The decrease in errors due to distinguishability dominates initially and eventually the increase in errors due to multiphoton contamination and loss takes over. This results in an optimal fusion size (star) which minimises the Pauli error rate.

in a Pauli error, while two noise events are needed for multiphoton contamination and loss to cause a fault. Therefore, the expected trend is that the reduction in Pauli errors due to the former should dominate initially as the size of the fusion grows from $n = 2$, causing an overall drop in the error rate, until the slower increase in the latter takes over and makes it rise again. As a result, there should be an optimal fusion size which minimises the Pauli error rate in the RSG scheme. These trade-offs and trends are represented in the sketch in fig. 2.14.

This is a very simplistic model which does not take into account any of the details of the structure of the resource state and the specific scheme used to generate it, both of which have a significant effect on the final error rate. In addition, the shape of the curve in fig. 2.14 is strongly influenced by the ratio of the rates of the different physical errors involved, and in particular the optimal fusion size changes depending on it. Therefore, the reasoning presented here should only be regarded as a starting point which motivates the investigation described in the rest of this thesis. The initial goal of the analysis presented in the following chapters was to verify this expectation, and to find ways of exploiting this trend in the design of improved RSG schemes. Along the way, we also explored new fusion circuit designs aimed at improving this scaling as well as the footprint of the resource state, and we considered the impact of using multi-way fusions not only for RSG, but also as part of the

fusion network.

2.4 Remainder of thesis outline

The overarching goal of the work described in the rest of this thesis is to test the intuition presented above that multi-way fusions have the potential to lead to performance improvements in LO-FBQC architectures in terms of both threshold and footprint.

We begin by investigating linear optical circuits which perform multi-way fusions. In Chapter 3, we consider the structure of known multi-way fusions and propose minor alterations which result in novel circuit constructions which have a higher success probability and/or perform entangling projections on subsets of the input qubits upon failure. Without using any additional ancilla photons, these fusions help mitigate the steep increase in resource overhead typically associated with multi-way fusions in linear optics.

Next we study the behaviour of linear optical fusion circuits in the presence of physical noise, particularly focusing on the leading error mechanisms behind unheralded Pauli errors, which cannot be removed via multiplexing. In Chapter 4, a framework is presented which allows to map the effect of errors of this kind into a Pauli error model for any multi-way fusion which admits a decomposition into T1 and T2 sub-circuits. We use this to analyse the contribution of different error mechanisms to the Pauli error rates of standard n-GHZ fusions of increasing size. This method is based on simple graphical and analytical calculations, providing the ability to easily estimate the expected behaviour of different circuits, which is an essential part of the architecture design process. In Chapter 5, a more sophisticated modelling framework is presented which does away with the limitation to fusion circuits with a specific structure. We apply this simulation-based method to a wider variety of multi-way fusion circuits and compare the performance of different types of 4-GHZ fusions introduced in Chapter 3, finding that some of the new circuits have reduced Pauli error rates in the considered parameter

regime. Methods to further reduce the probability of Pauli errors are also discussed.

After studying linear optical fusion circuits and their Pauli error properties, their impact of different aspects of an LO-FBQC architecture is analysed. Chapter 6 focuses on their use in the context of RSG. First, heuristics are developed to assess the resource overhead and performance of different schemes with respect to erasure and Pauli errors. These models are then used to obtain footprint-optimal RSG schemes for the two resource states used in the first FBQC proposal [43], the four-star and six-ring states. A comparison of generation protocols using only T2 fusions and ones with multi-way fusions is performed, showing that the latter can significantly outperform the former with respect to both footprint and Pauli error rates. Finally, Chapter 7 investigates the use of linear optical multi-way fusions as primitives of the computation in FBQC, between resource states in the fusion network. A model of partial erasure due to photon loss is used to assess various fusion circuits, and T2 fusions are found to still offer the best protection against this kind of error.

2.5 Conclusion

Photonic qubits have a number of properties which make them particularly well-suited for the realisation of a large-scale quantum computing architecture, including the fact that all the components required can be manufactured reliably in high-volume semiconductor fabrication facilities. The main drawback, on the other hand, is the inherently probabilistic nature of single photon sources and entangling operations in this platform. This results in large resource overheads, since multiplexing is required at each step to increase the success probability.

Footprint is one important aspect which determines the feasibility of building a useful quantum computer. The amount of physical noise the quantum computing architecture can tolerate, its threshold, is another fundamental metric. Both these properties are affected by the primitives used in the generation of a resource state and in the fusion network, and the goal of architecture design is to identify

the choices which allow to minimise resource requirements while increasing error tolerance.

The research described in the rest of this thesis is based on the intuition that multi-way fusions, i.e., fusions between more than two qubits, have the potential to improve the performance of linear optical FBQC architectures with respect to both footprint and threshold, particularly when used for resource state generation. We begin by presenting linear optical multi-way fusion circuits in Chapter 3, where we introduce novel schemes with improved entanglement generation properties. In Chapters 4 and 5 we introduce frameworks to analyse the performance of these fusions in the presence of physical noise and compare circuits in isolation. In Chapter 6 we apply the results obtained to evaluate the impact of multi-way fusion circuits on the generation of resource states, first in a general setting and then considering two specific examples from the literature. Finally, in Chapter 7 we consider the potential impact of using multi-way fusions between resource states in the fusion network.

Chapter 3

Linear Optical Multi-Way Fusion Circuits

This chapter starts by describing known linear optical circuits to fuse more than two qubits, and then introduces novel fusion circuits with improved properties, namely the ability to entangle subsets of input qubits upon failure and increased success probability. The construction of these new fusions can be generalised to arbitrary numbers of inputs, and consists simply in the addition of at most one layer of beamsplitters to the inputs of the original fusion circuits.

The two-qubit fusion circuits introduced in the previous chapter can be generalised into entangling operations acting on larger numbers of input qubits. We begin this chapter by introducing linear optical circuits which can be used to perform such operations, particularly focusing on projections onto n-qubit GHZ states. The use of these multi-way (or multi-qubit) fusions in LO-FBQC has not been studied as closely as schemes based on two-way fusions, but it has been shown that they can be used for the construction of fault-tolerant fusion complexes [54], and that they can, in principle, lead to significant overhead reductions in linear optical RSG [143]. In this chapter we take a first step towards furthering this line of research by studying the structure of n-GHZ analyzer circuits, which we refer to as *standard* n-GHZ fusions, and presenting new linear optical circuits which entangle the input qubits more efficiently.

Linear optical fusions of all sizes suffer from the disadvantage of being inherently non-deterministic,

and we will see that the success probability of standard n-GHZ fusions decreases exponentially with the number of input qubits. Therefore, their use in LO-FBQC architectures would result in a large increase in the resource overhead, because RSG fusions have to be multiplexed to generate resource states with sufficiently high probability, and a larger fusion network has to be built to deal with the partial erasure of information from failure outcomes. Many of the ancilla-based boosting schemes mentioned in the previous chapter for T2 fusions can be applied to the larger fusion circuits too, but the added cost of generating the ancillae and increased hardware requirements remain. In this chapter, we introduce two new all-linear-optical n-GHZ fusion circuit designs which perform entangling projections with higher probability without requiring additional resources. In both cases, the only modification to the standard n-GHZ fusion circuit is the addition of a few extra 50:50 beamsplitters. The first design projects onto the desired n-GHZ state with the same probability, but for $n > 2$ it turns some of the single-qubit measurement failure outcomes into projections onto smaller entangled states. The other design includes both circuits with unchanged success probability and entangling failure outcomes, and circuits with increased n-GHZ projection probability for $n \geq 4$.

The operations performed by all circuits described in this chapter have been characterised by performing Fock state simulations of the circuits acting on input qubits which are maximally entangled with separate unmeasured systems, and by postprocessing the simulation output to compute quantities of interest such as the stabilizer operators of the measurements and their probabilities. The steps of this analysis are described in more detail in Chapter 5.

3.1 Standard multi-way fusions

Since the introduction of the class of GHZ states, their generation and projection onto them have become central to many quantum information processing tasks, including some of the LOQC schemes previously described [108, 110]. The linear optical Bell state measurement (BSM) scheme (or T2

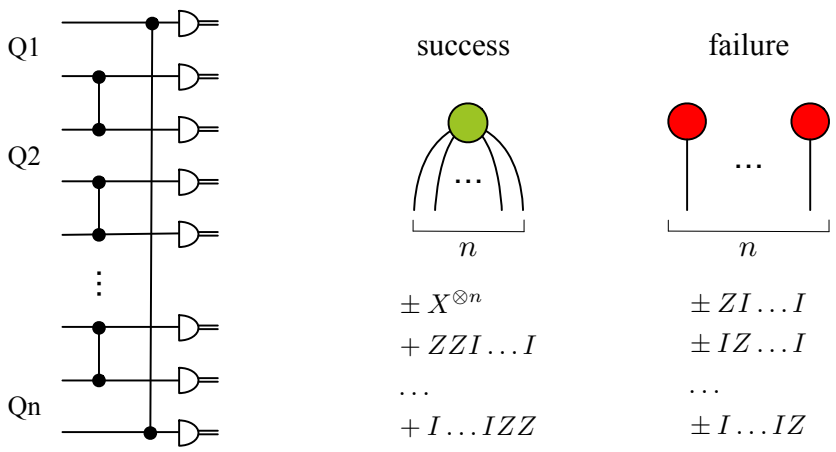

Figure 3.1: Circuit and ZX-diagram representations of the n-GHZ fusion. Rails from adjacent qubits are paired by n beamsplitters and detected. The fusion succeeds when each pair of detectors after a beamsplitter receives one photon, and in this case the input qubits are projected onto an n-GHZ state, represented by a green spider with n input legs. In all other cases, individual Z measurements are performed on all qubits.

fusion) was generalised to the unambiguous discrimination of GHZ states with the introduction of a *GHZ-state analyzer* acting on an arbitrary number of input qubits [151]. Here we refer to any circuit which probabilistically projects onto an n-qubit GHZ state as an *n-GHZ fusion*, and to the GHZ-state analyzer as a *standard* n-GHZ fusion. In the dual-rail qubit picture, the latter consists of a "loop" of beamsplitters which couple adjacent rails from neighbouring inputs followed by detectors, as shown in fig. 3.1.

Following the same procedure used to obtain the projectors associated with each detection pattern in Chapter 2, it can be seen that a successful entangling measurement is heralded by the detection of one photon after each beamsplitter, in which case a projection onto an n-GHZ state is applied:

$$F^{\text{succ}}_{n\text{-GHZ}} = \frac{1}{2}(|0\rangle^{\otimes n} \pm |1\rangle^{\otimes n})(\langle 0|^{\otimes n} \pm \langle 1|^{\otimes n}), \tag{3.1}$$

where the sign is determined by the specific detection pattern obtained. The stabilizer generators of this measurement are

$$S^{\text{succ}}_{n\text{-GHZ}} = \langle \pm X^{\otimes n}, \{+Z_1 Z_i \; \forall \; i \mid 2 \leq i \leq n\} \rangle, \tag{3.2}$$

so in the ZX-diagram notation this outcome is a green spider with n input legs, and the success

probability when the fusion acts on n maximally mixed qubits is $1/2^{n-1}$. In the remaining $1-1/2^{n-1}$ of cases, at least one of the detectors receives two photons and single qubit Z measurements are performed on all qubits. There are 2^n-2 projection operators onto all computational basis states except $|0\rangle^{\otimes n}$ and $|1\rangle^{\otimes n}$, described by the stabilizer generators

$$S^{\text{fail}}_{n\text{-GHZ}} = \langle \{ \pm Z_i \prod_{j\neq i} I_j \; \forall \; i \mid 1 \leq i \leq n \} \rangle. \tag{3.3}$$

These are all represented by n single-input red spiders. Similarly to the T2 fusion, this circuit has the property of heralding the presence of some kinds of errors, including photon loss, through the detection of unexpected patterns.

One way the structure of an n-GHZ fusion can be understood is as a concatenation of $(n-2)$ T1 fusions followed by one T2 fusion. The success outcome is obtained when all the $(n-1)$ underlying two-way fusions succeed, hence the exponential decrease in success probability with n. The example of a 4-GHZ fusion circuit is shown in fig. 3.2, from which it is apparent that this decomposition is not unique. The T1 fusions can be laid out in series, in parallel, or, for general n, as a combination of the two, and they can feed into one or both outputs of the final T2 fusion. The ZX-calculus representation is helpful to understand the set of possible options, which consist of any configuration of the constituent T1 and T2 spiders resulting in a diagram with n input ports and no outputs. The spider fusion rule of the ZX-calculus allows to map between decompositions. In the case of the 4-GHZ fusion, there are only two topologically distinct circuits shown in fig. 3.2.

A generalization of the T1 fusion to n qubits can be easily obtained from this modular view of the n-GHZ fusion by replacing the final T2 fusion with a T1 fusion, so that the circuit acts on n input qubits and outputs one when the projection is successful. Any n-GHZ fusion ZX-diagram can be transformed in this way simply by adding an output leg to the T2 fusion spider, so multiple configurations exist for a given n. These all correspond to the same linear optical circuit, which is the n-GHZ analyzer without

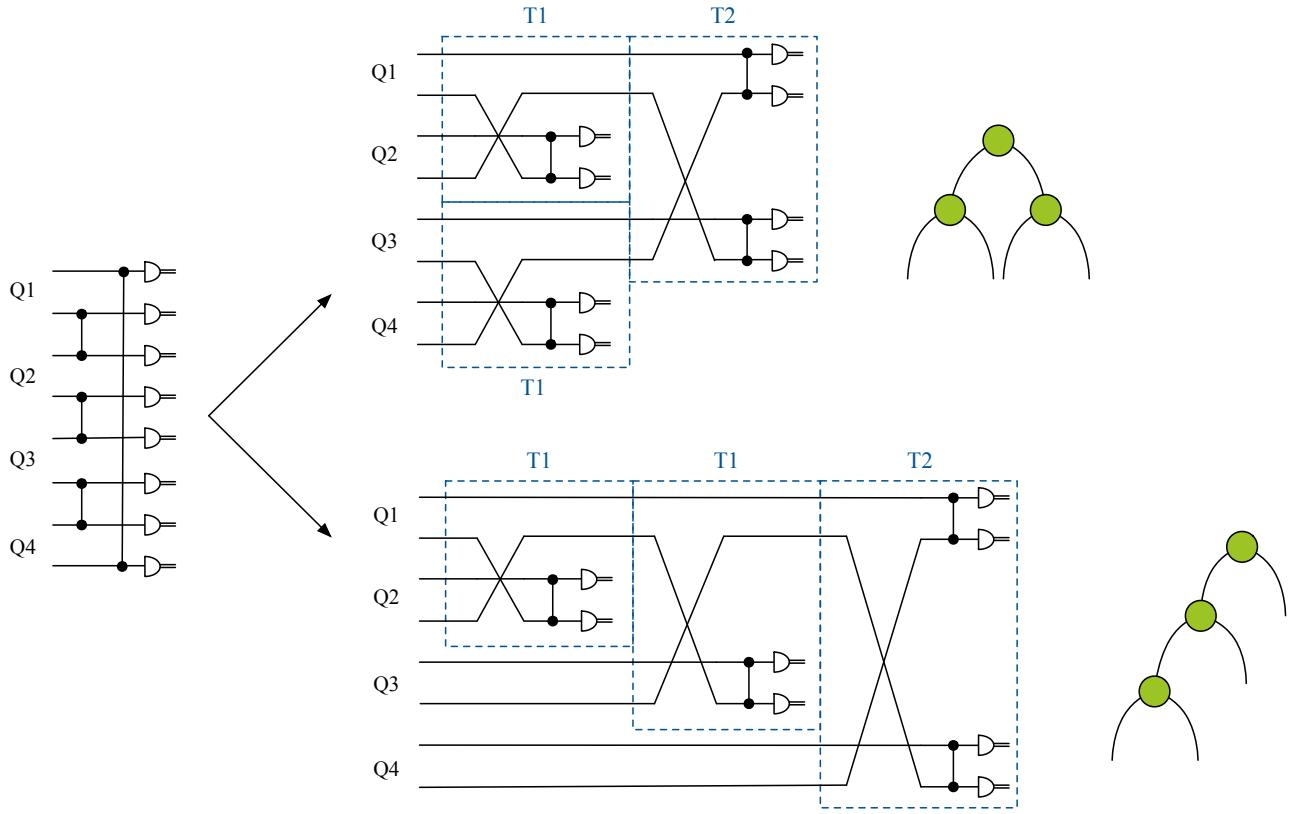

Figure 3.2: Decomposition of the 4-GHZ fusion (left) into T1 and T2 fusions. The circuit has two interpretations as a combination of three two-way fusions. One is a *tree* with two parallel T1 fusions feeding into a final T2 (top) and the other is a *chain* of serial T1 connected to one of the T2 inputs (bottom). The ZX-diagrams corresponding to the success outcomes are shown for the two cases.

one of the beamsplitters and the corresponding two detectors, as shown in fig. 3.3. This operation has been referred to as a T1 n-fusion in the literature [143]. The probability of success is the same as for the n-GHZ fusion and the corresponding operator takes the form

$$F^{\text{succ}}_{n-\text{T1}} = \frac{1}{\sqrt{2}}(|0\rangle\langle 0|^{\otimes n} \pm |1\rangle\langle 1|^{\otimes n}), \tag{3.4}$$

with stabilizer generators

$$S^{\text{succ}}_{n-\text{T1}} = \langle \pm X|X^{\otimes n}, +Z|ZI\ldots I, \ldots, +I\ldots IZZ\rangle. \tag{3.5}$$

In case of failure, the input qubits are individually measured in Z, but in contrast to the two-way T1 fusion, the state of the two output rails may or may not be in the qubit subspace depending on the

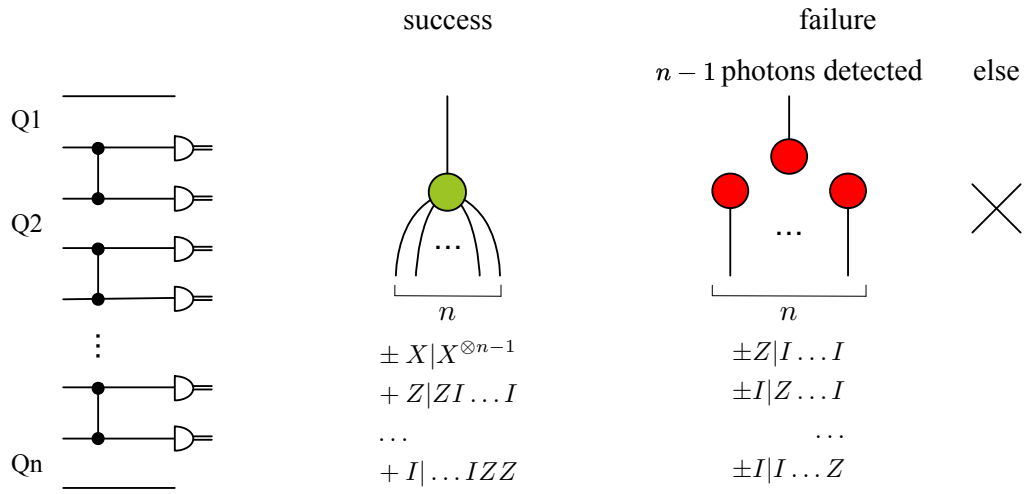

Figure 3.3: Generalised n-way T1 fusion, which takes n input qubits and outputs one upon success. The circuit is the n-GHZ analyzer where one beamsplitter and the associated detectors have been removed. Like the two-way T1, the success operation is represented by a green ZX-spider with $n-1$ inputs and one output. The circuit has two failure modes. In some cases, $n-1$ photons are detected in total and the output rails contain a qubit in an eigenstate of the Z operators, so the corresponding ZX-diagram consists of n single-leg red spiders. If more or fewer photons are detected, however, the output is not a qubit and the ZX notation cannot be used.

specific failure mode. In fact, in this case it is possible for the fusion to fail but still detect a total of $(n-1)$ photons, so that a single photon is output on the two undetected rails in an eigenstate of the Z operator.

The failure outcomes of the fully destructive T2 and n-GHZ fusions can be interpreted as partial erasure of the information that is available upon success [43]. For example, the failure operators of the T2 fusion in eq. 2.21 provide the value of the ZZ stabilizer operator. This means that when the fusion is performed, the ZZ information can be obtained with 100% probability in the absence of noise, but the XX information is only available 50% of the time, when the fusion succeeds. It is possible to choose what information to always get access to by applying unitaries to the input qubits. In fact, the operations performed by the T2 and n-GHZ fusions can be transformed into projections onto any state which is in the same *equivalence class* as a Bell pair or n-GHZ state under local unitaries and graph isomorphism by simply applying local unitaries on the input qubits [43]. For instance, according to the duality rule of ZX-diagrams, the color of the ZX-spiders associated with the fusions can be changed by applying single-qubit Hadamards to all the inputs.

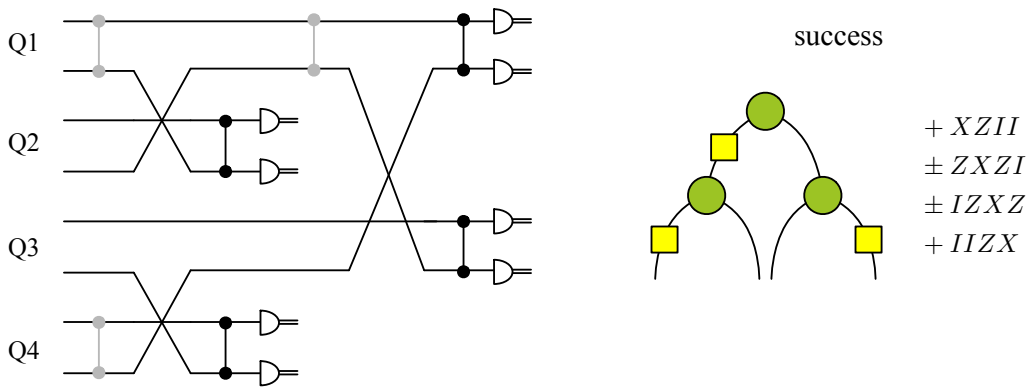

Figure 3.4: Four-line fusion, obtained by adding Hadamard rotations within the n-GHZ fusion circuit, shown in grey in the linear optical circuit and as yellow squares in the ZX-diagram. This is an example of how the modular structure of the circuit allows us to obtain circuits which project onto connected graph states outside of the n-GHZ equivalence class for $n > 3$. Projections onto any state with a tree ZX-diagram structure can be built by adding local Cliffords within the circuit in this way.

In the case of n-GHZ fusions with $n \leq 3$, this application of local unitaries allows us to obtain a projection onto all possible connected graph states, since there is only one equivalence class of this kind. However, for $n > 3$ multiple disjoint equivalence classes exist [132], so not all of them can be accessed in the same way. The decomposition into T1 and T2 fusions provides a way of expanding the set of possible projections by inserting single-qubit unitaries between two-qubit fusions within the circuit, allowing to project onto any state with a tree representation in the ZX-diagram notation. The example in fig. 3.4 is a circuit which projects onto a four-line, a member of a different equivalence class for $n = 4$ than the 4-GHZ. n-GHZ fusions are in principle sufficient for the generation of any stabilizer resource state, since a graph state can be generated by placing a $(d+1)$-GHZ projection circuit on every degree d node and a Bell pair on every edge and on the extra input, following the reason in Ref. [43]. The desired stabilizer resource state can then be obtained by applying local Cliffords to the unmeasured Bell pair qubits. However, the flexibility to project onto a more general class of states gives access to a wider array of options for the implementation of FBQC schemes, making non-GHZ fusions interesting to consider as part of the design process.

The modular structure of the n-GHZ fusions can also be exploited to increase the success probability through intermediate multiplexing. Each two-way fusion within the circuit can be repeated multiple times and an active switch network can be used to route the outputs of successful attempts to the next

operation. This helps combat the exponential decrease in success probability as the fusion size grows, but it increases the system complexity due to the switching and feedforward required, which in turns results in higher error rates on the input qubits. In order to determine the effective improvement given by intermediate multiplexing, an analysis of the trade-off between the reduction in resource overhead and the added sources of physical error is necessary. The addition of switching stages also introduces time-ordering, so different circuit decomposition are no longer equivalent and need to be analyzed separately. We focus on non-multiplexed circuits in this thesis.

3.2 Multi-way fusions with entangling failure outcomes

From the rules of the ZX-calculus described in Chapter 1, we know that a green spider can be turned into a red one by applying Hadamards on all its legs, and according to the identity rule, phaseless red and green spiders with two legs are equivalent. In the context of fusion measurements, this means that if a real $50:50$ beamsplitter is applied to each input qubit of the T2 fusion circuit in fig. 2.2b, the success outcomes remain the same up to a difference in the signs associated with the two stabilizer generators. In this new circuit with Hadamards on the inputs, a success operation is a measurement of $\langle +XX, \pm ZZ \rangle$, represented by a red spider, as opposed to $\langle \pm XX, +ZZ \rangle$ for the T2 fusion in fig. 2.2b. By the symmetry of ZX diagrams we consider in the context of an FBQC architecture, the possible differences in the phases of the stabilizer generators between the two circuits can be tracked separately, so the two can be considered interchangeable. This only applies to the success outcomes of the fusion, since the identity rule does not apply to the single-leg spiders corresponding to the failure outcomes. In fact, these are transformed from Z to X single qubit measurements by the addition of input Hadamards. For this reason, we refer to the original circuit in fig. 2.2b as a Z-fail T2 fusion and to the version with input Hadamards as an X-fail T2 fusion. Fig. 3.5 summarizes the outcomes of the latter.

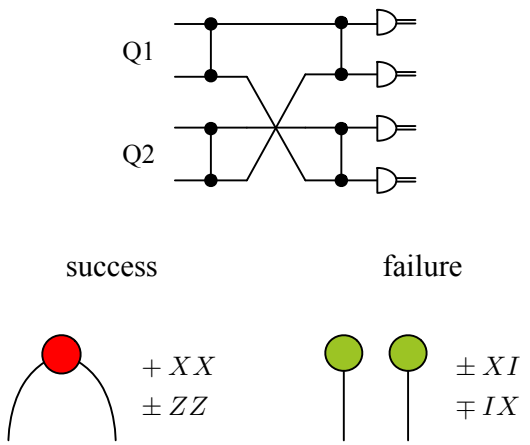

Figure 3.5: X-fail T2 fusion circuit, obtained by applying a Hadamard to each input qubit, and its outcomes. The stabilizer generators of the success outcomes are the same as the Z-fail T2 fusion up to signs, reflected in the equivalence of their ZX diagram representations. The failure outcome is different, consisting of two single-qubit X measurements in this case.

The invariance of the success outcomes can be used to obtain new n-GHZ fusion circuits from their decomposition into two-way fusions by replacing the T2 fusion with its X-fail version. In the case of success, the circuits are equivalent and different decompositions of a given n-way fusion perform the same projection. However, this equivalence does not hold for failure outcomes which are modified by the Hadamards. In the n-GHZ fusions with a Z-fail T2, failure of any of the constituent two-way fusions leads to single-qubit Z measurements on all the inputs. One way to understand why this is the case is to notice that when the T2 fusion fails, the result is a ZX diagram with green spiders connected to single-input red spiders, which disentangle all the input qubits due to the copy rule of the ZX-calculus. On the other hand, when an X-fail T2 fusion is used, green single-input spiders are introduced which do not have the same disentangling effect. As a result, some of the failure outcomes still entangle part of the input qubits. The structure of the network determines which qubits are entangled, so distinct decompositions are not equivalent in this case. Fig. 3.6 shows the outcomes associated with failure of the T2 fusion in the two four-way fusions in fig. 3.2. Other failure modes are harder to represent in this language due to T1 failures not being expressible in the ZX notation, but a similar mechanism applies in those cases as well. When considering all the possible failure outcomes, the version with a chain of T1 fusions (fig. 3.6b) performs a 3-GHZ projection on three of the qubits with probability $1/8$, and only performs four single qubit measurements with a reduced probability

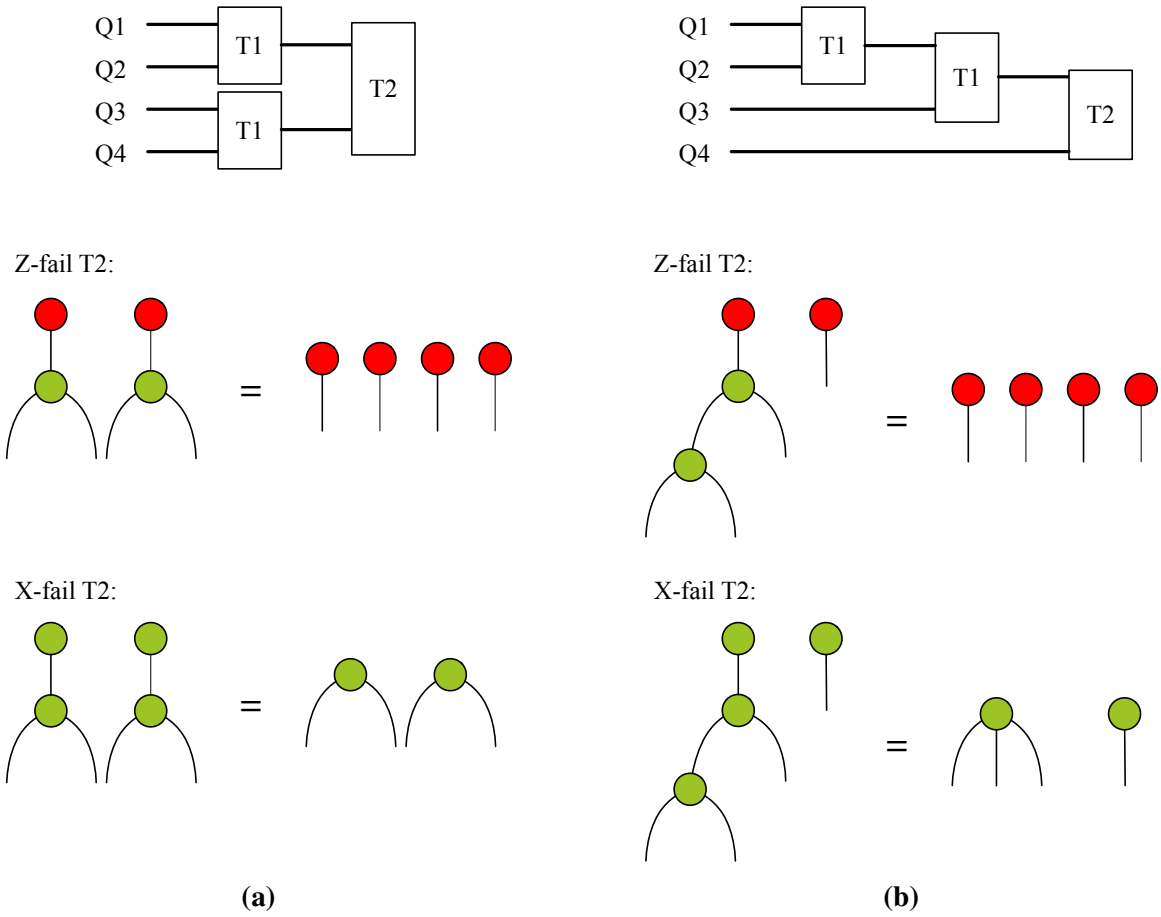

Figure 3.6: Outcomes of the two distinct versions of the 4-GHZ fusion associated with failure of the final T2 fusion with (X-fail) and without (Z-fail) input Hadamards. (a) Tree decomposition of the 4-GHZ fusion. In the Z-fail circuit, failure of the T2 introduces single-leg red spiders which result in four single-qubit Z measurements by the copy rule of the ZX calculus. In the X-fail circuit, the single-leg green spiders are merged with the ones they are connected to, leading to a projection onto two Bell pairs. (b) Chain decomposition of the 4-GHZ fusion. The T2 fusion in the Z-fail circuit similarly leads to a product measurement of all input qubits in Z. In the X-fail circuit, the result is a projection of three of the qubits onto a 3-GHZ and the measurement of the fourth in X.

of $3/4$. The tree version (fig. 3.6a) produces entanglement at a higher rate: with probability $1/8$ it projects two pairs of qubits onto two Bell pairs, and with probability $1/4$ it does one Bell state measurement and single qubit measurements on the other two qubits. This circuit projects onto at least one Bell pair with probability $3/4$, and in particular onto two Bell pairs or a 4-GHZ state $1/4$ of the time. Therefore, in a way this circuit can be considered to produce strictly more entanglement than a boosted T2 fusion, with the caveat that in this case the sets of qubits which are entangled by the fusion vary depending on the outcome.

As shown in fig. 3.7, when an n-GHZ fusion is represented as a monolithic linear optical circuit, the

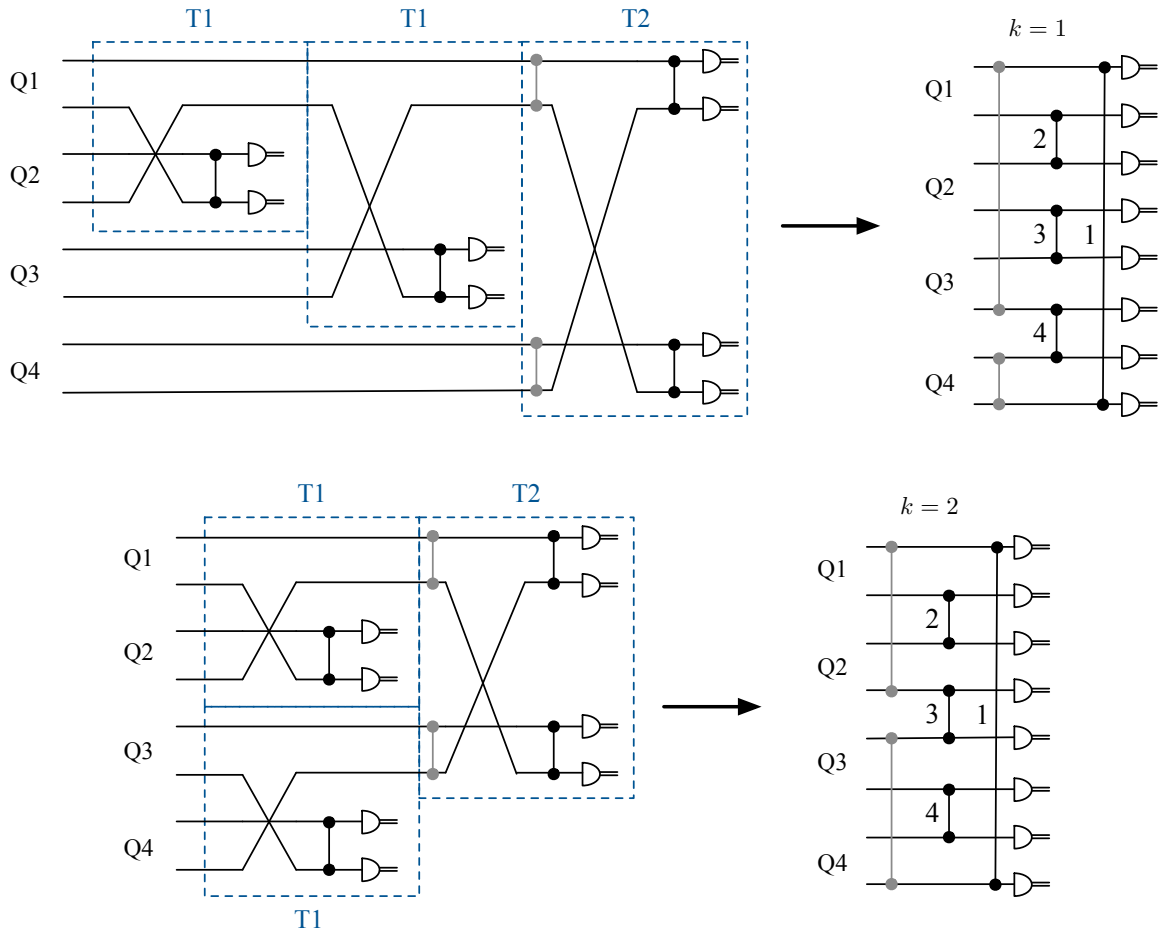

Figure 3.7: Representation of the two decompositions of the 4-GHZ fusion as a monolithic linear optical circuit with two extra 50 : 50 beamsplitters (grey). In the two cases, the two Hadamards are seen to couple different pairs of fundamental beamsplitters (black). In the chain fusion (top), the pair $(1,4)$ is coupled so the distance is $k = |1-4| \bmod 2 = 1$. In the tree fusion (bottom), the pair is $(1,3)$ and $k = 2$. For 4-GHZ fusions, these are representatives of the only two distinct types of fusions with two extra Hadamards.

two extra Hadamards are seen to couple the two input ports of a pair of *fundamental beamsplitters*, i.e., beamsplitters of the original n-GHZ analyzer circuit. Labelling these as $\{1,\ldots,n\}$, starting with the one between the top and bottom modes and then proceeding downwards, the pair of extra Hadamards can be identified by the indices of the beamsplitters it couples, (i,j). For example, $(1,4)$ in the top circuit in fig. 3.7 and $(1,3)$ in the bottom. Due to the cyclic symmetry of the n-GHZ fusion circuit, all circuits with the same distance $k = (|i-j| \bmod \lfloor n/2 \rfloor)$ between the coupled beamsplitters are equivalent up to swaps of the input qubits. So, for instance, in the case of a 4-GHZ fusion, Hadamard pairs between $(1,2),(2,3),(3,4)$ and $(1,4)$, for which $k = 1$, all give circuits which perform the same measurement operators up to permutations of qubits. For an n-GHZ fusion, $1 \leq k \leq \lfloor n/2 \rfloor$ so there

n	k	Projection type	Probability
3	0	3-GHZ	1/4
		$3 \times Z$	3/4
	1	3-GHZ	1/4
		BP + X	1/4
		$X + 2 \times Z$	1/2
4	0	4-GHZ	1/8
		$4 \times Z$	7/8
	1	4-GHZ	1/8
		3-GHZ + X	1/8
		$4 \times Z$	1/8
		$X + 3 \times Z$	5/8
	2	4-GHZ	1/8
		2×BP	1/8
		1×BP + $2 \times Z$	1/2
		$4 \times Z$	1/4
5	0	5-GHZ	1/16
		$5 \times Z$	15/16
	1	5-GHZ	1/16
		4-GHZ + X	1/16
		$5 \times Z$	3/16
		$X + 4 \times Z$	11/16
	2	5-GHZ	1/16
		1×3-GHZ + 1×BP	1/16
		1×3-GHZ + $2 \times Z$	1/8
		1×BP + $3 \times Z$	5/16
		$5 \times Z$	7/16

n	k	Projection type	Probability
6	0	6-GHZ	1/32
		$6 \times Z$	31/32
	1	6-GHZ	1/32
		5-GHZ + X	1/32
		$6 \times Z$	7/32
		$X + 5 \times Z$	23/32
	2	6-GHZ	1/32
		1×4-GHZ + 1×BP	1/32
		1×4-GHZ + $2 \times Z$	1/16
		1×BP + $4 \times Z$	11/32
		$6 \times Z$	17/32
	3	6-GHZ	1/32
		2×3-GHZ	1/32
		1×3-GHZ + $3 \times Z$	10/32
		$6 \times Z$	20/32

Table 3.1: Types of projections and their probabilities for different n-GHZ fusion circuits with $3 \leq n \leq 6$, assuming n maximally mixed input qubits. $k = 0$ denotes the standard n-GHZ fusion and $k > 0$ refers to circuits with one extra pair of beamsplitters coupling two of the fundamental beamsplitters. The locations of the extra beamsplitters are determined by the value of k, as described in the main text. The outcomes are grouped by the type of state they project onto, where X/Z indicates a single qubit measurement in that basis. In some cases, a specific entangling projection is performed on different sets of qubits with different probabilities, but it is represented as a single line in the table.

are $\lfloor n/2 \rfloor$ distinct types of circuits. $k = 0$ is used to denote the standard n-GHZ projection without extra Hadamards. Table 3.1 lists the variants of n-GHZ fusions of this kind for $3 \leq n \leq 6$ and the types of measurements they perform, calculated by performing Fock state simulations of the circuits and extracting the corresponding stabilizer operators (the simulation and analysis method is explained in more detail in Chapter 5).

The ability of some of the failure outcomes of these fusions to entangle the input states is a valuable resource which can in principle be used to generate resource states or perform computation with a

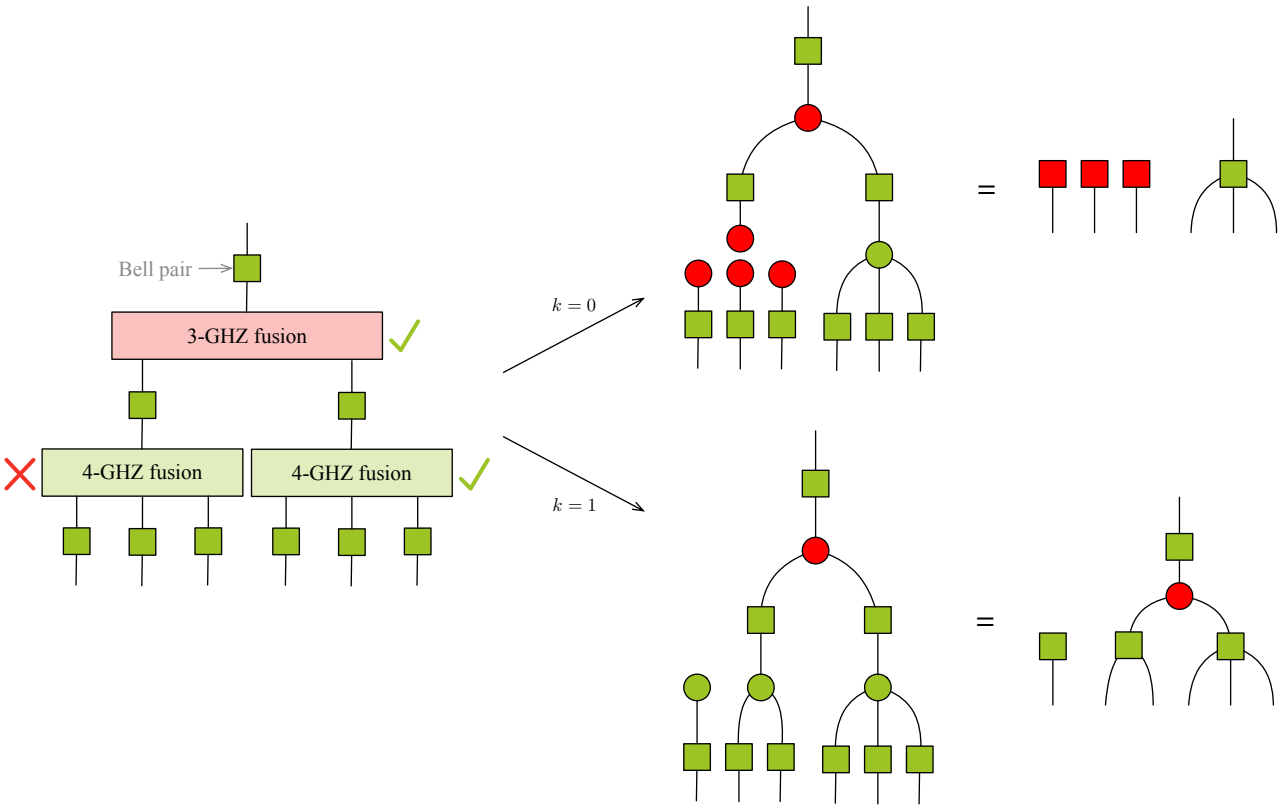

Figure 3.8: A scheme which generates a $(3Z, 2X)$ repetition code using Bell pair seed states and 3- and 4-GHZ fusions. The color of each fusion box in the diagram on the left reflects that of the corresponding success ZX spider. The ZX diagrams on the right hand-side show the encoding produced if one 4-GHZ fusion fails when the $k = 0$ (top) and $k = 1$ (bottom) version is used. Only the entangling failure outcome of the latter is represented. In the former case, the left branch of the tree is disentangled, resulting in a $(3Z, 1X)$ code. The entangling $k = 1$ fusion failure projects three of the qubits onto a 3-GHZ state, resulting in an irregular five-qubit repetition code which offers protection in both the Z and X bases.

fusion network more efficiently. As a simple illustrative example, consider a qubit with a tree local encoding protecting against erasure. One of the many ways to generate this kind of local code in LO-FBQC is to place an n-GHZ fusion operation at each internal n-valent node in the tree and a Bell pair on each edge, on the leaves and on the root. The code is obtained upon success of all fusions, and the color of the success ZX spider for each fusion, i.e., the presence or absence of Hadamard rotations on the inputs, determines the type of repetition at the corresponding level: green spiders give an $(n-1)$ repetition in Z and red ones in X. The diagram on the left in fig. 3.8 shows this generation scheme for a (2, 3)-encoded qubit. With standard n-GHZ fusions, if any of the failure outcomes are obtained, the inputs are disentangled and the corresponding branch is cut from the tree. However, if an entangling failure is obtained in one of the $k \neq 0$ fusions, the result can be just a reduction in the

branching factor which allows to retain more of the protection from the encoding. The right hand-side of fig. 3.8 shows this effect when one of the 4-GHZ fusions fails. If the $k = 0$ fusion is used, a code with no repetition in X is generated regardless of the failure mode. If the $k = 1$ fusion is used and the entangling failure outcome is obtained, the result is an asymmetric 5-qubit repetition code which still offers protection in both the Z and X basis. This is just one simple example of potential improvements to the efficiency and performance of an FBQC architecture afforded by the entangled failure outcomes of these n-GHZ fusion circuits. The investigation of other ways to make use of them and a quantitative assessment of their impact is left to future work.

3.3 Depth-2 multi-way fusions

The benefits of placing one pair of extra 50:50 beamsplitters at the input of an n-GHZ fusion lead to the question of whether adding more of them can improve the circuit further. By construction, with the addition of the two beamsplitters the depth of the fusion circuit becomes non-uniform: four modes cross two beasmplitters and the others just one. The structure can be rebalanced by also coupling the remaining fundamental beamsplitters in pairs. For the circuits in fig. 3.7, for example, this is done by adding two beamsplitters joining the pair $(2, 3)$ for $k = 1$ and $(2, 4)$ for $k = 2$, as shown in fig. 3.9. Note that in odd-n fusions this construction leads to one pair of modes encountering only one beamsplitter before detection, but we still include these kinds of circuits in our definition of depth-2 fusions for convenience. Each group of four coupled beamsplitters in this structure forms a special four-mode interferometer whose properties are at the heart of the operation of the *depth-2* fusions constructed in this way. We review this subcircuit next.

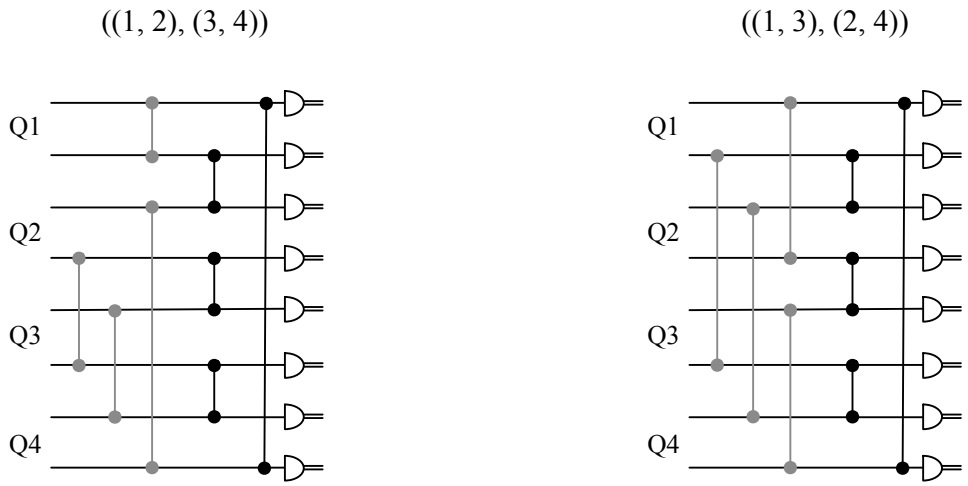

Figure 3.9: Depth-2 4-GHZ fusions obtained by adding two extra Hadamards to the $k = 1$ (left) and $k = 2$ (right) fusions. There is a third circuit with coupled beamsplitter pairs ((1, 4), (2, 3)), but its operation is equivalent to that of ((1, 2), (3, 4)) due to the cyclic symmetry of the n-GHZ fusion circuit. The outcomes of these fusions are reported in table 3.2.

3.3.1 The 4-Hadamard

The configuration of four beamsplitters which appears as a subcircuit of the depth-2 fusions is an interferometer with a real 4-Hadamard transfer matrix, up to permutations of rows and columns:

$$H_4 = \frac{1}{2}\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}. \tag{3.6}$$

We follow the graphical representation in [84] for this subcircuit, shown in fig. 3.10.

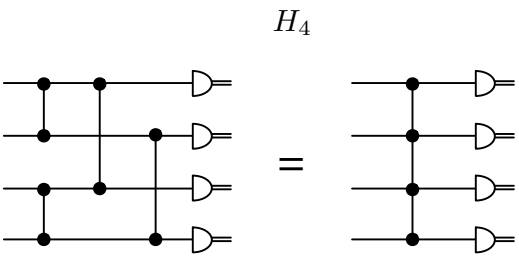

Figure 3.10: Graphical representation of the 4-Hadamard, H_4, subcircuit obtained when coupling two beamsplitters with another pair of beamsplitters.

Its action on the basis of four modes each occupied by zero or one photon consists of projections onto

six possible types of states, depending on the detection pattern obtained:

$$\begin{aligned}
&0 \text{ detected} \rightarrow |0000\rangle\rangle \\
&1 \text{ detected} \rightarrow |W^i_{4,1}\rangle\rangle \\
&2 \text{ detected (bunched)} \rightarrow |W^i_{4,2}\rangle\rangle \\
&2 \text{ detected (antibunched)} \rightarrow |\text{4-GHZ}^-_i\rangle\rangle \\
&3 \text{ detected} \rightarrow |W^i_{4,3}\rangle\rangle \\
&4 \text{ detected} \rightarrow |1111\rangle\rangle.
\end{aligned} \tag{3.7}$$

Here $|W_{N,n}\rangle\rangle$ are N-mode Dicke states [152] of the form:

$$|W_{N,n}\rangle\rangle = \sqrt{\binom{N}{n}} \sum_k \mathcal{P}_k |1\rangle\rangle^n |0\rangle\rangle^{N-n}, \tag{3.8}$$

where the sum is over all permutations $\mathcal{P}_k$ of the N modes which result in inequivalent terms. We refer to all these states as W-type states, based on the name given to the $|W_{N,1}\rangle\rangle$ state [153]. More precisely, the $|W_{4,n}\rangle\rangle$ states above are defined as:

$$\begin{aligned}
|W^i_{4,1}\rangle\rangle &= x_{i,1}|1000\rangle\rangle + x_{i,2}|0100\rangle\rangle + x_{i,3}|0010\rangle\rangle + x_{i,4}|0001\rangle\rangle, \\
|W^i_{4,2}\rangle\rangle &= y_{i,1}|0011\rangle\rangle + y_{i,2}|0101\rangle\rangle + y_{i,3}|0110\rangle\rangle + y_{i,4}|1001\rangle\rangle + y_{i,5}|1010\rangle\rangle + y_{i,6}|1100\rangle\rangle, \\
|W^i_{4,3}\rangle\rangle &= x_{i,1}|0111\rangle\rangle + x_{i,2}|1011\rangle\rangle + x_{i,3}|1101\rangle\rangle + x_{i,4}|1110\rangle\rangle
\end{aligned} \tag{3.9}$$

with coefficients:

$$x = H_4, \qquad y = \frac{1}{\sqrt{6}} \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & 1 & -1 & 1 \\ 1 & -1 & -1 & -1 & -1 & 1 \\ 1 & 1 & -1 & -1 & 1 & 1 \end{pmatrix}. \tag{3.10}$$

The 4-GHZ states are:

$$\begin{aligned}
|\text{4-GHZ}_1^-\rangle\rangle &= (|0011\rangle\rangle - |1100\rangle\rangle)/\sqrt{2}, \\
|\text{4-GHZ}_2^-\rangle\rangle &= (|1001\rangle\rangle - |0110\rangle\rangle)/\sqrt{2}, \\
|\text{4-GHZ}_3^-\rangle\rangle &= (|1010\rangle\rangle - |0101\rangle\rangle)/\sqrt{2}.
\end{aligned} \tag{3.11}$$

3.3.2 4-GHZ fusion with 4-Hadamards

The depth-2 4-GHZ fusions in fig. 3.9 can be seen as different combinations of two of these 4-Hadamard operations each acting on half the modes of the four input qubits. In these circuits, some of the two-photon projections from the two subcircuits combine to result in a successful 4-GHZ projection on the four input qubits. For example, consider the ((1, 3), (2, 4)) 4-GHZ fusion. Up to X rotations on two of the qubits, the circuit consists of one 4-Hadamard acting on all the '0' rails and one on all the '1' rails. Applying the first of these to a basis of four input qubits gives the Kraus operators:

$$\begin{aligned}
&0 \text{ detected} \to |1111\rangle\rangle\langle 1111|, \\
&1 \text{ detected} \to \sum_{w_{4,3}^i} |\overline{w_{4,3}^i}\rangle\rangle\langle w_{4,3}^i|, \\
&2 \text{ detected (bunched)} \to \sum_{w} |\overline{w_{4,2}^i}\rangle\rangle\langle w_{4,2}^i| \\
&2 \text{ detected (antibunched)} \to \sum_{\text{4-ghz}_i^-} |\overline{\text{4-ghz}_i^-}\rangle\rangle\langle \text{4-ghz}_i^-|, \\
&3 \text{ detected} \to \sum_{w_{4,1}^i} |\overline{w_{4,1}^i}\rangle\rangle\langle w_{4,1}^i|, \\
&4 \text{ detected} \to |0000\rangle\rangle\langle 0000|,
\end{aligned} \tag{3.12}$$

where $w_{4,j}^i$ and 4-ghz_i^- are the terms in the corresponding states in eq. 3.7 and $\overline{|\cdot\rangle}$ represents the complement of a term, e.g., $\overline{|1100\rangle} = |0011\rangle$. In these expressions, the kets are Fock-state operators

and the bras act on dual-rail qubits. When the second 4-Hadamard is applied, an outcome with $(4-n_d)$ detected photons is obtained, where n_d is the number of photons detected after the first. If $n_d \neq 2$, all the qubit terms projected onto by the first subcircuit survive in the second measurement, so a qubit projection of the same type as the corresponding single-rail operator in eq. 3.7 is performed. If $n_d = 2$, two different outcomes are possible. When both 4-Hadamards detect two bunched photons, all six terms in a $|W_{4,2}\rangle$ state survive and a qubit projection of this type is performed on the input qubits. On the other hand, if two antibunched photons are detected in either 4-Hadamard, only the two terms in a $|4\text{-GHZ}\rangle$-type state are left and a 4-GHZ projection is performed. The projection operators of the full fusion are:

$$\begin{aligned}
&(0,4) \text{ detected} \rightarrow |1111\rangle\langle 1111|\,,\\
&(1,3) \text{ detected} \rightarrow |W^i_{4,3}\rangle\langle W^i_{4,3}|\,,\\
&(2,2) \text{ detected (both bunched)} \rightarrow |W^i_{4,2}\rangle\langle W^i_{4,2}|\,,\\
&(2,2) \text{ detected (one or two antibunched)} \rightarrow |4\text{-GHZ}^{\pm}_i\rangle\langle 4\text{-GHZ}^{\pm}_i|\,,\\
&(3,1) \text{ detected} \rightarrow |W^i_{4,1}\rangle\langle W^i_{4,1}|\,,\\
&(4,0) \text{ detected} \rightarrow |0000\rangle\langle 0000|\,,
\end{aligned} \tag{3.13}$$

where $|4\text{-GHZ}^{+}_i\rangle$ are the states in eq. 3.11 with a $+$ sign between the two terms.

Assuming an unbiased four-qubit input basis, the probability of detecting two bunched photons after the first 4-Hadamard is $3/16$, and it is the same for two antibunched photons. Given that two photons were detected in the first 4-Hadamard, the other two give rise to a bunched or antibunched pattern after the second one with equal probability $1/2$. A 4-GHZ projection is performed when at least one of the two 4-Hadamards detects two antibunched photons, so the total success probability of this fusion is $\frac{3}{16}\times 1 + \frac{3}{16}\times\frac{1}{2} = \frac{9}{32} \sim 28\%$. This is more than double the success probability of other 4-GHZ fusions and it is also higher than that of a 3-GHZ fusion. So, unlike other linear optical boosting schemes,

this circuit achieves a significant increase in efficiency without the use of additional resources.

The main drawback of this fusion is that the majority of the failure outcomes, $19/32 \sim 59\%$, project onto a non-stabilizer W-type state. Stabilizer state projections allow to reuse the unmeasured qubits in subsequent measurement steps [106] and produce information which can be directly used for stabilizer fault-tolerance [43]. The states and information generated by non-stabilizer state projections, on the other hand, do not fit as naturally into the standard FBQC framework. So the increase in success probability comes at the cost of less desirable failure outcomes. Improved techniques to make use of non-stabilizer outcomes are likely necessary for this kind of circuit to be useful for ballistic fusions between resource states. On the other hand, this trade-off is likely to be favourable for RSG fusions which are multiplexed, since failed attempts can be discarded. The use of depth-2 fusions for RSG is discussed further in Chapter 6.

A similar kind of reasoning can be used to explain the outcomes of the ((1, 2), (3, 4)) 4-GHZ fusion, listed in table 3.2. In this case the success rate remains the same as the standard fusion, but the 4-Hadamard projection operators combine in such a way that all outcomes project onto pure stabilizer states. Entangling measurements are performed on some of the qubits upon failure with a probability $3\times$ higher than the $k = 1$ circuit. A 3-GHZ projection occurs at double the rate, and in addition a Bell state measurement is also performed some of the time.

3.3.3 n-GHZ fusion with 4-Hadamards

The construction of depth-2 circuits out of 4-Hadamard building blocks can be generalised to larger n-GHZ fusions. For odd n, the depth of the circuit cannot be uniform and two of the modes have to go through a single beamsplitter. For each n, there are $\frac{n!}{2^{\lfloor n/2 \rfloor} \cdot \lfloor n/2 \rfloor}$ circuits of this kind, which can be grouped into families whose members have equivalent outcomes up to swaps and X rotations of the input qubits. The families, with their success and non-stabilizer projection probabilities, are listed in table 3.2 for $4 \leq n \leq 8$, and the decomposition of the family representatives with highest success

n	Family representative	n-GHZ projection probability	Non-stabilizer projection probability
4	((1, 2), (3, 4))	1/8	0
	((1, 3), (2, 4))	9/32	19/32
5	((1, 2), (3, 4))	1/16	0
	((1, 2), (3, 5))	1/16	0
	((1, 3), (2, 4))	9/64	31/64
6	((1, 2), (3, 4), (5, 6))	1/32	0
	((1, 2), (3, 6), (4, 5))	1/32	0
	((1, 2), (3, 5), (4, 6))	9/128	31/64
	((1, 3), (2, 5), (4, 6))	23/256	77/256
7	((1, 2), (3, 4), (5, 6))	1/64	0
	((1, 2), (3, 4), (5, 7))	1/64	0
	((1, 2), (3, 6), (4, 5))	1/64	0
	((1, 2), (3, 7), (4, 6))	1/64	0
	((1, 3), (2, 4), (5, 7))	9/256	31/64
	((1, 2), (3, 5), (4, 6))	9/256	29/64
	((1, 2), (3, 5), (4, 7))	9/256	43/128
	((1, 3), (2, 5), (4, 6))	23/512	137/512
8	((1, 2), (3, 4), (5, 6), (7, 8))	1/128	0
	((1, 2), (3, 4), (5, 8), (6, 7))	1/128	0
	((1, 2), (3, 8), (4, 7), (5, 6))	1/128	0
	((1, 2), (3, 8), (4, 6), (5, 7))	9/512	31/64
	((1, 2), (3, 4), (5, 7), (6, 8))	9/512	29/64
	((1, 2), (3, 5), (4, 8), (6, 7))	9/512	43/128
	((1, 2), (3, 5), (4, 7), (6, 8))	23/1024	137/512
	((1, 3), (2, 6), (4, 8), (5, 7))	55/2048	361/2048
	((1, 3), (2, 5), (4, 7), (6, 8))	76/2048	33/64
	((1, 3), (2, 4), (5, 7), (6, 8))	81/2048	1447/2048

Table 3.2: Success and non-stabilizer projection probabilities for all depth-2 n-GHZ fusions for $4 \leq n \leq 8$, assuming n maximally mixed input qubits. For every n there is at least one family of circuits with higher n-GHZ projection probability than $1/2^{n-1}$ for the standard circuits, and the maximum relative boost increases with n, and the minimum non-stabilizer state projection probability decreases.

probability into H_4 building blocks is shown in fig. 3.11, up to X rotations on the input qubits such that as many of the 4-Hadamards as posssible only act on '0' or '1' rails.

The plots in fig. 3.12 summarize the key insights obtained from the data in table 3.2. For each n, there is at least one family with higher success probability compared to the standard n-GHZ fusion, and as n increases, so does the maximum relative boost in success probability. A trend in the way that the success probabilities increase with n can be obtained empirically from the data for the fusion sizes included in this analysis. For each fusion size n, there are different families of circuits which perform

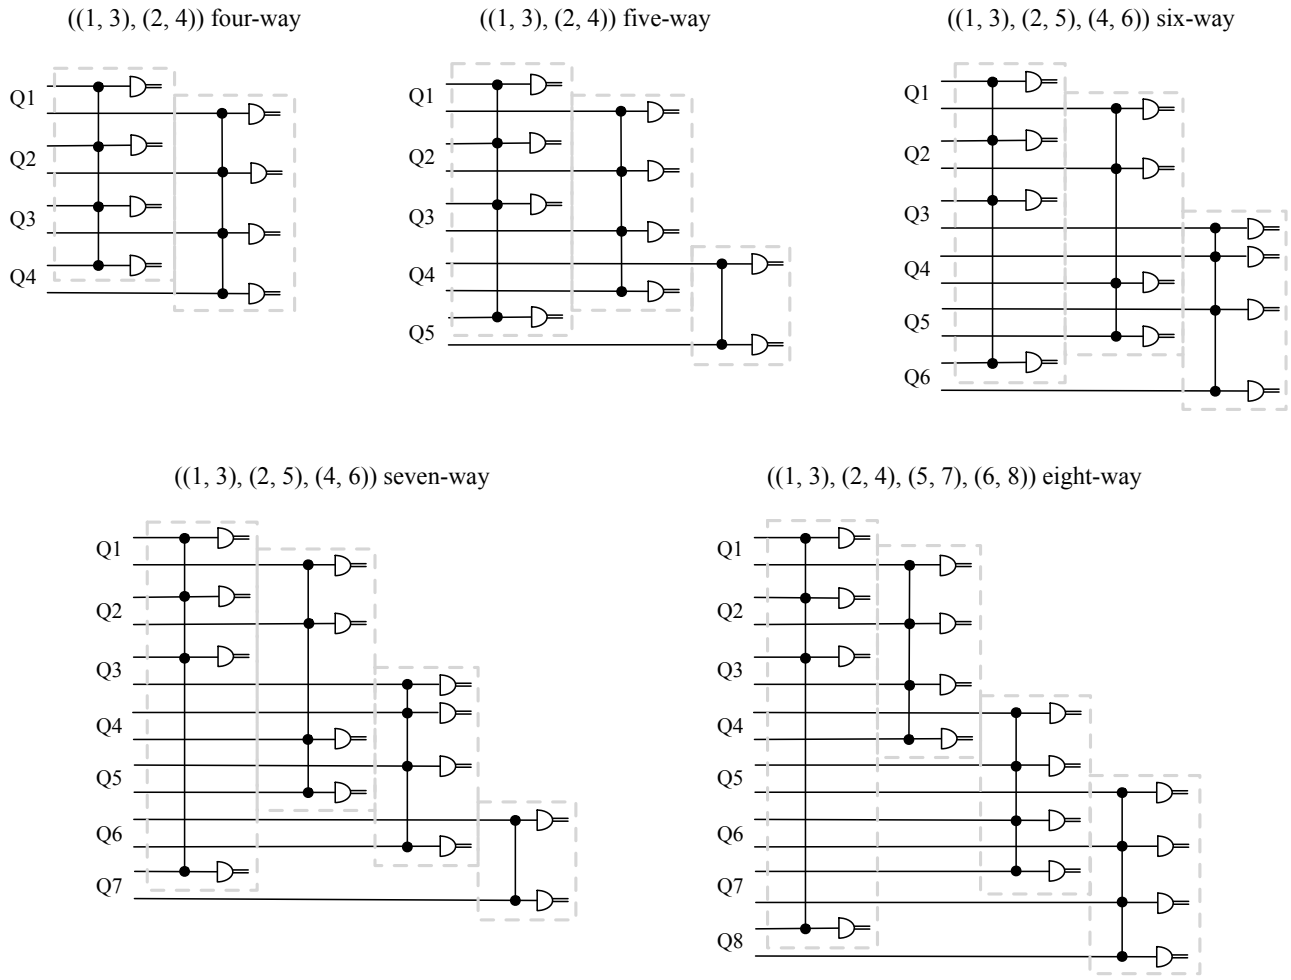

Figure 3.11: Decomposition of the representative circuits of the highest-probability depth-2 n-GHZ fusion families into H_4 interferometers. Fusions with odd n have two modes going through a 50:50 beamsplitter. Success outcomes involve two photons being detected after each H_4 and one after the H_2, where present.

an n-GHZ projection with probabilities of the form $x/2^{n+\frac{n_e}{2}-1}$, where n_e are all even numbers $n_e \leq n$ and the values of x depend on n_e: $x = 2$ for $n_e = 2$, $x = 9$ for $n_e = 4$, $x = 23$ for $n_e = 6$, and in the case of $n_e = 8$ there are three possible values $x \in \{55, 76, 81\}$. So, for example, for $n = 6$ we find families of fusions with success probabilities $1/32$ ($n_e = 2$), $9/128$ ($n_e = 4$), and $23/256$ ($n_e = 6$). For the given values of x, the scaling of the highest possible success probability is improved when going from an odd n to an even $(n+1)$, since this results in a reduction by less than a factor of 2. For example, for $n = 5$, the highest success probability is $9/64$ and for $n = 6$ it is $23/256$, i.e., only $\sim 1.6\times$ smaller. This mitigates the increase in resource cost associated with using larger fusions in FBQC, especially within the RSG where multiplexing can be used.

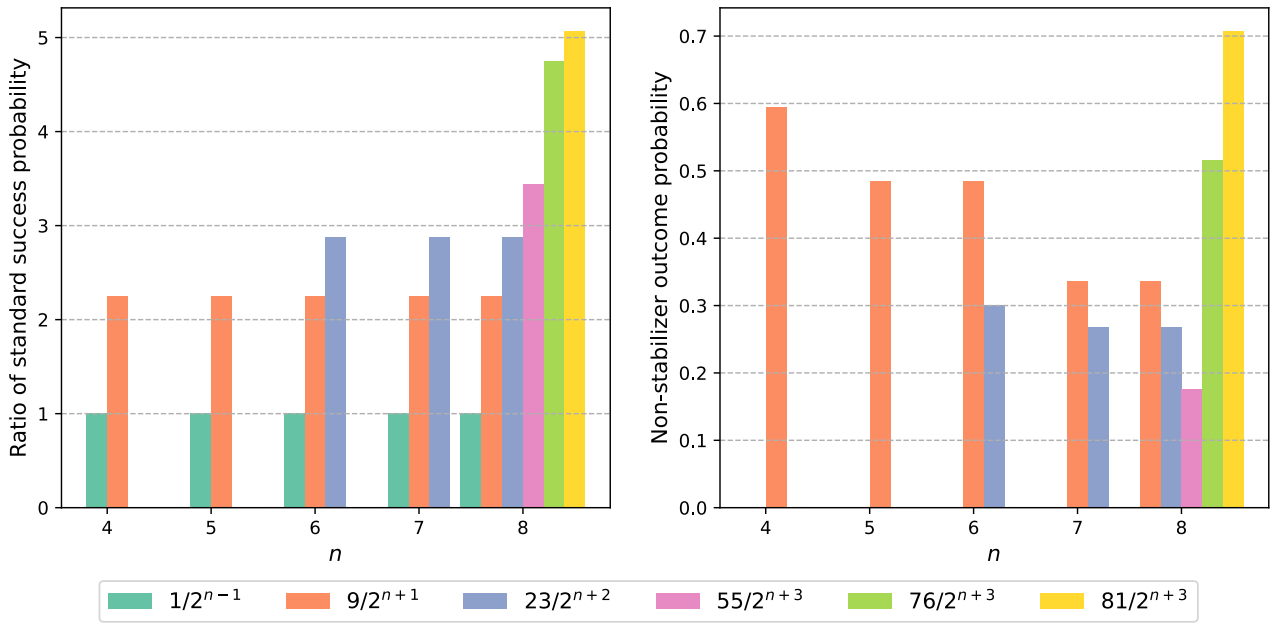

Figure 3.12: Comparison of the performance of different fusion families for $4 \leq n \leq 8$. Colors indicate circuit families, labelled by their n-GHZ projection probability. (Left) Ratio of the success probability to the standard $1/2^{n-1}$. As n increases, new circuits with a higher relative boost are found, with the maximum being $2.25\times$ for $n = 4$ and $5.0625\times$ for $n = 8$. (Right) Percentage of outcomes which project onto non-stabilizer states. Only the lowest possible rate is shown when multiple circuit families with the same success probability exist. As n increases, circuits with a decreasing ratio of non-stabilizer projections can be found, going from $\sim 59\%$ for $n = 4$ to $\sim 18\%$ for $n = 8$.

All these depth-2 fusions with increased success probability have some failure outcomes which project the inputs onto a non-stabilizer state. The probability of such outcomes varies for each family of fusions. As n increases, fusions with lower rates of non-stabilizer outcomes become available, with the lowest possible probability decreasing from 59% for $n = 4$ to 18% for $n = 8$. So, larger fusions have a higher proportion of stabilizer outcomes which can be used directly within the existing FBQC framework. However, the family with the lowest rate of non-stabilizer outcomes is not always the one with the highest success probability, so a choice has to be made to prioritize one aspect or the other when deisgning a scheme which uses these fusions.

In additon to the high success probability fusions discussed so far, for all n there are also depth-2 circuits with the standard $1/2^{n-1}$ success rate, but which only perform stabilizer state projections. In these fusions, the extra input Hadamards lead to additional entangling failure outcomes and multiple families with different sets of such operations exist. Table 3.3 describes these outcomes for circuits

n	Family representative	Projection type	Probability
4	((1, 2), (3, 4))	4-GHZ	$1/8$
		3-GHZ + X	$1/4$
		BP + X	$1/8$
		$2\times X + 2\times Z$	1/2
5	((1, 2), (3, 4))	5-GHZ	$1/32$
		4-GHZ + X	$1/8$
		3-GHZ + $2\times X$	$1/16$
		$X + 4\times Z$	$1/8$
		$2\times X + 3\times Z$	$1/8$
		$3\times X + 2\times Z$	$1/2$
	((1, 2), (3, 5))	5-GHZ	$1/32$
		4-GHZ + X	$1/16$
		3-GHZ + BP	$1/16$
		3-GHZ + $2\times Z$	$1/8$
		$2\times$ BP + X	$1/16$
		BP + $X + 2\times Z$	$3/8$
		$2\times X + 3\times Z$	$1/4$
6	((1, 2), (3, 4), (5, 6))	6-GHZ	$1/32$
		5-GHZ + X	3/32
		4-GHZ + $2\times X$	3/32
		3-GHZ + $3\times X$	1/32
		$2\times X + 4\times Z$	$1/4$
		$3\times X + 3\times Z$	$3/8$
	((1, 2), (3, 6), (4, 5))	6-GHZ	1/32
		5-GHZ + X	1/16
		4-GHZ + $2\times X$	1/32
		3-GHZ + 3-GHZ	1/32
		3-GHZ + BP + X	1/16
		3-GHZ + $X + Z$	1/4
		$2\times$BP + $2\times X$	$1/32$
		BP + $2\times X + 2\times Z$	1/4
		$2\times X + 4\times Z$	$1/4$

Table 3.3: Measurement outcomes of depth-2 n-GHZ fusion families with success probability $1/2^{n-1}$ for $4 \leq n \leq 6$. Fusions in these families only perform projections onto stabilizer states. All circuits have failure outcomes which entangle part of the input qubits, similarly to the fusions in table 3.1. Compared to those, these circuits produce some entanglement at a higher rate for a given n.

with $4 \leq n \leq 6$. It can be seen that these depth-2 fusions have a higher probability of generating partial entanglement compared to circuits of the same size with just two additional beamsplitters (c.f. table 3.1). For $n \geq 6$, fusions with higher success probability also have entangling stabilizer failure outcomes, which are listed in Appendix A.

In all these circuits, successful n-GHZ projections are heralded by the detection of two photons after

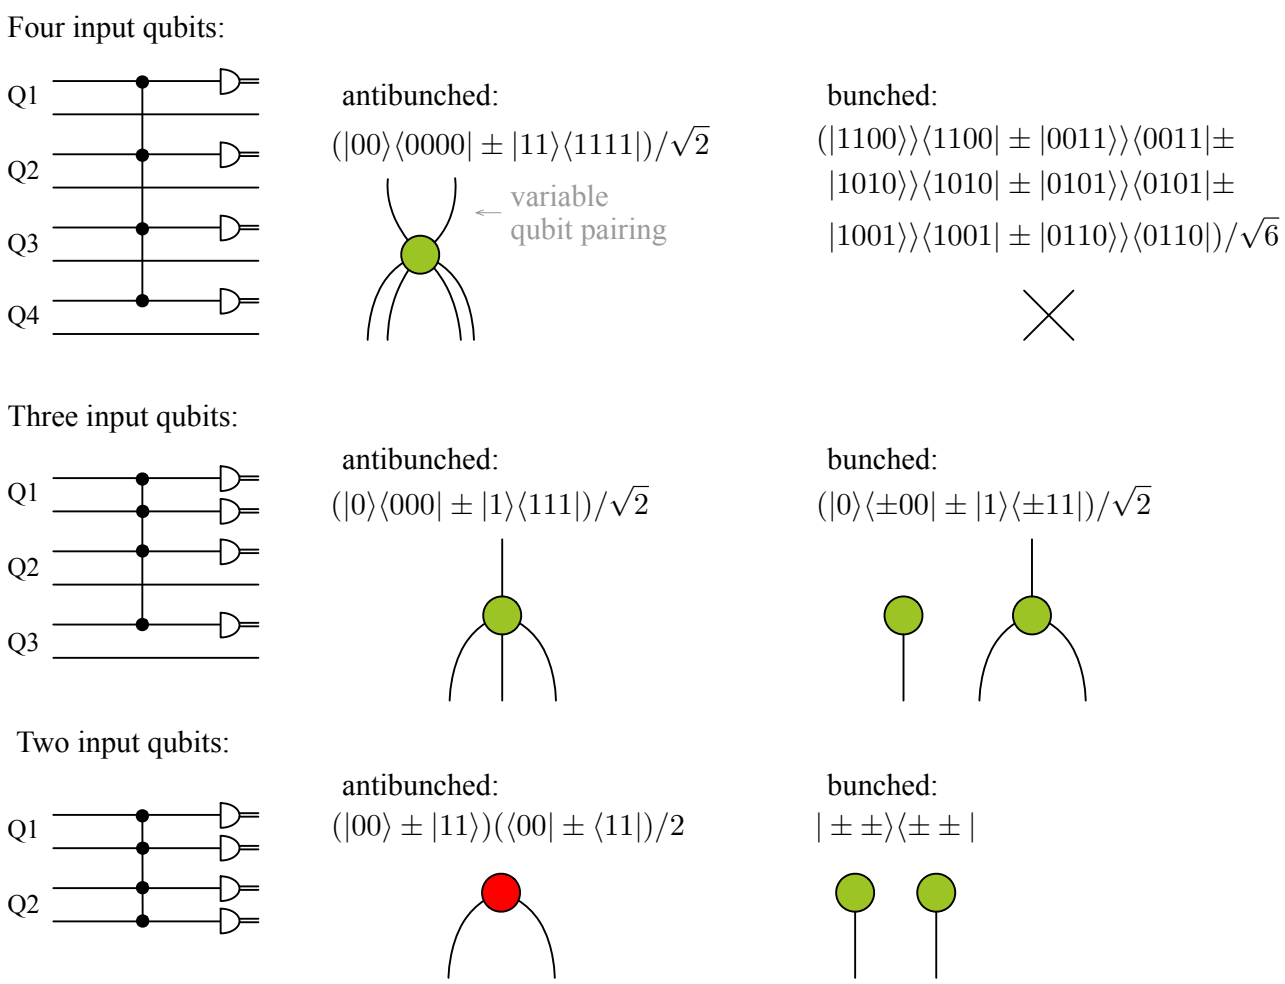

Figure 3.13: Two-photon detection outcomes of the 4-Hadamard applied to four, three and two input qubits. The operators and ZX spiders are represented up to X and Z rotations on the input and output qubits. The four and three input cases can be seen as analogous to a T1 fusion, since they perform partial measurements. When only two qubits are input, the circuit is simply an X-fail T2 fusion. Depth-2 n-GHZ fusions are constructed by concatenating 4-Hadamards of this form, up to swaps and X rotations of qubits, in different configuations.

each 4-Hadamard subcircuit, and one photon after the single beamsplitter in the case of odd n. Given the importance of these two-photon detection outcomes, it is helpful to understand the operations they perform when applied to different numbers of input qubits. Fig. 3.13 summarizes the operators and corresponding ZX-diagrams outcomes when the four input modes belong to four, three and two qubits. In the first two cases, the operation can be seen as analogous to a T1 fusion since it performs a partial measurement on the input qubits, while in the two-qubit case it is an X-fail T2 fusion. The success outcomes of depth-2 n-GHZ fusions can then be explained by combining these ZX-diagrams according to the structure of the circuits. In doing so, it is necessary to make use of the fact that the combination of a $W_{4,2}$-type and a GHZ-type measurement results in a GHZ-type measurement, as

seen in the case of the ((1, 3), (2, 4)) 4-GHZ fusion above. This means that the success outcomes of these fusions cannot be fully expressed in the ZX-diagram language, since the W-type measurements cannot be represented within this framework but still contribute to the overall success of the fusions. An extension of the language to $W_{4,2}$ spiders and the introduction of rules to combine these with ZX spiders is necessary to make a fully graphical representation possible.

3.4 Non-GHZ fusions

The success outcome of the fusions described so far is a projection onto an n-GHZ state. While these are in principle sufficient for FBQC, circuits which project onto non-equivalent fully connected graph states for $n > 3$ can allow more flexibility, for example in RSG design. In particular, here we focus on the n-qubit linear cluster, or n-line. For $n = 4$ this is a representative of the only other class of fully connected graph states which are not locally equivalent to a 4-GHZ.

3.4.1 n-line fusions with extra beamsplitters

As previously discussed, a circuit which projects onto an n-line can be obtained by adding Hadamards at the input and between the constituent T1 and T2 fusions of the standard n-GHZ analyzer. Removing the mode swaps in fig. 3.4, results in the representation of the circuit in fig. 3.14a. From this it is clear that the additional beamsplitters are placed such that one of the fundamental beamsplitters is coupled to all the other ones, i.e. $((1,2),(1,3),(1,4))$. In this case, the first two couplings are realised through the top ports of the beamsplitters, and the third through the bottom ports. Flipping this last one so that the top ports of $(1,4)$ are coupled is found to give another valid 4-line fusion circuit, shown in fig. 3.14b, which has the same success probability of $1/8$, but different properties. In particular, the circuit derived from the tree of T1 and T2 fusions only projects onto pure stabilizer states and it has one entangling failure outcome consisting of a projection of qubits 3 and 4 into a 2-qubit linear cluster, which happens with probability $3/8 = 37.5\%$. The same outcome occurs in the other 4-line fusion at

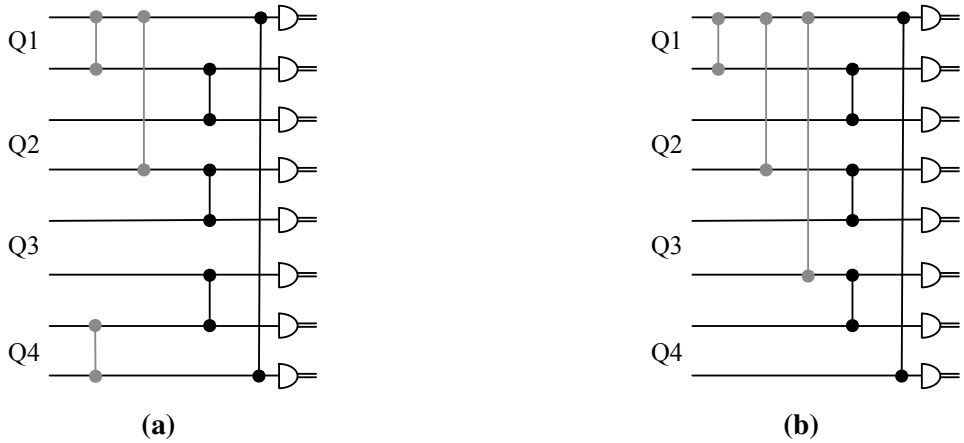

Figure 3.14: 4-line fusion circuits with extra beamsplitters (grey) added to the 4-GHZ analyzer (black). (a) Circuit obtained from the decomposition of the 4-GHZ fusion into a tree with two T1s feeding into a T2. The top mode has the maximum beamsplitter depth of 3. (b) Alternative circuit obtained by moving the bottom extra beamsplitter to couple the top ports of the same two fundamental beamsplitters. The maximum depth here is 4.

a lower rate of $1/4 = 25\%$, and in addition a projection onto a 3-line across qubits 2, 3 and 4 is done with probability $1/8 = 12.5\%$. This circuit, however, also has a $\sim 6\%$ probability of projecting onto a non-stabilizer state.

A general construction for n-line fusions with arbitrary n can be obtained from this $n = 4$ example. Starting from the n-GHZ analyzer, $n-1$ extra beamsplitters are added to couple one port of one of the fundamental beamsplitters, say 1, to the same port of all the other ones starting from the furthest mode and proceeding in order to the closest one. This gives a circuit where one of the modes crosses n beamsplitters before detection, like the 4-line fusion in fig. 3.14b. From this, other valid n-line projections can be obtained by "flipping" some of the extra beamsplitters so that they now couple the other two ports of the same fundamental beamsplitter. Starting closest to the n-GHZ analyzer and doing this to one additional extra beamsplitter each time in order gives $n-1$ n-line fusion circuits. Of these, only the $\lceil n/2 \rceil$ which have different maximum depths $\lfloor n/2 \rfloor + 1 \leq d_{\max} \leq n$ are not equivalent up to permutations of modes, due to the symmetry of the underlying n-GHZ analyzer.

All these circuits have the same success probability of $1/2^{n-1}$, but they have different entangling failure outcomes and probabilities of projecting onto a non-stabilizer state. The outcomes of these n-line fusions for $4 \leq n \leq 6$ are listed in table 3.4, where the circuits are identified by their $d_{\max}$ value. For

n	$d_{\max}$	Ent. fail prob	Non-stab. prob
4	2	2-line = 3/8	0%
		Total = 37.5%	
	3	2-line = 1/4 3-line = 1/8	5.9%
		Total = 37.5%	
5	3	2-line = 1/2 2×2-line = 3/16	0%
		Total ∼ 68.8%	
	4	2-line = 39/128 3-line = 3/16	2.9%
		Total ∼ 49.2%	
	5	2-line = 277/1024 3-line = 1/8 4-line = 1/16	10.2%
		Total ∼ 45.8%	
6	4	2-line = 341/1024 3-line = 39/256 2×2-line = 1/8 2-line + 3-line = 3/32	1.5%
		Total ∼ 70.4%	
	5	2-line = 481/2048 3-line = 307/2048 4-line = 3/32 2×2-line = 7/256	8%
		Total ∼ 50.6%	
	6	2-line = 3919/16384 3-line = 277/2048 4-line = 1/16 5-line = 1/32 2×2-line = 51/2048	14%
		Total ∼ 49.3%	

Table 3.4: Summary of entangling and non-stabilizer state projection outcomes for n-line projection circuits with different maximum depth $d_{\max}$ for $4 \leq n \leq 6$. The total probability of each type of entangling failure outcome is reported, but this may consist of projections of different subsets of qubits onto the given state.

$n > 4$, the total probability of an entangling failure outcome decreases with $d_{\max}$, while the probability of a non-stabilizer state projection increases with it. Interestingly, for $n = 6$ none of the fusions can avoid non-stabilizer failure outcomes, whereas this is possible for $n < 6$. This points to the circuits with smallest $d_{\max} = \lfloor n/2 \rfloor + 1$ being the most appealing to use in an FBQC protocol, since they maximise the amount of potentially useful entanglement generated while minimizing non-stabilizer outcome rates. Higher-$d_{\max}$ n-line fusions, however, have the advantage that some of their failures entangle a larger subset of the qubits, which can be desirable in specific schemes. The decision of which of these line fusions to use will ultimately depend on the details of the scheme being considered.

3.4.2 4-line fusion with 4-DFTs

4-Hadamard building blocks were found to give rise to new n-GHZ fusion circuits with improved success probability and rate of entangling failure outcomes. Therefore, it is natural to ask whether a similar construction exists for n-line projections using a different primitive subcircuit. H_4 belongs to the class of complex Hadamard matrices, defined as square orthogonal matrices whose elements

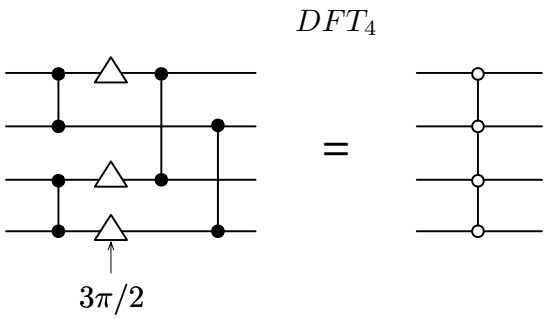

Figure 3.15: Realisation of the DFT_4 matrix as a linear optical circuit with 50:50 beamsplitters and passive $3\pi/2$ phase shifters. The notation on the right-hand side is analogous to the one used for the H_4 in fig. 3.10.

have unit modulus. Matrices of this type and their properties have been studied since the late 1800's [154, 155] and while the problem of identifying all classes of matrices which are not equivalent up to phases and permutations of rows and columns for all sizes is still open, a categorisation of all known examples is available [156]. For size 4 there is a continuous family of inequivalent complex Hadamard matrices specified by one parameter, $a \in [0, \pi)$ [157]:

$$F_4(a) = \frac{1}{2}\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & ie^{ia} & -1 & -ie^{ia} \\ 1 & -1 & 1 & -1 \\ 1 & -ie^{ia} & -1 & ie^{ia} \end{pmatrix}. \tag{3.14}$$

Up to a swap of rows, $F(\pi/2)$ is H_4 and $F(0)$ is the four-mode Fourier matrix:

$$DFT_4 = \frac{1}{2}\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & -1 & i \\ 1 & -1 & 1 & -1 \\ 1 & i & -1 & -i \end{pmatrix}, \tag{3.15}$$

whose implementation as a linear optical circuit is shown in fig. 3.15. We note that this matrix is equivalent to the one used in Ref. [84] to generate two-qubit linear cluster states.

When DFT_4 followed by detection is applied to a uniform basis of four single rails occupied by either zero or one photon each, the resulting projection operators are similar to those in eq. 3.7 when 0, 1, 3 or 4 photons are detected, except that in this case the matrix of coefficients in eq. 3.9 is $x = DFT_4$. When two photons are detected, on the other hand, three possible types of operations take place. If the photons bunch at one detector, the input rails are projected onto a state similar to $|W_{4,2}\rangle\rangle$ in eq. 3.9 with additional phases in the $y_{i,j}$. In the antibunched case, the two patterns (1, 0, 1, 0) and (0, 1, 0, 1) correspond to a projection onto one of the single-rail 4-GHZ states $(|1010\rangle\rangle \pm |0101\rangle\rangle)/\sqrt{2}$. We refer to these as the "ghz" patterns. The other four patterns, project onto states of the form:

$$|L^{\pm}\rangle\rangle = \frac{1}{2}\left[(|0011\rangle\rangle - |1100\rangle\rangle) \pm e^{i\pi/2}(|1001\rangle\rangle + |0110\rangle\rangle)\right], \tag{3.16}$$

which are equivalent to single-rail 4-lines on rails 2–0–1–3 up to local Cliffords H and S. For this reason, we call these "line" detection patterns.

When two of these are applied in parallel to the '0' and '1' rails of four dual-rail input qubits, as in the ((1, 3), (2, 4)) depth-2 4-GHZ fusion, the following projection operators are obtained:

$$\begin{aligned}
&(0,4) \text{ detected} \to |1111\rangle\langle 1111|, \\
&(1,3) \text{ detected} \to |\tilde{W}^i_{4,3}\rangle\langle \tilde{W}^i_{4,3}|, \\
&(2,2) \text{ detected (both bunched)} \to |\tilde{W}^i_{4,2}\rangle\langle \tilde{W}^i_{4,2}|, \\
&(2,2) \text{ detected (ghz + ghz or bunched)} \to |\text{4-GHZ}^{\pm}_i\rangle\langle \text{4-GHZ}^{\pm}_i|, \\
&(2,2) \text{ detected (line + bunched)} \to |L^{\pm}\rangle\langle L^{\pm}|, \\
&(2,2) \text{ detected (line + line)} \to |\psi^{\pm}\rangle\langle\psi^{\pm}|_{(1,3)} \otimes |\psi^{\pm}\rangle\langle\psi^{\pm}|_{(2,4)}, \\
&(3,1) \text{ detected} \to |\tilde{W}^i_{4,1}\rangle\langle \tilde{W}^i_{4,1}|, \\
&(4,0) \text{ detected} \to |0000\rangle\langle 0000|,
\end{aligned} \tag{3.17}$$

where these states are the dual-rail versions of the single-rail ones described above, the $|\tilde{W}^i_{4,x}\rangle$ indicate

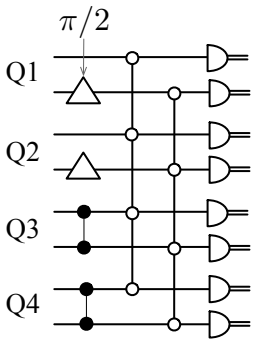

Figure 3.16: 4-line fusion circuit constructed out of two parallel DFT_4 circuits acting on the '0' and '1' rails of the input qubits, preceded by local H and S Cliffords. This fusion projects onto a 4-line across qubits 2–0–1–3 with the standard probability $1/2^{4-1} = 1/8$. It also has two failure outcomes which project onto entangled pure stabilizer states, a state locally equivalent to a 4-GHZ or two Bell pairs.

states similar to the ones in eq. 3.9 with updated coefficients, and $|\psi^{\pm}\rangle_{(i,j)}$ is a Bell pair across qubits (i, j) and the Bell pairs on the two sets of qubits are the same.

In this circuit, the way the single-rail operators combine gives rise to three distinct entangled stabilizer state projections when two photons are detected after each DFT_4. Similarly to the H_4 case, when two of the photons bunch at one detector and the other two do not, the overall projection operator is determined by the antibunched pattern, which filters out terms in the $|\tilde{W}^i_{4,2}\rangle$ state to give a projection onto a 4-GHZ or an $|L\rangle$ state with probabilities $3/32 \sim 9.4\%$ and $1/8 = 12.5\%$ respectively Interestingly, when two "line" patterns are obtained, cancellation of terms in the single-rail operators leads to a projection onto two separate Bell pairs with probability $1/16 \sim 6.3\%$. The $|\tilde{W}^i_{4,2}\rangle$ outcomes occur with probability $3/32 \sim 9.4\%$, the $|\tilde{W}^i_{4,1}\rangle$ and $|\tilde{W}^i_{4,3}\rangle$ ones have a cumulative probability of 50%, and the product state projections $1/8 = 12.5\%$. So, applying the local H and Z operations at the input of the fusion gives an alternative circuit, shown in fig. 3.16, to project onto a four-line state with the usual $1/2^{4-1} = 1/8$ probability.

While the 4-Hadamard allows to design n-GHZ fusion circuits with favourable properties, including significantly higher success probability and more entangling failure outcomes, the DFT_4 does not seem to lead to the same improvements for line projections. The main advantage of this 4-line fusion is the fact that it has a failure outcome which entangles all four input qubits, which has not been

observed in other circuits. However, since the success probability is unchanged and the rate of non-stabilizer failure outcomes is significant, it is not clear that this circuit would benefit the performance of an FBQC scheme overall. An indication that this subcircuit may hold more potential for fusion is given by the observation that it can also be used to perform partially destructive fusion, analogous to a T1 and to the application of a single H_4 circuit to dual-rail qubits described in fig. 3.13. In fact, when the DFT_4 circuit followed by detection is applied to half the rails of four maximally mixed input qubits, two outcomes are obtained such that two qubits are defined on the four output rails. With probability $1/16 \sim 6.3\%$ the Kraus operator corresponds to a 6-GHZ across the four input and two output qubits, similar to the antibunched outcome in fig. 3.13, and with probability $1/8 = 12.5\%$ it corresponds to a 6-line across the qubits up to local H and S operations. Since this thesis focuses on fully destructive fusions, we leave the investigation of this use of the DFT_4 circuit to future work.

3.5 Conclusion

The generalization of the T2 fusion to multiple input qubits is the GHZ-state analyzer, a fully-destructive linear optical circuit which projects its n inputs onto a GHZ state upon success and measures them individually in the computational basis the rest of the time. We refer to this as a "standard" n-GHZ fusion. The success probability of this type of fusion is $1/2^{n-1}$, which decreases exponentially with n, making large fusions very expensive in terms of resource cost.

Starting from the GHZ-state analyzer circuit, we constructed novel fusions which allow to perform the same operation but achieve a higher rate of entanglement generation. By adding only two 50:50 beamsplitters between pairs on input modes, families of circuits are found with failure outcomes which project subsets of inputs onto GHZ states. The success probability of these fusions is the same as for the GHZ-state analyzer, but useful entanglement is generated even in the event of failure. The addition of more extra beamsplitters gives more n-GHZ fusions with entangling failure outcomes, as

well as families of circuits which achieve significantly higher success probabilities, e.g., a $2.25\times$ boost for $n = 4$, without using additional ancillary resources. A depth-2 fusion circuit was also described which performs a projection onto either a 4-GHZ state or a state locally equivalent to a four-qubit linear cluster with total probability $7/32$. The increased entanglement generation rate of these new fusion circuits substantially reduces the otherwise prohibitive cost of multi-way fusions, opening a path for larger fusions to be used in the design of RSG schemes and fusion networks with decreased resource requirements. The impact of some of these fusions on the footprint of RSG schemes in particular is analysed in Chapter 6.

The ability of the linear optical fusions described in this chapter to perform entangling operations with higher probability than previously known circuits naturally leads to the question of whether further improvements are possible. In general, limits for the success probability of entangling linear optical circuits have only been established for a few specific examples including Bell state measurements [123, 129, 130], the non-linear sign-shift gate [158, 159], and the controlled-phase (CZ) gate [160]. In this work, we have focused on a constructive approach to find new fusion circuits with improved entanglement generation properties, for which upper bounds on the achievable success probability do not need to be known. An investigation of the success and entanglement generation limits for multi-way fusions would be an interesting and valuable extension of this work, with the potential to shed light on further improvements to the constructions described here.

Chapter 4

Modelling Errors in Standard n-GHZ Fusions

This chapter describes the dominant sources of noise which contribute to logical errors in a linear optical fusion-based quantum computing architecture: photon loss, distinguishability, and multiphoton contamination. Known methods to model the impact of each of these errors on dual-rail photonic qubits are described, from which Pauli error models for two-way fusions are derived. A novel framework is introduced which allows to compose these into Pauli error models for larger multi-way fusions. The results obtained for fusions on up to five qubits are presented and used to gain insight into the scaling of errors due to different physical mechanisms with fusion size.

So far the description of linear optical fusions has been limited to the case of ideal input qubits being acted on by a perfect linear optical circuit. In practice, however, every qubit and operation is affected by physical noise. Since all qubits are eventually measured by fusions, it is important to understand how noise affects their measurement outcomes, and to include its effect in any evaluation of the performance of a realistic FBQC architecture. Fusion measurement errors due to noise can be heralded or unheralded, with the former causing erasure of the measurement outcomes, and the latter turning into Pauli errors. This and the next two chapters focus on unheralded errors, since these cannot be discarded via multiplexing in the RSG process. Erasure errors are considered separately in Chapter 7.

The methods presented here aim to provide a description of the impact of the main sources of unheralded errors in the form of a probabilistic and stochastic Pauli error model on the outcomes of fully-destructive fusions. In particular, photon loss, distinguishability and multiphoton contamination are considered. The effect of each of the first two on linear optical entanglement generation has been studied in isolation [161–165], but here we set out to obtain a comprehensive model of errors on linear optical fusions which simultaneously includes all three error mechanisms. A framework is introduced to map physical errors into Pauli error models on the input qubits for the success outcomes of standard n-GHZ fusions. Models to describe the effect of photon loss, distinguishability and multiphoton contamination on dual-rail photon qubits are presented and used to derive the impact of these errors on the operation of T1 and T2 fusions. Then, the decomposition of standard n-GHZ fusions into constituent two-way fusions presented in Chapter 3 is used to obtain Pauli error models for them via composition. Models for fusions with up to five input qubits are derived using this framework and compared to obtain an insight into the scaling of errors due to different physical noise mechanisms with fusion size.

4.1 From physical noise to qubit errors

As discussed in Chapter 2, errors on dual-rail qubits undergoing a linear optical fusion can be classified as *heralded* or *unheralded*. In our analysis, we are interested in understanding the impact of these error mechanisms on fusion measurements to leading order, so we consider models which capture the occurrence of a single one of the types of errors described above on the input qubits, e.g., the loss of a single photon, one of the inputs being distinguishable, and so on. In reality, many more hardware imperfections and combinations of errors than the ones considered here will be present in a real-world machine. In the silicon photonic platform, a few examples are variability in the reflectivity of the directional couplers, differences in the length of optical fibres or on-chip waveguides, instability of the phases applied by active phase shifters, and many more [69]. Technological advances in the de-

sign and fabrication of hardware devices in recent years have significantly improved the performance specifications and reliability of many linear optical components, making many of these contributions increasingly less impactful [77]. Nonetheless, a truly accurate model of a realistic LO-FBQC architecture can only be obtained by taking into account all these sources of error, which is a very complex and computationally intensive task. By restricting the analysis to the dominant error mechanisms, we can obtain valuable insight into the expected behaviour of fusion circuits and more easily build an intuition to guide the design of improved quantum computing architectures.

4.1.1 Loss

Loss is the main type of error which affects photonic qubits, mostly caused by the absorption or scattering of photons as they propagate through the system, as well as inefficiency of the detectors. While every component along the path of a photon can introduce loss, the main contributors are active phase shifters, photon detectors and chip-fibre interfaces [90].

A common way to model the effect of loss is to imagine that a perfect beamsplitter couples the photon rail to a vacuum ancilla mode which is then traced out [166, 167]. The reflectivity of the beamsplitter captures the loss rate. In this picture, the action of the loss channel on a state with a fixed total number of photons, n, is given by:

$$\rho \to L(\rho) = \sum_{k=0}^{n} \binom{n}{k} l^k (1-l)^{n-k} \rho_k, \tag{4.1}$$

where k is the total number of lost photons, l is the loss rate, and ρ_k is the state obtained by projecting the ancilla modes onto the k-photon subspace and tracing them out.

This channel describes *balanced* loss across all the rails the state ρ is defined on. One of its properties is that it commutes with linear unitary operations on the rails, so any loss due to components seen by all rails in a linear optical circuit can be considered to act on its inputs or outputs [167]. This means that all balanced losses encountered by a qubit before or within a fusion circuit can be commuted to its

input, and the circuit can then be modelled as lossless. In an LO-FBQC architecture, this accurately captures all of the largest sources of loss listed above. The depth of other passive elements, such as directional couplers, may not be uniform in some circuits. However, their contribution to the overall loss is significantly smaller [77] and in practice it may be preferable to artificially balance their losses in fabrication anyway. Photon loss due to imperfect photon number resolution at the detectors can be captured by the model in eq. 4.1 under the assumption that photon-number-resolving detectors are implemented by splitting the input photons between a large number of separate detectors, such that the probability of any of them receiving more than one photon is vanishingly low [168]. Moreover, these kinds of errors are not relevant to many of the success detection patterns considered in our analysis, which consist of single photons being detected separately.

Another simplification in a leading-order analysis is to only consider events in which one of the photons entering a fusion is lost. Consider a generic m-qubit pure state $|\psi\rangle = a_0|A_0\rangle|10\rangle\rangle_i + a_1|A_1\rangle|01\rangle\rangle_i$, where the state of qubit i has been explicitly written out in Fock state notation. If the probability of that qubit being lost is p_l, the state $\rho = |\psi\rangle\langle\psi|$ is transformed as:

$$\rho \to L_i(\rho) = (1-p_l)\rho + p_l \times \rho_{\text{loss}}, \tag{4.2}$$

where the second term corresponds to the event that the photon is lost and the resulting state is:

$$\rho_{\text{loss}} = \frac{1}{2}\left(|a_0|^2|A_0\rangle\langle A_0| + |a_1|^2|A_1\rangle\langle A_1|\right) \otimes |00\rangle\rangle\langle\langle 00|. \tag{4.3}$$

So, if a loss event takes place, the affected qubit is left in a vacuum state and the rest of the state is the one obtained by tracing out the lost qubit.

4.1.2 Distinguishability

Linear optical fusions rely on HOM interference of photons to perform entangling measurements on the input qubits. The internal states of the photons influence the quality of the measurement, since ideal interference can only take place between photons which are fully indistinguishable in all degrees of freedom except the one used to define the qubit state.[1] In any physical realisation of multiphoton interference, distinguishability between photons can be induced by imperfections in the sources used to generate them or in the linear optical circuits along their paths. Heralded single photon sources based on SFWM produce pairs of photons with entangled joint spectral amplitude, resulting in non-unit purity of the signal photon upon detection of the herald. In addition, the wavelength and spectral shape of photons emitted by different sources cannot be made to match perfectly. Techniques to mitigate these effects have been proposed and experimentally demonstrated, allowing to achieve on-chip HOM visibilities of 99.5% [77]. Beyond the source, noise in the linear optical circuits used to propagate and interfere the photons can also alter their internal state, e.g., dispersion in waveguides or length mismatch in optical fibres, causing distinguishability. However, the magnitude of these effects is significantly smaller than source imperfections [171]. Because of this, in our leading-order analysis we consider a single distinguishability rate for all photons regardless of their worldline, under the assumption that differences between their paths have a negligible impact.

In order to model distinguishability in a dual-rail qubit, we express its state as $\rho_{\mathrm{L}} \otimes \rho_{\mathrm{I}}$, where the two terms represent the logical (L) and internal (I) states respectively. The former is an arbitrary qubit state, and the latter can be any state in the Hilbert space spanned by a "good" state $|g\rangle$, i.e., the reference internal state targeted for every single photon, and an orthogonal "bad" state $|b_i\rangle$, such that $\langle g|b_i\rangle = 0$ for each qubit i. Since detectors cannot resolve the internal state of the photons and

[1] This holds in our model, which assumes that photons are produced by independent single photon sources, and therefore have separable internal states, and that detectors do not have access to the internal state of the photons, as described in Chapter 2. If these assumptions are broken, e.g., by considering two photons whose internal state is not separable [169], or detectors with the capability to resolve the internal state [170], it is possible for photons which are not fully indistinguishable to give rise to perfect interference.

orthogonal components do not interfere, coherences in the state ρ_{I} do not have an observable effect, so we can use the following diagonal mixture to describe the internal photonic state:

$$\rho_{\mathrm{I}} = p_g |g\rangle\langle g| + p_b |b_i\rangle\langle b_i| \,, \tag{4.4}$$

where p_g and p_b depend on the rate of distinguishability. Since we are interested in leading-order error mechanisms, we consider the scenario of only one of the input qubits of a fusion being distinguishable, i.e., having $p_b \neq 0$ above. In this case, a single "bad" state $|b\rangle$ is sufficient since two distinguishable photons never enter the circuit at the same time. Although it is not part of our work, the scenario of multiple distinguishable photons has been studied in detail, particularly under the so-called Orthogonal Bad Bits Model (OBBM), which assumes that the $|b_i\rangle$ for different photons are orthogonal to each other, i.e., $\langle b_i|b_j\rangle = \delta_{ij}$ [162, 163].

Given the form of eq. 4.4, it is sufficient to consider the "good" and "bad" parts of the state separately. The former does not alter the action of a fusion circuit, so to understand the impact of distinguishability, only the latter needs to be modelled. This can be achieved by applying the following channel to one of the input qubits of the fusion:

$$\begin{aligned} \rho = \rho_{\mathrm{L}} \otimes \rho_{\mathrm{I}} \to D(\rho) = &(\mathbb{I} \otimes (|b\rangle\langle g|)\,(\rho_{\mathrm{L}} \otimes \rho_{\mathrm{I}})\,(\mathbb{I} \otimes (|g\rangle\langle b|)) + \\ &(\mathbb{I} \otimes (|b\rangle\langle b|)\,(\rho_{\mathrm{L}} \otimes \rho_{\mathrm{I}})\,(\mathbb{I} \otimes (|b\rangle\langle b|))\,. \end{aligned} \tag{4.5}$$

This maps the internal state of the qubit to $|b\rangle\langle b|$, so that it is fully distinguishable from the other fusion inputs.

4.1.3 Multiphoton contamination

There are several mechanisms which can result in the two rails of a dual-rail qubit containing more than one photon, an error referred to as *multiphoton contamination*. One of the dominant causes of

this is loss or imperfect photon number resolution in the heralding arm of the single photon source, which causes to mistake the emission of multiple photons for that of a single photon [90]. Beyond the single photon source, the heralding pattern in a Seed State Generation (SSG) circuit can be similarly misinterpreted as a result of dark counts at the detectors, so that more photons than expected are output [134]. Photons can also be added to a qubit rail within a multiplexing network, for instance, due to faulty components causing imperfect interference.

Two different models of multiphoton contamination on dual-rail qubits are considered in this work. They both only account for first-order error events, which result in two photons being present in the qubit rails. The simpler *maximally-mixed model* is used for the analytical derivation of error channels for standard n-GHZ fusions in this chapter, while the results presented in the following chapters are based on the more realistic *photon addition model*.

Maximally-mixed model

Under this model, a dual-rail qubit affected by multiphoton contamination is traced out and replaced by a uniform mixture of the three two-photon terms $\{|20\rangle\rangle, |02\rangle\rangle, |11\rangle\rangle\}$. This corresponds to the following channel on an ideal single qubit state:

$$\rho_1 \to M(\rho_1) = \frac{1}{3}\left(|20\rangle\rangle\langle\langle 20| + |02\rangle\rangle\langle\langle 02| + |11\rangle\rangle\langle\langle 11|\right). \tag{4.6}$$

Consider the multi-qubit pure state $|\psi\rangle = a_0|A_0\rangle|10\rangle\rangle + a_1|A_1\rangle|01\rangle\rangle$. If multiphoton contamination occurs on qubit i, whose state is explicitly written out in Fock notation, the resulting state is:

$$\begin{aligned} &\rho = |\psi\rangle\langle\psi| \to \\ &M_i(\rho) = \rho_{\text{multi}}^{\text{mix}} = \left(|a_0|^2|A_0\rangle\langle A_0| + |a_1|^2|A_1\rangle\langle A_1|\right) \otimes \frac{1}{3}\left(|20\rangle\rangle\langle\langle 20| + |02\rangle\rangle\langle\langle 02| + |11\rangle\rangle\langle\langle 11|\right). \end{aligned} \tag{4.7}$$

This has the same form as the state obtained in the case of loss, in eq. 4.2, except that the faulty qubit now contains two photons in a random configuration. This can be viewed as the simplest description of multiphoton contamination, since there are no coherences left between the multiphoton qubit and the rest of the state, or between the two-qubit terms. The lack of these cross-terms greatly simplifies the analysis of the operation of fusion circuits in the presence of this error, as we will see in sec. 4.2 below. Moreover, the map in eq. 4.7 turns any single-qubit state into the maximally mixed two-photon state, which is invariant under any passive linear optical transformation, so it can be commuted through the worldline of the qubit and applied at the input of the fusion circuit, regardless of where the multiphoton contamination takes place. This property also means that the structure of the two-photon state cannot be manipulated by linear optical circuits, ensuring that no spurious effects are observed in the results of the analysis which are specific only to this model.

Photon addition model

A more physically-motivated approach to modelling multiphoton contamination is to treat it similarly to loss, by coupling the r rails of the input state to an ancilla state. In this case, the ancilla is in the maximally mixed state of one photon in r rails:

$$\rho_A = \frac{1}{r}(|10\ldots0\rangle\rangle\langle\langle 10\ldots0| + \ldots + |0\ldots01\rangle\rangle\langle\langle 0\ldots01|). \tag{4.8}$$

Since we only focus on first-order errors in our analysis, we assume that this photon is indistinguishable from other photons in the system. Consider the pure multi-qubit state $\rho = |\psi\rangle\langle\psi|$ previously defined. Coupling one of its qubits to ρ_A and postselecting on one photon being added to its rails

gives the following transformation:

$$\rho = |\psi\rangle\langle\psi| \rightarrow$$
$$\rho^{\text{add}}_{\text{multi}} = \left(\sqrt{\frac{2}{3}}a_0|A_0\rangle|20\rangle\rangle + \sqrt{\frac{1}{3}}a_1|A_1\rangle|11\rangle\rangle\right)\left(\sqrt{\frac{2}{3}}a_0\langle A_0|\langle\langle 20| + \sqrt{\frac{1}{3}}a_1\langle A_1|\langle\langle 11|\right) +$$
$$\left(\sqrt{\frac{2}{3}}a_1|A_1\rangle|02\rangle\rangle + \sqrt{\frac{1}{3}}a_0|A_0\rangle|11\rangle\rangle\right)\left(\sqrt{\frac{2}{3}}a_1\langle A_1|\langle\langle 02| + \sqrt{\frac{1}{3}}a_0\langle A_0|\langle\langle 11|\right). \tag{4.9}$$

This is a mixture of states corresponding to the addition of one photon in the first (top line) and second qubit rail (bottom line). When the two photons are bunched in the same rail the original state of the qubit before the contamination event can still be inferred. The state also retains coherence between the $|A_0\rangle$ and $|A_1\rangle$ terms, which plays an important role when this multiphoton state is subject to loss, as discussed below. Like loss, this channel can be shown to commute with unitary transformations which involve all qubit rails, so all multiphoton contamination along the path of a qubit can be applied at the input of the fusion it is measured in. Errors due to detector false clicks in SSG circuits or fusions cannot be described in this way, but their rate is assumed to be much lower than other sources and left out of this leading-order analysis [77].

4.1.4 Combination of multiphoton contamination and loss

Loss and multiphoton contamination change the number of photons entering a fully destructive n-way fusion, resulting in heralded errors when they occur on their own. On the other hand, when both these errors take place at the same time, the total number of photons in the circuit is conserved, so underalded errors can take place. Due to the rate of loss being much higher than other physical errors, this is considered to be one of the main mechanisms causing Pauli errors alongside distinguishability, despite it being a second-order effect. The two errors can affect different input qubits or the same one. The former case is modelled by applying the chosen multiphoton map to one qubit and the loss

map to the other before propagating the inputs through the circuit. The Pauli error rate in the latter scenario can be calculated analytically directly by applying the loss channel to the multiphoton qubit state.

When applied to the two-photon state obtained in the maximally-mixed model, loss transforms it into a maximally-mixed single-qubit state:

$$\rho^{\text{mix}}_{\text{multi}} \to \rho^{\text{mix}}_{\text{multi+loss}} = \left(|a_0|^2 |A_0\rangle\langle A_0| + |a_1|^2 |A_1\rangle\langle A_1|\right) \otimes \frac{1}{2}\left(|10\rangle\rangle\langle\langle 10| + |01\rangle\rangle\langle\langle 01|\right). \tag{4.10}$$

The overlap of $\rho^{\text{mix}}_{\text{multi+loss}}$ with the ideal state ρ is:

$$\langle\psi|\,\rho^{\text{mix}}_{\text{multi+loss}}\,|\psi\rangle = \frac{1}{2}\left(|a_0|^4 + |a_1|^4\right). \tag{4.11}$$

The probability of a Pauli error is $1 - \langle\psi|\,\rho^{\text{mix}}_{\text{multi+loss}}\,|\psi\rangle$. For example, if $a_0 = a_1 = \sqrt{1/2}$, the error probability is $3/4$, equally distributed across the three Paulis.

On the other hand, when the photon addition model is considered and the loss map is applied to the two-photon terms in eq. 4.9, the result is:

$$\begin{aligned}
&\rho^{\text{add}}_{\text{multi}} \to \\
&\rho^{\text{add}}_{\text{multi+loss}} = \quad \frac{5}{6}\left(\sqrt{\frac{4}{5}}a_0|A_0\rangle|10\rangle\rangle + \sqrt{\frac{1}{5}}a_1|A_1\rangle|01\rangle\rangle\right)\left(\sqrt{\frac{4}{5}}a_0\langle A_0|\langle\langle 10| + \sqrt{\frac{1}{5}}a_1\langle A_1|\langle\langle 01||\right) + \\
&\qquad\qquad \frac{5}{6}\left(\sqrt{\frac{4}{5}}a_1|A_1\rangle|01\rangle\rangle + \sqrt{\frac{1}{5}}a_0|A_0\rangle|10\rangle\rangle\right)\left(\sqrt{\frac{4}{5}}a_1\langle A_1|\langle\langle 01| + \sqrt{\frac{1}{5}}a_0\langle A_0|\langle\langle 10|\right) + \\
&\qquad\qquad \frac{1}{6}\left(|a_1|^2|A_1\rangle|10\rangle\rangle\langle A_1|\langle\langle 10| + |a_0|^2|A_0\rangle|01\rangle\rangle\langle A_0|\langle\langle 01|\right).
\end{aligned} \tag{4.12}$$

In this case the overlap with the ideal state ρ is

$$\langle\psi|\,\rho^{\text{add}}_{\text{multi+loss}}\,|\psi\rangle = \frac{5}{6}\left(\frac{4|a_0|^4}{5}+\frac{4|a_0|^2|a_1|^2}{5}+\frac{|a_1|^4}{5}\right)+\frac{5}{6}\left(\frac{|a_0|^4}{5}+\frac{4|a_0|^2|a_1|^2}{5}+\frac{4|a_1|^4}{5}\right), \quad (4.13)$$

which means that the infidelity for $a_0 = a_1 = 1/\sqrt{2}$ is $1/4$, which is three times lower compared to the error rate for the maximally-mixed model. In this case, the state in eq. 4.12 can be rewritten as:

$$\rho^{\text{add}}_{\text{multi+loss}} = \frac{3}{4}|\psi\rangle\langle\psi| + \frac{1}{12}|\psi^X\rangle\langle\psi^X| + \frac{1}{12}|\psi^Y\rangle\langle\psi^Y| + \frac{1}{12}|\psi^Z\rangle\langle\psi^Z|, \quad (4.14)$$

where $|\psi\rangle = \frac{1}{\sqrt{2}}(|A_0\rangle\rangle|10\rangle\rangle + |A_1\rangle\rangle|01\rangle\rangle)$ is the starting state as previously seen, and the other states are defined as $|\psi^P\rangle = (I\otimes P)\,|\psi\rangle$ where $P \in \{X,Y,Z\}$ is a Pauli rotation on the qubit whose state is explicitly written out. So we see that in this case the effect of multiphoton contamination followed by loss is the application of a Pauli rotation on the state with total probability $1/4$, uniformly distributed accross the three Pauli operators.

Eq. 4.12 describes the case in which loss occurs after multiphoton contamination. It is also possible for the order of these two error events to be reversed, for example if a photon is added within a multiplexing network to a qubit which has already suffered a loss. In that case the final state is the same as $\rho^{\text{mix}}_{\text{multi+loss}}$. Under the assumption that heralding at the source is the dominant contribution to the total multiphoton error rate, loss is more likely to be the second of the two errors to affect the qubit. This justifies mainly focusing on the case captured by eq. 4.12 in a leading-order analysis based on the photon addition model, with the understanding that a combination of $\rho^{\text{add}}_{\text{multi+loss}}$ and $\rho^{\text{mix}}_{\text{multi+loss}}$ will be present in a realistic architecture.

4.1.5 Qubit Pauli errors

Before reaching the input of a fusion, photons are interfered and propagated through linear optical networks, where imperfections can alter the qubit state. The three physical noise mechanisms described

so far can affect the operation of not only fusions, but also Seed State Generation (SSG) circuits, resulting in unheralded rotations on the output seed states [134, 162–165]. In addition, noise in the multiplexing or routing networks encountered by the qubits after the SSG stage can also introduce unwanted rotations of the state.

In general, physical noise can lead to coherent errors, which are known to be more difficult to analyse and simulate than probabilistic, stochastic Pauli errors [172]. It is therefore desirable to be able to describe the input qubits of a fusion with the latter type of error model. The technique of *Pauli twirling* allows to do this by applying random Pauli operations to each qubit, which transforms the errors in such a way that they can be described as stochastic Pauli operations [173]. In our analysis, it is assumed that this method is used such that it is sufficient to describe errors accumulated by the qubits before they reach the inputs of a fusion as a probability distribution of Pauli errors, i.e., with a Pauli error model. In general, such a model allows to describe each error as a Pauli dephasing channel on the input qubit:

$$\rho \to \Delta_P(\rho) = p_p \times \frac{\rho + P\rho P}{2}, \tag{4.15}$$

where P is any single- or multi-qubit Pauli operator and p_P is the probability of the error occurring. In our analysis we only consider single-qubit Pauli errors which may occur on each fusion input qubit, under the assumption that the qubits entering a fusion belong to independent states and have no correlated errors between them.

4.2 Pauli errors in linear optical fusions

The physical mechanisms described above can cause both heralded and unheralded errors at a fusion. In this and the following chapter we focus on understanding the behaviour of fusion circuits with respect to the latter, since we are mostly interested in their analysis in the context of RSG, where the former can be mitigated via multiplexing. Unheralded errors can lead to Pauli errors on the fusion

outcomes, and the rate at which they do so depends, in part, on the size and structure of the fusion circuit. To understand how different fusions respond to these types of errors, we aim to translate the fusion operation obtained when the physical error channels described above are applied to the inputs of an ideal fusion circuit into an equivalent description of the ideal fusion preceded by Pauli errors on the input qubits. This effective Pauli error model allows to compare different fusion circuits with respect to the same types of error metrics, which are also relevant to the assessment of their impact on the performance of LO-FBQC architectures, as described in detail in Chapter 6.

Among the types of errors listed in the previous section, distinguishability and combinations of multi-photon contamination and loss are the main sources of Pauli errors on the fusion outcomes, in addition to Pauli errors accumulated by the input qubits along their path to the fusion. The translation into a Pauli error model is only required for the former, since the latter are already expressed in this form. Here we introduce methods to perform this translation and to obtain a description of the effective operation performed by a fusion in the presence of noise as a probability distribution of Pauli dephasing channels acting on the input qubits. We start from the derivation of the Pauli error models for two-way T1 and T2 fusions, and then introduce methods to compose them into models for larger standard n-GHZ fusions.

4.2.1 Two-way fusions

T1 and T2 fusions are the building blocks of standard n-GHZ fusions, as shown in fig. 3.2. It is therefore useful to understand the effect of faulty inputs on their success measurement operators as a way to obtain insight into the behaviour of larger fusion circuits.

Distinguishability

The behaviour of linear optical two-way fusions in the presence of distinguishability has been thoroughly investigated [162–165, 174]. Here we consider the scenario in which one of the input qubits

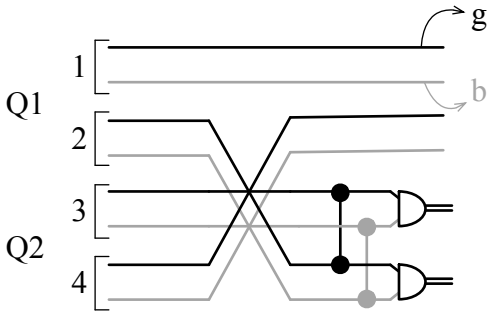

Figure 4.1: Circuit used to analyse the action of a T1 fusion circuit in the presence of distinguishability. The circuit is duplicated, with the grey copy acting on the fully distinguishable photon (b) in parallel with the black copy acting on the ideal photon (g). Linear optical operations are the same in both cases. The two copies of each detected rail are recombined at the detectors, which are assumed not to have access to the internal state of the photons.

of the fusion is ideal while the other one is fully distinguishable, i.e., the channel in eq. 4.5 is applied to it. The action of a fusion circuit on this input can be analysed by propagating the distinguishable photon through its own parallel copy of the circuit and recombining corresponding modes from the two qubits at the detectors, tracing out the internal state of the photons. This reflects the assumptions that the components of linear optical circuits have the same action on the logical part of the state regardless of its internal state, and that single photon detectors only count the number of photons in a rail without registering any information about their internal state.

Fig. 4.1 shows the diagram used to model a T1 fusion. The action of this circuit on a two-qubit basis where the second input is distinguishable is given by:

$$\begin{aligned}
a^\dagger_{1,g}a^\dagger_{3,b} &\to \frac{1}{\sqrt{2}}a^\dagger_{1,g}\left(a^\dagger_{3,b}+a^\dagger_{4,b}\right),\\
a^\dagger_{1,g}a^\dagger_{4,b} &\to a^\dagger_{1,g}a^\dagger_{2,b},\\
a^\dagger_{2,g}a^\dagger_{3,b} &\to \frac{1}{2}\left(a^\dagger_{3,g}-a^\dagger_{4g}\right)\left(a^\dagger_{3,b}+a^\dagger_{4,b}\right),\\
a^\dagger_{2,g}a^\dagger_{4,b} &\to \frac{1}{\sqrt{2}}a^\dagger_{2,b}\left(a^\dagger_{3,g}-a^\dagger_{4,g}\right),
\end{aligned} \tag{4.16}$$

where the subscripts g and b refer to the two copies of the rail acting separately on the two internal

states. The first and last line correspond to the success outcomes with one photon detected, giving the same success probability of $1/2$ for independent input qubits like in the ideal case. From these, depending on whether the ideal or the indistinguishable photon is detected, we obtain the operators:

$$F_{\mathrm{T1},1}^{\mathrm{dist}} = \{|0\rangle\langle 00| \otimes |g\rangle\langle gb|, |1\rangle\langle 11| \otimes |b\rangle\langle gb|\}. \tag{4.17}$$

The logical operators and the internal state of the output qubit are correlated by these outcomes. Applying single-qubit X rotations to the two inputs and the output of the T1 fusion results in a circuit with the same action on an ideal input in the case of success, but with opposite correlations with the internal state when one of the qubits is distinguishable:

$$F_{\mathrm{T1},2}^{\mathrm{dist}} = \{|0\rangle\langle 00| \otimes |b\rangle\langle gb|, |1\rangle\langle 11| \otimes |g\rangle\langle gb|\}. \tag{4.18}$$

In the rest of this analysis a T1 fusion is modelled as an equal mixture of these two versions. While this is a helpful step to simplify the modelling of larger fusions and obtain an understanding of their behaviour, in practice only one of the two circuits can be built, which results in correlations being present. The more general modelling framework described in the Chapter 5 allows to account for this and removes this simplification.

Averaging over the two sets of operators gives the channel:

$$\begin{aligned}
\rho_{\mathrm{in}} \rightarrow & \frac{1}{4}\sum_i F_i \rho_{\mathrm{in}} F_i^\dagger = \\
= & \frac{1}{4}|0\rangle\langle 00| \otimes (|g\rangle\langle gb|\,\rho_{\mathrm{in}}\,|gb\rangle\langle g| + |b\rangle\langle gb|\,\rho_{\mathrm{in}}\,|gb\rangle\langle b|) \otimes |00\rangle\langle 0| + \\
& \frac{1}{4}|1\rangle\langle 11| \otimes (|g\rangle\langle gb|\,\rho_{\mathrm{in}}\,|gb\rangle\langle g| + |b\rangle\langle gb|\,\rho_{\mathrm{in}}\,|gb\rangle\langle b|) \otimes |11\rangle\langle 1|,
\end{aligned} \tag{4.19}$$

where the $F_i \in F_{\mathrm{T1},1} \cup F_{\mathrm{T1},2}$.Consider the action of a dephasing Z channel on input qubit 1 of an ideal

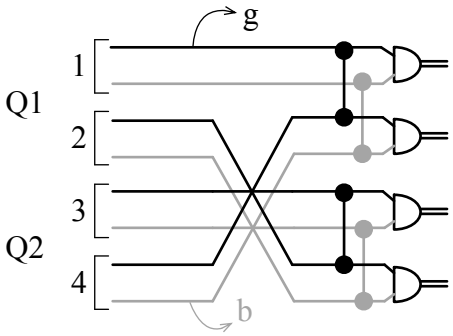

Figure 4.2: Circuit used to analyse the action of a T2 fusion circuit in the presence of distinguishability.

T1 fusion:

$$\begin{aligned}
&\Delta_{Z_1}\left(\frac{1}{2}\left(|0\rangle\langle 00|+|1\rangle\langle 11|\right)\rho_{\text{in}}\left(|00\rangle\langle 0|+|11\rangle\langle 1|\right)\right)= \\
&=\frac{1}{4}\left[\left(|0\rangle\langle 00|+|1\rangle\langle 11|\right)\rho_{\text{in}}\left(|00\rangle\langle 0|+|11\rangle\langle 1|\right)+\left(|0\rangle\langle 00|-|1\rangle\langle 11|\right)\rho_{\text{in}}\left(|00\rangle\langle 0|-|11\rangle\langle 1|\right)\right]= \\
&=\frac{1}{2}\left(|0\rangle\langle 00|\,\rho_{\text{in}}|00\rangle\langle 0|+|1\rangle\langle 11|\,\rho_{\text{in}}|11\rangle\langle 1|\right).
\end{aligned} \tag{4.20}$$

This can be combined with the distinguishability channel in eq. 4.5 to rewrite eq. 4.19 as a channel acting on ideal inputs:

$$\begin{aligned}
\rho_{\text{in}} \to &\left(\frac{1}{2}\left(I_o+D_o\right)\otimes\Delta_{Z_1}\otimes I_2\right)\left(\frac{1}{2}\left(|0\rangle\langle 00|+|1\rangle\langle 11|\right)\rho_{\text{in}}\left(|00\rangle\langle 0|+|11\rangle\langle 1|\right)\right)= \\
&=\left(\frac{1}{2}\left(I_o+D_o\right)\otimes I_1\otimes I_2\right)\left(\frac{1}{2}\left(|0\rangle\langle 00|\,\rho_{\text{in}}|00\rangle\langle 0|+|1\rangle\langle 11|\,\rho_{\text{in}}|11\rangle\langle 1|\right)\right)= \\
&=\frac{1}{4}|0\rangle\langle 00|\otimes\left(|g\rangle\langle gg|\,\rho_{\text{in}}|gg\rangle\langle g|+|b\rangle\langle gg|\,\rho_{\text{in}}|gg\rangle\langle b|\right)|00\rangle\langle 0|+ \\
&\frac{1}{4}|1\rangle\langle 11|\otimes\left(|g\rangle\langle gg|\,\rho_{\text{in}}|gg\rangle\langle g|+|b\rangle\langle gg|\,\rho_{\text{in}}|gg\rangle\langle b|\right)|11\rangle\langle 1|,
\end{aligned} \tag{4.21}$$

where the label o refers to the output qubit of the T1 fusion and $1,2$ are the two input qubits. From this we see that distinguishability on one of the input qubits can be mapped to an equal mixture of two channels: a Δ_Z on one input and the combination of an input Δ_Z with a distinguishability channel D on the output qubit, as seen in the first line above.

A similar analysis can be conducted for the T2 fusion using the circuit in fig. 4.2. In this case the input basis for the case that the second qubit is distinguishable is evolved as:

$$\begin{aligned}
a^\dagger_{1,g}a^\dagger_{3,b} &\to \frac{1}{2}\left(a^\dagger_{1,g}+a^\dagger_{2,g}\right)\left(a^\dagger_{3,b}+a^\dagger_{4,b}\right),\\
a^\dagger_{1,g}a^\dagger_{4,b} &\to \frac{1}{2}\left(a^\dagger_{1,g}+a^\dagger_{2,g}\right)\left(a^\dagger_{1,b}-a^\dagger_{2,b}\right),\\
a^\dagger_{2,g}a^\dagger_{3,b} &\to \frac{1}{2}\left(a^\dagger_{3,g}-a^\dagger_{4g}\right)\left(a^\dagger_{3,b}+a^\dagger_{4,b}\right),\\
a^\dagger_{2,g}a^\dagger_{4,b} &\to \frac{1}{2}\left(a^\dagger_{3,g}-a^\dagger_{4,g}\right)\left(a^\dagger_{1,b}-a^\dagger_{2,b}\right).
\end{aligned} \tag{4.22}$$

The top and bottom lines correspond to the success outcomes, where one photon is detected across detectors 1 and 2, and one across detectors 3 and 4. The success probability does not change compared to the ideal case. Depending on which of the detected photons is distinguishable, a projection is performed onto either $|00\rangle\langle 00|$ or $|11\rangle\langle 11|$ with equal probability, i.e., the operator which is applied is:

$$F^{\text{dist}}_{\text{T2}} = \frac{1}{2}\left(|00\rangle\langle 00| + |11\rangle\langle 11|\right). \tag{4.23}$$

Similarly to eq. 4.20, this can be seen to be equivalent to an ideal T2 fusion with a Z dephasing channel on one of its input qubits. This reproduces a known result from Refs. [162, 163].

Multiphoton contamination and loss

The effect of multiphoton contamination and loss affecting the same qubit was described in the previous section, where the state of the faulty qubit was found to be fully depolarized in eq. 4.10. This can be described as the application of two dephasing channels, Δ_Z and Δ_X, to the qubit before the fusion, since:

$$\Delta_Z\left(\Delta_X(\rho)\right) = \Delta_Z\left(\frac{\rho + X\rho X}{2}\right) = \frac{\rho + X\rho X + Y\rho Y + Z\rho Z}{4}. \tag{4.24}$$

This can be expressed more generally as a composition rule for dephasing channels:

$$\Delta_{P_1} \circ \Delta_{P_2} = \frac{1}{2}(\Delta_{P_1} + \Delta_{P_2} + \Delta_{P_1 P_2} - I), \tag{4.25}$$

from which we also see that $\Delta_P \circ \Delta_P = \Delta_P$.

The effect of multiphoton contamination and loss on two different qubits can be analysed in a similar way to the case of distinguishability, except that in this case it is not necessary to consider parallel copies of the fusion circuit. In a T1 fusion, loss and multiphoton contamination can separately lead to a success outcome, because the measurement is not fully destructive. Consider the action of the circuit when the top qubit is in a vacuum state and the bottom one is ideal. The fusion can only succeed if the photon is in the third rail, meaning that the operator $|00\rangle\rangle\langle\langle 00|\langle 0|$ is applied. This can be turned into an operator on a dual-rail qubit input subspace by realising that none of the stabilizer generators of the fusion measurement are preserved. This is equivalent to an effective model in which the loss channel is applied to the output qubit, which is in fact in a $|00\rangle\rangle$ state, removing any stabilizer operator with support on it, and an X dephasing channel acts on one of the two input qubits, randomizing the $I|ZZ$ generator.

When one of the inputs, e.g., the top qubit, is in the maximally mixed two-photon state as a result of multiphoton contamination, the term with two bunched photons in the detected mode cannot produce a success detection pattern, while the other two can combine with each of the basis states of the other ideal qubit to lead to success. This means that the success probability of the fusion is reduced to $2/3$ of the ideal one, since each ideal basis state has probability $1/2$ and each two-photon term $1/3$. The measurement operators applied by the T1 fusion in fig. 4.1 are $\{|20\rangle\rangle\langle\langle 20|\langle 0|, |11\rangle\rangle\langle\langle 11|\langle 1|\}$. As in the case of loss, this corresponds to an effective model on a two-qubit input space where all stabilizer generators of the fusion are lost, which can be described as the application of the multiphoton contamination channel on the output and a Δ_X channel on one of the inputs.

If one of the input qubits is lost and the other one suffers multiphoton contamination, only the $|11\rangle\rangle$ two-photon term can contribute to a one-photon success detection pattern. So, the success probability is reduced to $1/3$ in this case too. One photon exits the circuit in either rail depending on the version of the T1 fusion circuit considered. Once again, all stabilizer fusion outcomes are lost because the two input qubits are missing. Since the output rails are in the qubit subspace, with one photon between them, in this case the effective model which captures this is one in which Δ_X channels are applied to one of the inputs and the output qubit, which randomizes all ZZ stabilizer generators, and a Δ_Z channel acts on one of the qubits, randomizing the value of the XXX operator.

T2 fusions are fully destructive, so loss or multiphoton contamination alone cannot cause Pauli errors, because they alter the number of photons which reach the detectors, resulting in deterministic failure. Therefore, only the scenario in which the two errors occur on the two input qubits needs to be considered. The success probability of the fusion is reduced to $1/3$, because the $|11\rangle\rangle$ two-photon term is the only one which produces success detection patterns. Since both input qubits are leaked outside of the dual-rail single-photon subspace, the values of both the ZZ and XX stabilizer operators are lost, which corresponds to a model in which a fully depolarizing $\Delta_Z \circ \Delta_X$ channel is applied to one of the inputs.

We introduce a graphical representation of the mapping from physical to Pauli error channels on the fusion inputs based on the ZX-calculus language. Circles of different colours are added to the legs of the T1 and T2 spiders to represent the action of different error channels: grey for distinguishability, yellow for loss, purple for multiphoton contamination (under the maximally mixed model), red for Δ_Z, and blue for Δ_X. The translations of the different error combinations on the inputs into effective Pauli error models on the input qubits of both fusions are summarised graphically in fig. 4.3. Although a specific assignment of dephasing channels on the inputs and output is shown, it is possible to find equivalent choices using the stabilizer generators of the spiders. For example, consider the Z dephasing channel on the input of a T2 fusion in the top right diagram in fig. 4.3. Multiplying the

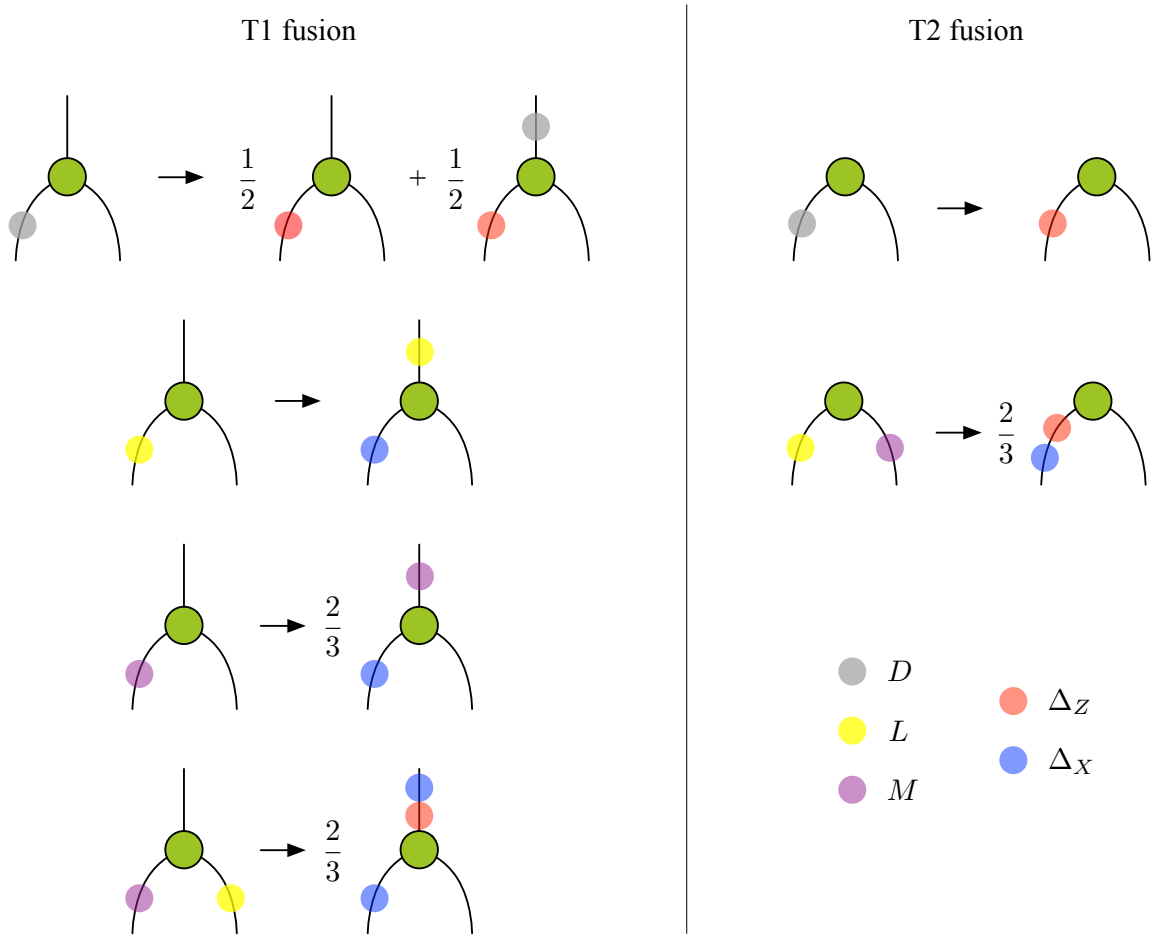

Figure 4.3: Mapping of physical error channels on the inputs of T1 and T2 fusions to effective Pauli error channels, plus physical errors on the output for the former. Arrows show the correspondence between physical error channels and Pauli error channels on input qubits. Circles of different colours are used to represent each error channel: grey for distinguishability, yellow for loss, purple for multi-photon contamination (maximally mixed model), red for Z dephasing, and blue for X dephasing. The $1/2$ factors in the top left diagram represent the mixture of two channels, and the $2/3$ factors in other diagrams represent a reduction in the fusion success probability.

input channel by the stabilizer generator ZZ gives:

$$(Z_1Z_2)\Delta_{Z_1}(\rho) = \frac{(Z_1Z_2)\rho(Z_1Z_2) + Z_2\rho Z_2}{2} = \frac{\rho + Z_2\rho Z_2}{2} = \Delta_{Z_2}(\rho), \tag{4.26}$$

where the subscripts 1 and 2 refer to the two input qubits and $(Z_1Z_2)\rho(Z_1Z_2) = \rho$ was used. Therefore, choosing the dephasing channel to act on the other input leg gives an equivalent model.

The mappings obtained for the T2 fusion can be easily translated into a Pauli error model for the smallest fully destructive linear optical fusion. The possible input distinguishability configurations are {(dist, ideal), (ideal, dist)}, where each tuple represents a set of inputs and the label “dist” is used

for a distinguishable qubit. From the top right mapping in fig. 4.3, these both correspond to dephasing Z channels on the faulty input qubit, so we have:

$$\rho_{\text{in}} \to p_d \left(\Delta_{Z_1} + \Delta_{Z_2}\right)(\rho_{\text{in}}) = p_d \left(\rho_{\text{in}} + \frac{1}{2} Z_1 \rho_{\text{in}} Z_1 + \frac{1}{2} Z_2 \rho_{\text{in}} Z_2\right), \tag{4.27}$$

where p_d is the input distinguishability rate. Similarly, there are $n(n-1) = 2$ configurations of multiphoton contamination and loss on the two input qubits, {(multi, loss), (loss, multi)}, both of which correspond to a $2/3$ reduction in success probability and a fully depolarizing input channel as shown in the bottom right diagram in fig. 4.3, resulting in the map:

$$\begin{aligned} \rho_{\text{in}} \to & \frac{2}{3} p_m p_l \left(\left(\Delta_{Z_1} \circ \Delta_{X_1}\right) + \left(\Delta_{Z_2} \circ \Delta_{X_2}\right)\right)(\rho_{\text{in}}) = \\ & = \frac{2}{3} p_m p_l \left[\rho_{\text{in}} + \frac{1}{4}\left(X_1 \rho_{\text{in}} X_1 + Y_1 \rho_{\text{in}} Y_1 + Z_1 \rho_{\text{in}} Z_1 + X_2 \rho_{\text{in}} X_2 + Y_2 \rho_{\text{in}} Y_2 + Z_2 \rho_{\text{in}} Z_2\right)\right], \end{aligned} \tag{4.28}$$

where p_m and p_l are the input multiphoton contamination and loss rates. The same map, but without the $2/3$ factor, is obtained when these two errors affect the same qubit, per eq. 4.24, and Pauli errors already present on the input correspond to dephasing channels applied with probability p_P, where $P \in \{X, Y, Z\}$, as discussed above. All these error channels can be combined and translated into a probability distribution of flips on the stabilizer generators of the fusion success operator:

$$\begin{aligned} \Pr(-XX, +ZZ|\text{success}) &= p_d + \frac{1}{3} p_m p_l + \frac{1}{2} p_m p_l + 2p_Z, \\ \Pr(+XX, -ZZ|\text{success}) &= \frac{1}{3} p_m p_l + \frac{1}{2} p_m p_l + 2p_X, \\ \Pr(-XX, -ZZ|\text{success}) &= \frac{1}{3} p_m p_l + \frac{1}{2} p_m p_l + 2p_Y, \end{aligned} \tag{4.29}$$

where a $-$ sign next to a stabilizer generator is used to represent a flip with respect to its ideal value. From this we see that while combinations of multiphoton contamination and loss errors cause every possible flip with the same probability, distinguishability only flips the XX operator. Up to the values of the input Pauli rates, this means that the distribution of Pauli errors is biased towards this operator.

Eq. 4.29 is a Pauli error model for the smallest fully destructive fusion. Similar models can be obtained for larger n-GHZ fusions by composing the error channels derived above for T1 and T2 fusions according to the structure of the circuit, as described next.

4.2.2 Standard n-GHZ fusions

In Chapter 3, the view of standard n-GHZ fusions as concatenations of $n-1$ T1 fusions and one T2 fusion was introduced. In this picture, the larger fusion succeeds when a success outcome is obtained at all constituent two-way fusions. This structure is particularly helpful for error modelling, since it allows to derive the response of larger fusions to physical input errors through the combination of the two-way fusions models. The graphical notation in fig. 4.3 can be used to propagate physical errors through the ZX-diagrams of n-GHZ fusions and translate them into effective Pauli error channels on their input qubits. Using this framework, we derive leading-order Pauli error models for n-GHZ fusions with $3 \leq n \leq 5$ and use it to obtain an insight into the scaling of the impact of errors from different sources with the size of the fusion.

There are multiple possible decompositions of a standard n-GHZ fusion circuit into constituent two-way fusions, each corresponding to a different ZX-diagram. Diagrams can have inequivalent topologies, like in the case of the two $n = 4$ examples in fig. 3.2, or be identical up to relabelling of input qubits. In both cases, the Pauli error channels obtained when the two-way fusion models are combined can be different. Fig. 3.2 shows that the linear optical circuits corresponding to the two topologies differ in terms of the mode crossings they require. When the fusion is fabricated, one of the variants needs to be chosen, and imperfections in the crossing networks results in different error models for the two configurations. While crossing errors are not included in our analysis, we base our modelling on the version of the circuit with the most uniform and lowest average number of crossings per mode, since this balances and reduces errors due to these components. The corresponding ZX-diagram is the most balanced, "tree-like" version for any n. For this topology, n equivalent diagrams exist which

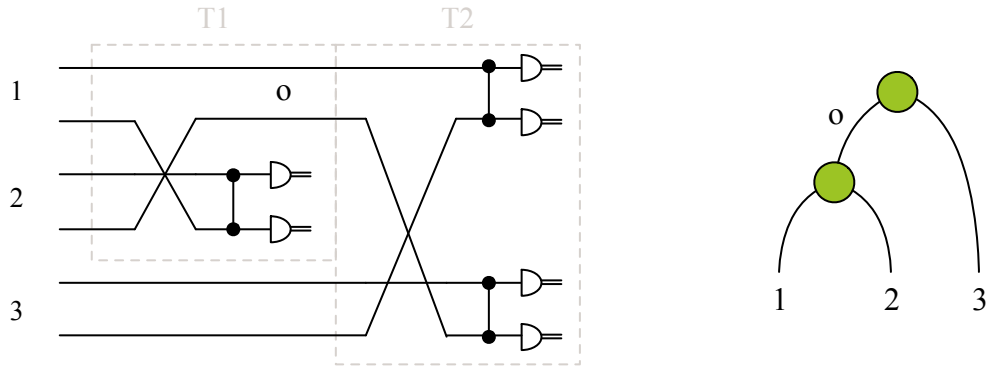

Figure 4.4: Linear optical circuit and ZX-diagram representation of a possible decomposition of the standard 3-GHZ fusion circuit into T1 and T2 fusion. The ZX-diagram corresponds to the success outcomes only.

can be obtained by applying cyclic permutations to the input, reflecting the symmetry of the n-GHZ analyzer circuit. We average over these by propagating errors through one representative and then symmetrising the resulting Pauli error channel across all equivalent combinations of input qubits.

3-GHZ fusion

A standard 3-GHZ fusion can only be decomposed as a T1 fusion feeding into one input of a T2 fusion, with one representative shown in circuit and ZX-diagram form in fig. 4.4, where the latter represents the success outcome only. Since this is a fully destructive fusion, to leading order, unheralded errors can only occur as a result of one input being distinguishable or of a combination of multiphoton contamination and loss on different qubits. Pauli errors already present on the input qubits also contribute to the effective Pauli error channel for the fusion, but they are independent of the structure of the circuit and can be added to it directly, as in the case of the T2 fusion above.

Consider the case of one of the input qubits being distinguishable. There are $n = 3$ configurations of this kind, which we denote as {(dist, ideal, ideal), (ideal, dist, ideal), (ideal, ideal, dist)}, where the three entries in a tuple represent the three qubits from left to right in the ZX-diagram and the faulty qubit is labelled as “dist”. In the case of the first configuration, the T1 fusion map (top left in fig. 4.3) can be used to push the distinguishability channel through the T1 spider, resulting in a Δ_{Z_1} channel on the first input qubit and a distinguishable output half of the time. Then, applying the

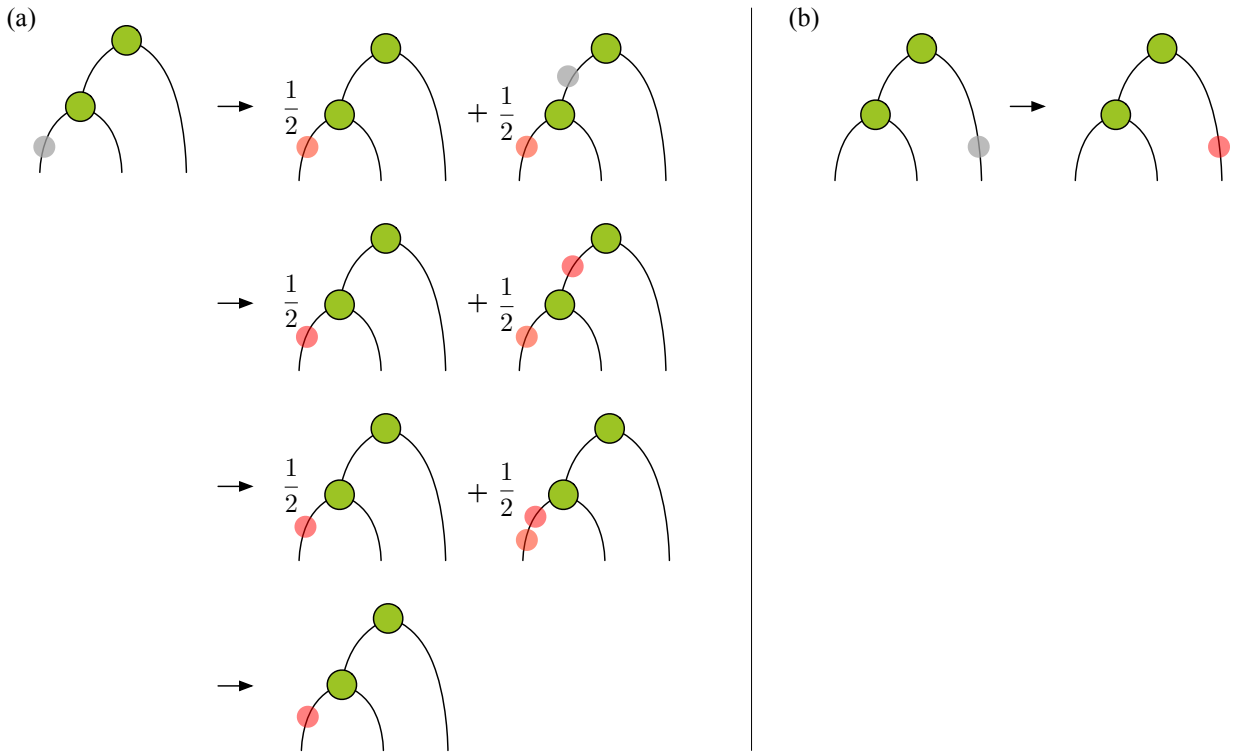

Figure 4.5: Graphical derivation of the Pauli error channels due to input distinguishability in the 3-GHZ fusion via composition of the T1 and T2 fusion maps in fig. 4.3. The (ideal, dist, ideal) input configuration is not shown because the derivation is identical to (a), except that the Δ_Z channel can be chosen to apply to qubit 2.

map for distinguishability at a T2 fusion (top right in fig. 4.3) transforms the $D_o(\rho)$ channel into a Δ_{Z_o} channel. Finally, conjugating this dephasing channel by the $Z_o|Z_1I$ stabilizer of the T1 spider results in the channel $\Delta_{Z_1} \circ \Delta_{Z_1} = \Delta_{Z_1}$, so that adding up the two halves of the calculation gives Δ_{Z_1} overall. The steps of the derivation are shown graphically in fig. 4.5(a). Since pairwise Z operators as stabilizer generators of an n-GHZ fusion, e.g., ZZI, ZIZ for this 3-GHZ fusion, all Δ_Z channels are equivalent. Therefore, any individual Δ_{Z_i} can be replaced by a symmetric channel:

$$\Delta_Z = \frac{1}{n} \sum_{i=1}^{n} \Delta_{Z_i}. \tag{4.30}$$

The (ideal, dist, ideal) configuration follows the same logic as (dist, ideal, ideal), also resulting in a Δ_Z channel. The third configuration can be modelled simply by applying the T2 fusion distinguishability map, as shown in fig. 4.5(b), resulting again in a Δ_Z. Combining these all three input configurations

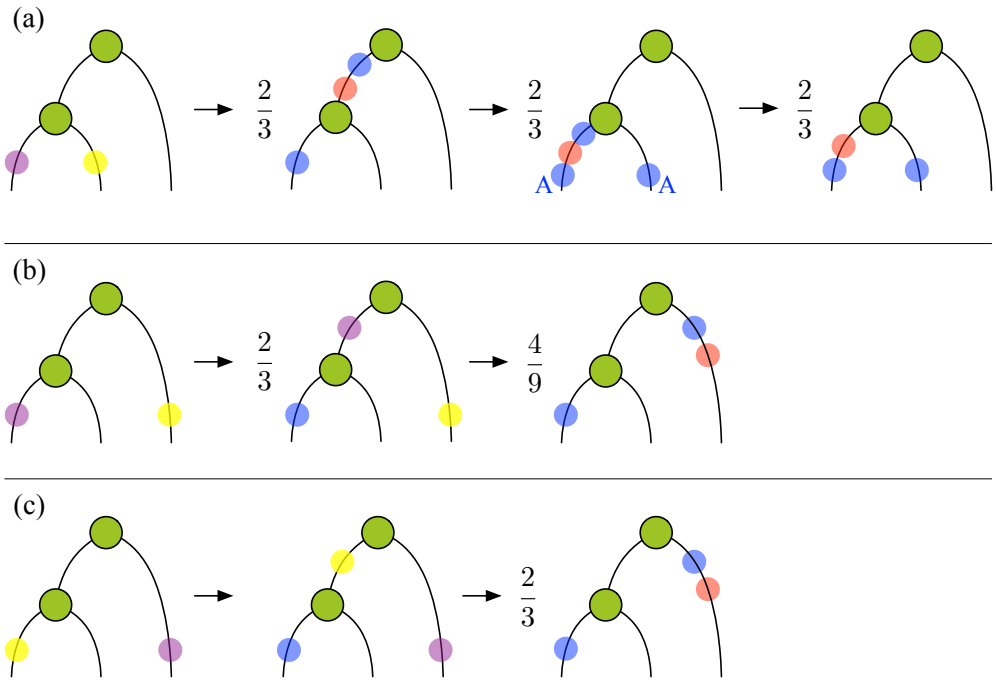

Figure 4.6: Graphical derivation of the Pauli error channels due to multiphoton contamination and loss on different inputs of the 3-GHZ fusion. Only three of the six possible input error configurations are shown which correspond to the composition of different two-way fusion maps. Each of the other three follows a similar derivation to one of these, up to the single Δ_X channel being applied to different qubits. In (a), two blue circles labelled with the same letter "A" indicate the $\Delta_{X_1X_2} \neq \Delta_{X_1} \circ \Delta_{X_2}$ channel.

then gives the channel:

$$\rho_{\text{in}} \to p_d \times 3 \times \Delta_Z(\rho_{\text{in}}) = \sum_{i=1}^{3} \Delta_{Z_i}(\rho_{\text{in}}) = p_d \left(\frac{3}{2}\rho_{\text{in}} + \frac{1}{2}Z_1\rho_{\text{in}}Z_1 + \frac{1}{2}Z_2\rho_{\text{in}}Z_2 + \frac{1}{2}Z_3\rho_{\text{in}}Z_3 \right). \quad (4.31)$$

Given three inputs, there are $n(n-1) = 6$ configurations of multiphoton contamination and loss on different qubits, for which we use labels "multi" and "loss" respectively. The propagation of errors in the case that the two inputs of the T1 fusion are faulty is shown in fig. 4.6(a). Conjugating the Δ_{Z_o} channel by the $Z_o|Z_1I$ stabilizer turns it into Δ_{Z_1}, and doing the same with the Δ_{X_o} channel and the $X_o|X_1X_2$ stabilizer operator gives $\Delta_{X_1X_2}$. Using eq. 4.25, this can be rewritten as:

$$\Delta_{X_1X_2} = \frac{2}{3}(2\Delta_{X_1} \circ \Delta_{X_2} - \Delta_{X_1} - \Delta_{X_2} + I), \quad (4.32)$$

so when this is combined with the Δ_{X_1} already present on the input, we get:

$$\Delta_{X_1} \circ \Delta_{X_1 X_2} = (2\Delta_{X_1} \circ \Delta_{X_1} \circ \Delta_{X_2} - \Delta_{X_1} \circ \Delta_{X_1} - \Delta_{X_1} \circ \Delta_{X_2} + \Delta_{X_1}) = \Delta_{X_1} \circ \Delta_{X_2}, \tag{4.33}$$

where we have used the identity $\Delta_P \circ \Delta_P = \Delta_P$. So, the mapping for these two configurations is:

$$\text{(multi, loss, ideal)} = \text{(loss, multi, ideal)} \to \frac{2}{3}\Delta_{X_1} \circ \Delta_{X_2} \circ \Delta_Z, \tag{4.34}$$

where the $2/3$ factor comes from the reduction in the success probability of the T1 fusion.

When one of the T1 fusion inputs suffers multiphoton contamination and qubit 3 is lost, applying the multiphoton T1 fusion map followed by the T2 fusion map results in the following mapping:

$$\begin{aligned} \text{(multi, ideal, loss)} &\to \frac{4}{9}\Delta_{X_1} \circ \Delta_{X_3} \circ \Delta_Z, \\ \text{(ideal, multi, loss)} &\to \frac{4}{9}\Delta_{X_2} \circ \Delta_{X_3} \circ \Delta_Z, \end{aligned} \tag{4.35}$$

where the $4/9$ factor here reflects the reduction in the success probability of both two-way fusions. This is shown for the first configuration in fig. 4.6(b). The same applies to the (loss, ideal, multi) and (ideal, loss, multi) configurations, except that in this case only the T2 fusion succeeds less often, so the factor is $2/3$, as seen in fig. 4.6(c).

The 3-GHZ ZX-diagram considered so far is a single representative of three topologically equivalent diagrams, with the other two pairing input qubits (2, 3) and (3, 1). This is due to the cyclic symmetry of the n-GHZ analyzer circuit. In order to obtain a model which averages across all three possible decompositions, the channels in eqs. 4.34 and 4.35 need to be symmetrized. This can be done by noticing that all of them have the form $\Delta_{X_i} \circ \Delta_{X_{(i+1) \pmod 3}} \circ \Delta_Z$, so that each operator can be rewritten

as an average of all three $i \in \{1,2,3\}$:

$$\Delta_{X_i} \circ \Delta_{X_{(i+1) \bmod 3}} \circ \Delta_Z \to \frac{1}{3} \sum_{j=1}^{3} \left(\Delta_{X_j} \circ \Delta_{X_{(j+1) \bmod 3}} \right) \circ \Delta_Z \tag{4.36}$$

Performing this replacement and combining the channels obtained for all six configurations results in the following total Pauli channel:

$$\begin{aligned}
\rho_{\text{in}} \to & p_m p_l \left(4 \times \frac{2}{3} + 2 \times \frac{4}{9} \right) \times \frac{1}{3} \sum_{i=1}^{3} \left(\Delta_{X_i} \circ \Delta_{X_{(i+1) \bmod 3}} \right) \circ \Delta_Z(\rho_{\text{in}}) = \\
= & \frac{16}{27} p_m p_l \sum_{i=1}^{3} \left(\Delta_{X_i} + \Delta_{X_{(i+1) \bmod 3}} + \Delta_{X_{(i+2) \bmod 3}} - I \right) \circ \Delta_Z(\rho_{\text{in}}) = \\
= & \frac{8}{9} p_m p_l \left(\Delta_{X_1} + \Delta_{Y_1} + \frac{1}{3}\Delta_{Z_1} + \Delta_{X_2} + \Delta_{Y_2} + \frac{1}{3}\Delta_{Z_2} + \Delta_{X_3} + \Delta_{Y_3} + \frac{1}{3}\Delta_{Z_3} - 3I \right) (\rho_{\text{in}}) = \\
= & \frac{4}{9} p_m p_l \left(\rho_{\text{in}} + X_1 \rho_{\text{in}} X_1 + Y_1 \rho_{\text{in}} Y_1 + X_2 \rho_{\text{in}} X_2 + Y_2 \rho_{\text{in}} Y_2 + X_3 \rho_{\text{in}} X_3 + Y_3 \rho_{\text{in}} Y_3 \right) + \\
& \frac{4}{27} p_m p_l \left(Z_1 \rho_{\text{in}} Z_1 + Z_2 \rho_{\text{in}} Z_2 + Z_3 \rho_{\text{in}} Z_3 \right),
\end{aligned} \tag{4.37}$$

where we have used the fact that for a 3-GHZ fusion $X_i X_{i+1 \pmod 3} = X_i X_{i+2 \pmod 3}$ since XXX is a stabilizer generator. Combining the Pauli error channels in eqs. 4.31 and 4.37 with depolarizing errors from multiphoton contamination and loss on the same input qubit and other input Pauli errors, as was done for the T2 fusions above, gives the following probability distribution of flips on the stabilizer

generators of the 3-GHZ projection:

$$\begin{aligned}
\Pr(-XXX,+ZZI,+ZIZ|\text{success}) &= \frac{3}{2}p_d + \frac{4}{9}p_m p_l + \frac{3}{4}p_m p_l + 3p_Z,\\
\Pr(+XXX,-ZZI,+ZIZ|\text{success}) &= \frac{4}{9}p_m p_l + \frac{1}{4}p_m p_l + p_X,\\
\Pr(+XXX,+ZZI,-ZIZ|\text{success}) &= \frac{4}{9}p_m p_l + \frac{1}{4}p_m p_l + p_X,\\
\Pr(+XXX,-ZZI,-ZIZ|\text{success}) &= \frac{4}{9}p_m p_l + \frac{1}{4}p_m p_l + p_X,\\
\Pr(-XXX,-ZZI,+ZIZ|\text{success}) &= \frac{4}{9}p_m p_l + \frac{1}{4}p_m p_l + p_Y,\\
\Pr(-XXX,+ZZI,-ZIZ|\text{success}) &= \frac{4}{9}p_m p_l + \frac{1}{4}p_m p_l + p_Y,\\
\Pr(-XXX,-ZZI,-ZIZ|\text{success}) &= \frac{4}{9}p_m p_l + \frac{1}{4}p_m p_l + p_Y.
\end{aligned} \tag{4.38}$$

As for the T2 fusion, a single flip of the XXX stabilizer is the most likely type of error, up to the Pauli rates on the input qubits, due to the contribution from distinguishability and the $3\times$ factor from the depolarizing noise on each input qubit. The two factors proportional to $p_m p_l$ capture the error rate due to combinations of multiphoton contamination and loss on different input qubits (first) and on the same input qubit (second).

Larger fusions

The same type of analysis can be repeated for larger n-GHZ fusions. As explained above, we choose to propagate errors through the most balanced versions of their ZX-diagram symmetrizing across combinations of inputs as in the case of the 3-GHZ fusion above.

Two types of input error configuration encountered for $n = 3$ can be directly generalized to larger standard GHZ fusions without the need to perform a full derivation. The derivation of the error channel associated with input distinguishability applies to any concatenation of green fusion spiders,

regardless of the number and exact configuration. Therefore eq. 4.31 can be extended to any n:

$$\rho_{\text{in}} \to p_d \times n \times \Delta_Z(\rho_{\text{in}}) \to p_d \sum_{i=1}^{n} \Delta_{Z_i}(\rho_{\text{in}}) = p_d \left(\frac{n}{2}\rho_{\text{in}} + \frac{1}{2}\sum_{i=1}^{n} Z_i \rho_{\text{in}} Z_i \right), \tag{4.39}$$

meaning that in all these circuits input distinguishability can only lead to Z errors, flipping the $X^{\otimes n}$ stabilizer generator with a total probability of $\frac{n}{2}p_d$. This agrees with the result obtained in the analysis of distinguishability in n-GHZ analyzers in Ref. [165]. In addition, the effect of configurations with multiphoton contamination and loss on the two inputs of the same T1 fusion is also independent of the fusion size and structure, since only the first two-way fusion spider is affected. The corresponding Pauli error model always has the form of eq. 4.36 and a success probability reduction factor of $\frac{2}{3}$:

$$\begin{aligned} \frac{2}{3}p_m p_l \left(\Delta_{X_i} \circ \Delta_{X_{i+1}} \circ \Delta_Z\right) =& \frac{1}{3}p_m p_l \left(\Delta_{X_i} + \Delta_{X_{i+1}} + \Delta_{X_i X_{i+1}} - I\right) \circ \Delta_Z = \\ =& \frac{1}{6}p_m p_l \left(\Delta_{X_i} + \Delta_{Y_i} + \Delta_{X_{i+1}} + \Delta_{Y_{i+1}} + \Delta_{X_i X_{i+1}} + \frac{1}{2}\left(\Delta_{Y_i X_{i+1}} + \Delta_{X_i Y_{i+1}}\right) + \Delta_Z - 5I \right), \end{aligned} \tag{4.40}$$

where the index $i+1$ is used to represent $(i+1)$ mod n for simplicity, and $\frac{1}{2}\left(\Delta_{Y_i X_{i+1}} + \Delta_{X_i Y_{i+1}}\right)$ is an average of the Z operator being composed with X_i and X_{i+1}. In the case of $n > 3$, $X_i X_{i+1}$ cannot be converted into a single-qubit X operator, so correlated error channels appear. This is an important property of larger n-way fusions, whose Pauli error channels can contain correlated errors of weight up to $\lfloor n/2 \rfloor$.

Other configurations of multiphoton contamination and loss at the input lead to similar Pauli channels consisting of combinations of Δ_X and one Δ_Z on adjacent qubits, which require propagating errors through the diagrams to be extracted. Once all channels are obtained for a given representative, they are symmetrized by averaging over all equivalent sets of adjacent qubits. Then, combining these averaged channels with the Z dephasing due to distinguishability (eq. 4.39), the depolarizing noise from multiphoton contamination and loss on the same qubit, and the dephasing Pauli errors present

on the input qubits, gives a total Pauli error channel which can be turned into a probability distribution of flips on the stabilizer generators of the n-GHZ fusion. The steps in the derivation of the Pauli error models for 4- and 5-GHZ fusions are described in Appendix B.

These models allow to calculate metrics to quantify and compare the performance of the different circuits in terms of Pauli errors. The *total Pauli error rate* is the total probability that any of the stabilizer generators of the fusion success outcome are flipped, i.e., the sum of all the flip probabilities in the distribution. This is the infidelity of the operation performed to the ideal fusion success outcome, which quantifies the sensitivity of a fusion to input physical errors. The *per-qubit marginal Pauli error rate* is the probability that each qubit suffers a Pauli X, Y or Z error, obtained by marginalizing the total Pauli error channel. This can be useful to understand the effect of a fusion on a specific input without the complexity introduced by correlated errors. In addition, it can provide insight into the degree of bias in the Pauli error distribution, which can in principle be exploited in the design of the architecture. Finally, marginal probabilities can also be more representative of the error rate seen by a regular decoder which does not have the capability to take correlations into account. The comparison of these metrics for standard n-GHZ fusions with $2 \leq n \leq 5$ is shown in figs. 4.7 and 4.8 respectively, where the values used for the physical error parameters are:

$$\begin{aligned} p_d &= 0.1\%, \\ p_m &= 0.3\%, \\ p_l &= 15\%, \\ p_X &= 0.01\% \\ p_Y &= 0.01\% \\ p_Z &= 0.03\%, \end{aligned} \tag{4.41}$$

chosen to be representative of projected performance of future photonic hardware. Cascaded res-

onator HSPS designs have been proposed which enable low spectral variability between different sources and high photon purity, resulting in very low photon distinguishability rates, of the order of p_d above [77]. To first order, the rate of multiphoton contamination is given by the probability that successful heralding at the source is due to a two-photon emission at the source followed by a loss at the herald detector. Major advances have been made with respect to the efficiency and number-resolution capabilities of SNSPDs [77, 175], with median on-chip detection efficiencies of $\sim 97\%$ accompanied by resolution of up to four photons. At a single-photon emission probability of 5%, this corresponds roughly to the multiphoton contamination rate $p_m = 0.3\%$ given above. p_l is the total loss rate on the fused qubits, including contributions from all elements along the photon worldline. In practice this will vary depending on the location of the fusion in the architecture. The $p_l = 15\%$ used here is a conservative estimate of the loss experienced by a photon which goes through two fibre delays, two electro-optic fast phase shifters, and eventually reaches a detector within a fusion, based on the projected component performance values in Ref. [77]. Finally, the rates of Pauli errors on input qubits are chosen to capture the fact that the effect of photon distinguishability at the input of an SSG circuit such as the 4-photon, 8-mode Bell state generator can be represented by a Pauli channel applied to one of the qubits, with the probability of a Z flip being higher than X or Y [162, 163]. This is not the only contribution to Pauli errors on the fusion input qubits, but assuming that other mechanisms result in unbiased errors, we expect p_Z to be higher than p_X and p_Y. A specific set of physical error rates was chosen for this analysis, and an effort was made to consider values currently expected to be representative of a realistic operational regime, such that informative conclusions applicable to the design of a real-world machine can be drawn. However, the framework presented here can be applied to any set of parameters, as long as the leading-order approximations are valid, and in practice it is necessary to update these choices based on the specifics of the system under consideration to obtain useful insights. One example of a different set of parameters is considered in Appendix C, where the impact of variations in the physical error rates on the results in this and the following chapter are

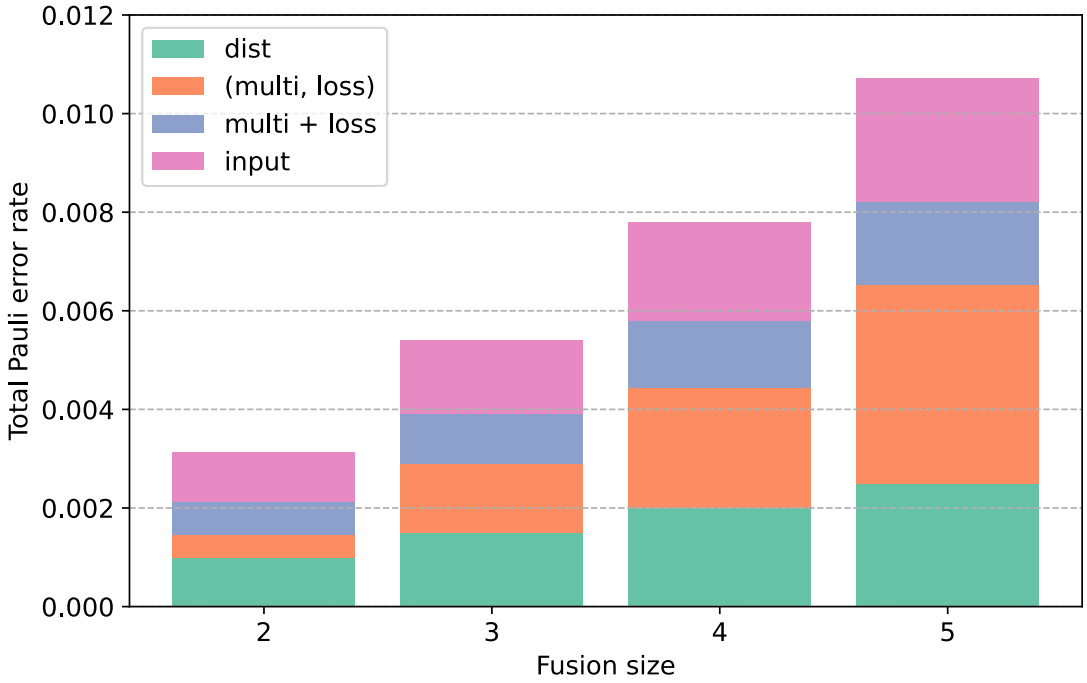

Figure 4.7: Total Pauli error rate for standard n-GHZ fusions with $2 \leq n \leq 5$. The physical error parameters are given in eq. 4.41. This is the probability that the success outcome suffers any Pauli flip. Colours are used to represent contributions from different types of input error configurations.

discussed.

The breakdown of different contributions to the total probabilities is shown in the plots, with the following labels: "dist" for distinguishability, "(multi, loss)" and "multi + loss" for combinations of multiphoton contamination and loss on different input qubits or the same one respectively, and "input" for the sum of all Pauli errors already present on the qubits when they reach the fusion. The total amount of Pauli error is seen to increase with fusion size in both plots, as expected due to the number of possible input error configurations growing with n. Distinguishability and input Pauli errors, including those due to multiphoton contamination and loss on the same qubit, can all occur in n different ways and result in single-qubit error channels. Therefore, their associated Pauli error probability also scales linearly with n, as reflected in the linear increase in the height of the green, blue and pink bars in the plot in fig. 4.7, and in their height being constant in fig. 4.8 which shows a per-qubit error probability. The number of "(multi, loss)" configurations, on the other hand, is $n(n-1)$, resulting in an almost quadratic scaling. The success probability of the fusion is reduced by different

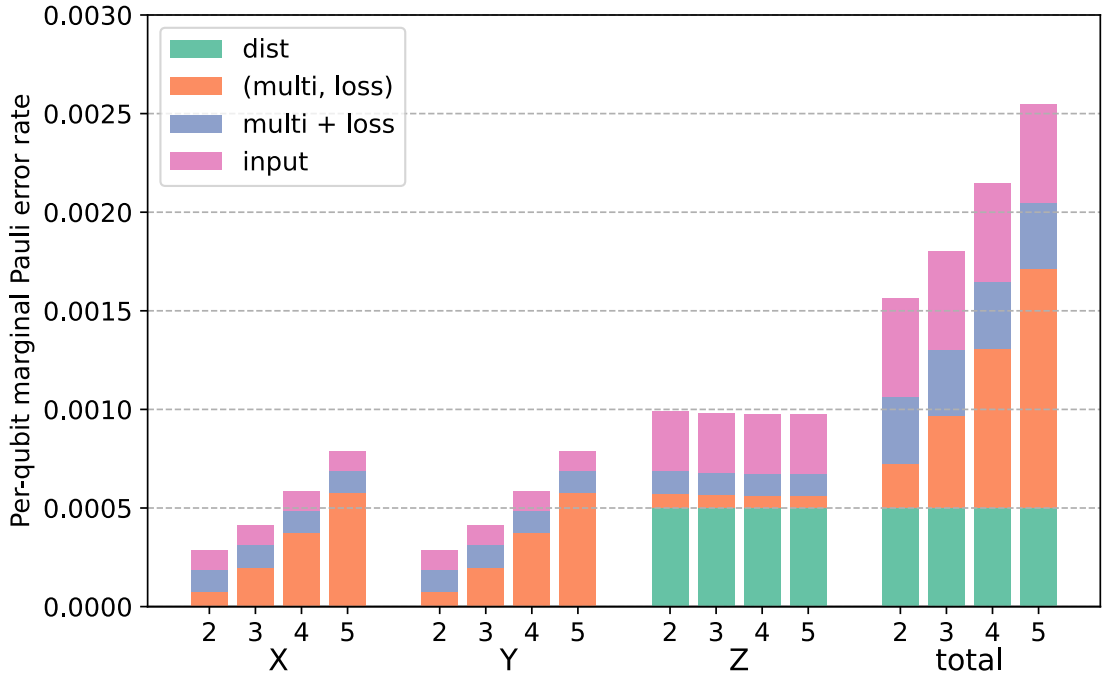

Figure 4.8: Per-qubit marginal Pauli error rate for standard n-GHZ fusions with $2 \leq n \leq 5$. The physical error parameters are given in eq. 4.41. This is obtained by marginalizing the Pauli error model and averaging over all input qubits. Colours are used to represent contributions from different types of error configurations.

amounts depending on the specific input error configuration, and on average the reduction is higher for larger fusions, which can be seen in the slight decrease in the orange Z bars. When looking at the total error rates, however, it is clear that this contribution is the fastest growing with n in both plots. The difference in the height of the "total" orange bars appears to grow significantly faster for the per-qubit marginal Pauli error rate than for the total Pauli error rate. This is due to correlated errors, which occur for $n > 3$ and are double-counted in the marginalization process, contributing to a larger increase in the overall probability.

4.2.3 Non-i.i.d. fusion Pauli errors

The performance of fault-tolerant schemes is typically evaluated under a phenomenological model of *independent and identically distributed (i.i.d.)* Pauli errors, the simplest and least biased model in the absence of additional information about the expected distribution of errors. In the context of FBQC architectures [43, 112, 113, 115, 117], this assumes that each of the stabilizer operators associated

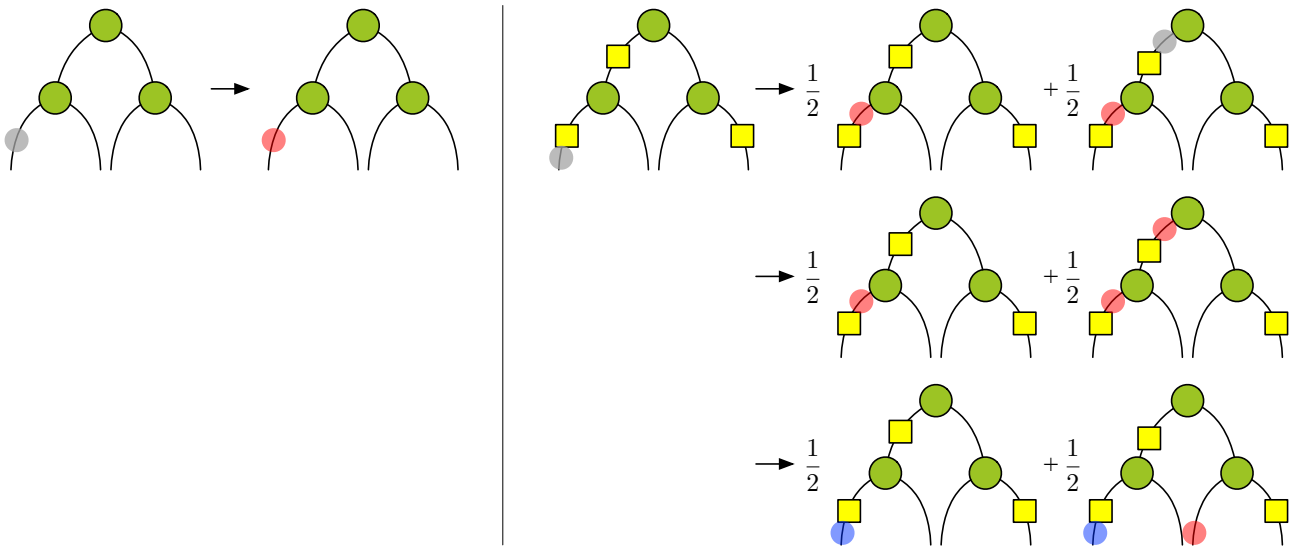

Figure 4.9: Pauli error channels associated with input distinguishability in the standard 4-GHZ fusion (left) and the 4-line fusion (right). From the previous sections it is known that a single-qubit Δ_Z channel is obtained for the former, whereas a Δ_X and a $\Delta_Z \circ \Delta_X$ channel arise for the latter. The distinguishability channel (grey) is assumed to commute through Hadamard nodes, and X and Z dephasing channels are exchanged across them.

with the success outcome has the same probability of suffering a flip. The framework described above takes a step towards a more realistic model of the errors caused by physical mechanisms in a real-world LO-FBQC architecture, and the error channels and metrics calculated point towards this differing significantly from an i.i.d. Pauli error model. More detailed error models are ultimately required to accurately characterize the performance of any fusion-based scheme, but these can be substantially more complex and computationally intensive to extract. A further step in this direction is described in the next chapter, where a more general, simulation-based modelling framework is introduced. On the other hand, while the approach described in this chapter has a few important limitations, it is at the same time detailed enough that the response of fusions to different sources of physical noise can be separated out, and it is simple enough that calculations can be performed analytically. The performance of different fusion circuits can thus be rapidly assessed and an intuition can be built for the impact of their structure on the resulting error model, providing valuable guidance in the design process.

As an example of the kind of insight that can be obtained using the methods described above, consider

the extraction of a Pauli error model for a 4-line fusion, which can be decomposed into T1 and T2 fusions in the same way as a 4-GHZ fusion, with the addition of a few Hadamard rotations. By propagating error channels through the corresponding ZX diagram, it is straightforward to identify differences between the Pauli error models for the two circuits. For instance, fig. 4.9 shows the derivation of the Pauli error channels corresponding to a single distinguishable input qubit in the standard 4-GHZ fusion (left) and in the 4-line fusion (right). In this calculation the distinguishability channel commutes through a Hadamard node, which is assumed not to alter the internal state of the qubit, and Δ_Z and Δ_X channels are exchanged through it. As discussed in the previous section, in the 4-GHZ fusion distinguishability leads to a single-qubit Z dephasing channel. On the other hand, the effect on the 4-line fusion can be either a single-qubit Δ_X channel on the first qubit or a two-qubit $\Delta_{X_1} \circ \Delta_{Z_3}$ channel, meaning that in this case not only single-qubit flips can occur, but also correlated two-qubit errors. Thus, we see that this method provides an insight into important differences between fusion circuits through a simple graphical calculation.

4.3 Conclusion

In order to understand whether multi-way fusions can be useful for the design of LO-FBQC schemes with improved performance, it is essential to understand their behaviour in the presence of physical noise. In this chapter, we introduced a framework to model the impact of leading-order sources of error in the linear optical platform on n-GHZ fusions constructed as concatenations of T1 and T2 fusions. In particular, we focus on the description of unheralded errors as Pauli error channels on the input qubits, since, unlike heralded errors, these cannot be filtered out through multiplexing in the architecture.

The results of the modelling presented here confirm our initial expectations described in Chapter 2: Pauli errors due to distinguishability are found to increase linearly with n, while combinations of

multiphoton contamination and loss scale more rapidly, roughly as n^2. The question of whether this scaling makes larger fusions beneficial for RSG is investigated in Chapter 6.

The modelling procedure described here is a powerful tool to obtain Pauli error models for fusions through simple graphical and analytic calculations, but it has a couple of important limitations: it is restricted to circuits which can be decomposed into two-way fusions only, it is based on an abstract T1 fusion model which does not capture the specific structure of a realistic linear optical multi-way fusion circuit, and it uses a simple multiphoton contamination error model. Therefore, before turning our focus to the use of fusions in the context of linear optical RSG, in the next chapter we introduce a new error modelling framework which aims to address these limitations.

Chapter 5

Pauli Errors in General Linear Optical Fusions

This chapter presents a general framework to obtain Pauli error model descriptions of the impact of physical noise on linear optical fusion circuits, overcoming the main limitations of the approach used in Chapter 4. This allows us to extend the error analysis to the new fusion circuits introduced in Chapter 3, not only those which admit a decomposition into two-way fusions, and to include more realistic physical circuit and noise models. With this, it is possible to assess whether the new fusions maintain favourable properties even in the presence of physical noise.

The Pauli error models for multi-way fusions subject to physical noise described in Chapter 4 provided an insight into the expected ratios and scaling of errors caused by different error mechanisms. The simple framework used to extract the models, however, presents some important limitations. The ZX-diagrams it propagates errors through fundamentally do not capture the details of the linear optical circuits they correspond to, and in fact the T1 fusion diagram is assumed to represent an average operation which is practically challenging to realise, and not likely to reflect the implementation used in a linear optical multi-way fusion circuit within an LOQC architecture. While this greatly simplifies the derivations and allows us to obtain a high level understanding of the performance of larger fusions, an accurate model can only be obtained by taking into account the specific structure of their linear optical implementation. Moreover, this methodology can only be applied to fusion circuits which can

be decomposed into constituent two-way fusions, such as standard n-GHZ fusions, but it cannot be used to model more general fusion circuits which do not fit easily into this description, like some of the depth-2 fusions described in Chapter 3.

In order to be able to analyse the impact of Pauli errors at different multi-way fusions on a realistic LO-FBQC architecture, and on the process of RSG in particular, it is necessary to extend this modelling framework to include other types of fusion circuits, which also requires it to be more closely based on the structure of their linear optical circuits. In this chapter we present a more general and accurate modelling method, based on Fock-state simulations of specific fusion circuits and postprocessing of the output of those simulations to extract Pauli error models. One important difference with the previous analysis is that here the photon addition model of multiphoton contamination is used, since this is considered to be more representative of the physical error mechanism and it does not introduce additional complexity. This framework is used to obtain models for standard n-GHZ fusions as well as other multi-way fusion circuits from Chapter 3, and a comparison is performed based on the total Pauli error rate and per-qubit marginal error rate metrics introduced in the previous chapter.

5.1 General modelling framework

5.1.1 Fock state simulations

The starting point of this modelling procedure consists of a series of Fock state simulations capturing the relevant combinations of physical mechanisms which can lead to unheralded errors at the fusion: distinguishability on one qubit, and the combination of multiphoton contamination and loss on two different qubits. In principle, these simulations can include higher-order error mechanisms as well. However, from the physical error parameter values in eq. 4.41 it can be seen that in our regime of interest the probability of a second-order distinguishability event is three orders of magnitude lower

than a first-order event, and for an n-input fusion there are only $n(n-1)$ configurations of this kind. Similarly, double multiphoton contamination and loss events are more than $2000\times$ less likely than single ones under the same set of parameters in eq. 4.41. Therefore, the Pauli error rate is dominated by leading-order contributions, so we choose to restrict our analysis in favour of making the output results computationally easier to obtain and simpler to parse. In the same spirit, the combinations of multiphoton contamination and loss on a single qubit and Pauli errors already present on the fusion inputs are also excluded from our simulations, since their effect is described by Pauli error channels which can be added in postprocessing. Under the assumption that all errors are either commuted to the inputs of the fusion or they are sufficiently small to be omitted from a leading-order analysis, these simulations involve an ideal fusion circuit with imperfect inputs.

The most general scenario is one in which all the input qubits are independent and equally likely to be in either computational state before they suffer any errors. A description of the channel due to the combined action of qubit errors and the fusion circuit on these inputs can be obtained by making use of the Choi-Jamiołkowski isomorphism [176, 177]. This relates a quantum channel $\mathcal{E}_S$ on a d-dimensional system S to an operator $\mathcal{E}_S \otimes \mathbb{I}_A$ on $S \otimes A$, where A is an ancilla system maximally entangled to S. Taking the initial state of S and A to be $\rho_{S,A} = |\Phi\rangle\langle\Phi|$, where $|\Phi\rangle = \frac{1}{\sqrt{d}} \sum_{i=0}^{d-1} |i\rangle_S |i\rangle_A$, and applying $\mathcal{E}_S \otimes \mathbb{I}_A$ to it allows to infer the action of the channel $\mathcal{E}_S$ directly from the joint state. In the case of the fully destructive fusions considered here, the S system is removed and the channel is obtained from the state of the ancilla A. Thanks to this, the operation of an n-way fusion on faulty input qubits can be extracted by preparing n ideal Bell pairs, applying any desired error channels to one qubit from each of them, and finally propagating those qubits through the fusion circuit. The channel corresponding to a detection pattern can then be reconstructed from the corresponding state of the untouched Bell pair qubits.

Separate simulations are performed for each of the n input configurations with one distinguishable qubit, and $n(n-1)$ with one lost and one multiphoton qubit. Each configuration is denoted by $\mathbf{Q}_i$, a

vector specifying the error transformation applied to each qubit, e.g., (ideal, ideal, multi, loss) for a four-way fusion. The corresponding set of channels for all possible detection patterns $x \in X$ gives a description of the action of the fusion on an input $\mathbf{Q}_i$ as a *quantum instrument* $\{\mathcal{E}_{x|\mathbf{Q}_i}\}_X$.

5.1.2 Pauli error models

Once the quantum instrument of a fusion circuit with a given input error configuration is available, the constituent channels are processed individually. For each of these, the full stabilizer group of the measurement operator is found by applying each of the 4^n n-qubit Paulis to it and identifying the ones which leave it unchanged up to a global π phase. Operators which are pure stabilizer state projections on all n qubits result in stabilizer groups with 2^n elements. The modelling steps outlined in the rest of this section apply to measurements of this kind and lead to a description of errors as Pauli flips of the ideal operators. It is assumed that Pauli twirling [172] is applied to the input qubits of the fusion, as described in Chapter 4, such that a distribution of Pauli operators is a valid representation of the error channel.

The quantum channels can be turned into *Pauli error models*, which describe them as a probability distribution of Pauli channels acting on the input qubits. This is done by comparing $\mathcal{E}_{x|\mathbf{Q}_i}$ with the ideal channel obtained from an error-free simulation, $\mathcal{E}_{x|\mathrm{id}}$. Every stabilizer state can be represented as a graph state up to local Clifford operations [178], with a procedure to perform this translation outlined in Ref. [179]. According to *Proposition 3* from Ref. [132], the set of states obtained by applying all combinations of single-qubit Z rotations to an n-qubit graph state is a basis for all n-qubit states. This means that the set of all possible Pauli errors, $\{P_j\}$, on the ideal channel operator can be obtained by mapping it to a graph state, defining the *local Z operators* (Z conjugated by the local Clifford) for each of its qubits, and finding all the possible combinations of these. From this it follows

that the non-ideal quantum channel can be expressed as:

$$\mathcal{E}_{x|\mathbf{Q}_i} = \sum_j \Pr(P_j, x|\mathbf{Q_i}) \left(P_j \mathcal{E}_{x|\mathrm{id}} P_j\right), \tag{5.1}$$

where the probabilities associated with each Pauli error P_j can be obtained as:

$$\Pr(P_j, x|\mathbf{Q_i}) = \mathrm{Tr}(\mathcal{E}_{x|\mathbf{Q}_i} P_j \mathcal{E}_{x|\mathrm{id}} P_j). \tag{5.2}$$

A Pauli error model consists of the distribution of these probabilities for every possible Pauli error P_j, and a separate model of this kind is obtained for each channel.

5.1.3 Fusion models

In a quantum instrument, different detection patterns x may correspond to the same quantum channel up to known Pauli rotations, so the corresponding measurement operators can be described by the same set of stabilizer generators up to phases. Since these phases do not change the entanglement structure of the operation, these detection patterns are considered to give equivalent outcomes and belong to a group denoted by X_O. The group is associated with an "ideal outcome" O, consisting of the shared set of stabilizer generators of the channel operator with all $+$ signs. For example, the success outcomes of a 3-GHZ fusion all have $O = \langle +XXX, +ZZI, +ZIZ \rangle$.

We are interested in calculating the error probabilities for each X_O as a whole, which requires combining the Pauli error models corresponding to all equivalent outcomes $x \in X_O$. For each detection pattern x, the ideal outcome O is identified. Then the $\Pr(P_j, x|\mathbf{Q_i})$ are translated into a distribution of probabilities $\Pr(\tilde{O}, x|\mathbf{Q}_i)$, where each $\tilde{O}$ is an outcome obtained by flipping the sign of all generators in O which anticommute with P_j. For example, any single-qubit Pauli Z error on a 3-GHZ fusion corresponds to $\tilde{O} = \langle -XXX, +ZZI, +ZIZ \rangle$. This is the same notation used in the probability distributions in Chapter 4, where a $-$ sign indicates an undesired flip of a stabilizer. Summing these

probabilities for all the detection patterns associated with the same O then gives a distribution of conditional probabilities which describe the error rates on the fusion outcome:

$$\Pr(\tilde{O}, X_O|\mathbf{Q}_i) = \sum_{x \in X_O} \Pr(\tilde{O}, x|\mathbf{Q}_i). \tag{5.3}$$

These conditional probabilities describe the action of a fusion given a specific configuration of input errors. For our purposes, it is often sufficient to have a more coarse-grained description of the behaviour of the fusion given a type of error, e.g., one distinguishable input qubit. We obtain this by combining the probabilities for different $\mathbf{Q}_i$ corresponding to the same type of error:

$$\Pr(\tilde{O}, X_O|E_i) = \sum_{\mathbf{Q}_i \in E_i} \Pr(\tilde{O}, X_O|\mathbf{Q}_i) \Pr(\mathbf{Q}_i), \tag{5.4}$$

where $E_i = \{\mathbf{Q}_i\}$ is the set of all input configurations for the given type of error. For example, for a 3-way fusion, E_{dist} is {(dist, ideal, ideal), (ideal, dist, ideal), (ideal, ideal, dist)}. In principle different $\Pr(\mathbf{Q}_i)$ could be used for each input configuration, but typically it is assumed that all $\mathbf{Q}_i$ are equally likely so that the expression in eq. 5.4 is a simple average. The set of the $\Pr(\tilde{O}, X_O|E_i)$ for all the E_i of interest, which we refer to as a *fusion model*, contains all the information needed to describe errors on the fusion as a distribution of Pauli errors.

The rate at which Pauli errors affect a fusion outcome is calculated starting from the fusion model coefficients $\Pr(\tilde{O}, X_O|E_i)$ from eq. 5.4. Bayes' theorem can be used to reverse the conditional probability:

$$\Pr(\tilde{O}, E_i|X_O) = \frac{\Pr(\tilde{O}, X_O|E_i) \Pr(E_i)}{\Pr(X_O)}, \tag{5.5}$$

where $\Pr(E_i)$ is the probability of an input error E_i occurring and $\Pr(X_O)$ the total probability of the

circuit heralding a pattern in X_O. The latter is calculated as:

$$\Pr(X_O) = \sum_{E_i} \sum_{\tilde{O}} \Pr(\tilde{O}, X_O | E_i) \Pr(E_i) = \sum_{E_i} \Pr(X_O | E_i) \Pr(E_i), \tag{5.6}$$

where $\Pr(X_O|E_i)$ is the probability of obtaining outcome X_O in the presence of input error E_i. For example, we saw in Chapter 4 that this is $2/3$ for the success outcome of a T2 fusion in the presence of multiphoton contamination and loss. A Pauli error model for the fusion, expressed as a conditional probability distribution of Pauli flips as in eqs. 4.29 and 4.38, is then simply obtained by summing over all error types of interest:

$$\Pr(\tilde{O}|X_O) = \sum_{E_i} \Pr(\tilde{O}, E_i | X_O). \tag{5.7}$$

Under the assumption that the dominant contribution to multiphoton contamination and loss on the same qubit is one in which the latter occurs after the former, this type of error causes a fully depolarizing channel on the input qubit with probability $1/4$, as seen in Chapter 4, so each $\Pr(\tilde{O}|X_O)$ has a $\frac{1}{12} p_m p_l$ term added for each single-qubit Pauli error consistent with its set of stabilizer flips. In the following analysis we only consider this contribution, but it would be straightforward to include the case of the two errors occurring in opposite order simply by adding an extra type of error "loss + multi" and the associated probabilities. Similarly, a p_P term is added for each input Pauli error P.

5.1.4 Pauli error metrics

In Chapter 4, two error metrics were introduced to quantify and compare the Pauli error rates of different fusions: the total Pauli error rate and the per-qubit marginal error rate. The former is obtained by summing over all $\tilde{O}$ in eq. 5.7:

$$\Pr(\text{Pauli}|X_O) = \sum_{\tilde{O}} \Pr(\tilde{O}|X_O). \tag{5.8}$$

The contribution from each type of input error is obtained by doing the same for each error term separately, i.e.:

$$\Pr(\text{Pauli}, E_i|X_O) = \sum_{\tilde{O}\neq O} \Pr(\tilde{O}, E_i|X_O). \tag{5.9}$$

The second metric is the *per-qubit marginal Pauli rate*. The first step to obtain it is to translate each set of stabilizer generator flips in $\tilde{O}$ into Pauli errors on qubits. Care is taken to only include the lowest-weight Paulis compatible with each $\tilde{O}$, and also to symmetrize the errors over the input qubits. For example, the 3-GHZ fusion outcome $\tilde{O} = \langle -XXX, +ZZI, +ZIZ\rangle$ is turned into three equally likely single-qubit Paulis, (ZII, IZI, IIZ). This procedure transforms the distribution of probabilities $\Pr(\tilde{O}, E_i|X_O)$ into $\Pr(P_j, E_i|X_O)$, which can then be marginalized to obtain the probabilities of each qubit suffering a Pauli error $P \in \{X, Y, Z\}$. Then, the marginal per-qubit Pauli probability due to each type of error, $\Pr(P, E_i|X_O)$, is obtained by summing over all P_j's which contain the Pauli operator P and dividing over the number of qubits n. The total marginal Pauli error probability for a given fusion outcome is then obtained as:

$$\Pr(\text{Pauli (m)}|X_O) = \sum_P \left(\sum_{E_i} \Pr(P, E_i|X_O) \right). \tag{5.10}$$

The steps of the framework described so far apply to any fusion outcome X_O. In our analysis, we focus on the success outcome, $X_O = \text{succ}$, since we are mostly interested in the impact of Pauli errors in RSG, where it is assumed that multiplexing is used to discard failed fusion attempts.

5.2 Pauli errors in standard n-GHZ fusions revisited

The first application of the modelling framework just described is to revisit the analysis of standard n-GHZ fusions, to verify the conclusions drawn in Chapter 4. In order to calculate the Pauli error models, the total probability of each type of input error, $\Pr(E_i)$, has to be defined. This is the sum of the probabilities of all error configurations $\mathbf{Q}_j$ in E_i, which depends on the physical error model

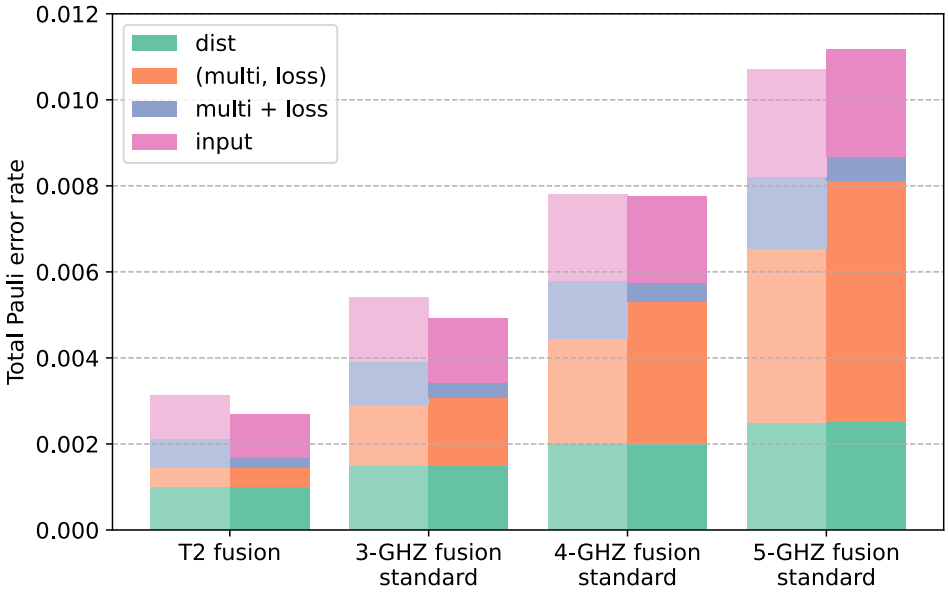

Figure 5.1: Total Pauli error rate broken down into contributions from different input errors for the success outcomes of standard n-GHZ fusions with $2 \leq n \leq 5$. Darker colours correspond to values obtained using the modelling framework described in this chapter, while lighter colours are used for the values from fig. 4.7 in Chapter 4. The values of the physical error rates are the same as previously used, given in eq. 4.41. The total Pauli error rate increases with n due to the growth of all contributions. Most of them scale linearly with n, except the combination of multiphoton contamination and loss on different qubits (orange), which scales as $n^2 - n$.

considered. Here we assume that the rate of each physical error is the same for all input qubits, which gives the following $\Pr(E_i)$:

$$\begin{aligned}
\Pr(E_i = \text{dist}) &= n \times p_d, \\
\Pr(E_i = (\text{multi, loss})) &= n(n-1) \times p_l p_m, \\
\Pr(E_i = (\text{multi + loss})) &= n \times p_l p_m \\
\Pr(E_i = X) &= n \times p_X, \\
\Pr(E_i = Y) &= n \times p_Y, \\
\Pr(E_i = Z) &= n \times p_Z,
\end{aligned} \tag{5.11}$$

where p_d, p_l, p_m, p_X, p_Y and p_Z are the total per-qubit distinguishability, loss, multiphoton and single-qubit Pauli rates commuted to the input of the fusion. Their values are the same as in the previous chapter, given in eq. 4.41.

Fig. 5.1 shows the total Pauli error rates obtained in this modelling framework for n-GHZ fusions with $2 \leq n \leq 5$ (darker colours), side-by-side with the values from fig. 4.7 (lighter colours). The same breakdown of the probabilities into contributions from different error mechanisms is shown here, with the "input" label referring to any Pauli already present on the input qubits. The total Pauli error rates found here are similar to those obtained in Chapter 4, although systematically higher for all n. The contributions from distinguishability and input errors are identical, but there are noticeable differences in the impact of multiphoton contamination and loss. The photon addition model used here only leads to a Pauli error on the input qubit with probability $1/4$ when combined with loss, compared to $3/4$ with the maximally mixed model previously considered. This reduction is reflected in the much smaller height of the blue bars. At the same time, the Pauli error rate associated with these two physical errors affecting different qubits is higher in this model. The reason for this is a difference in the reduction in success probability in the two cases. Here it is found that all standard n-GHZ fusions have their success rate reduced by $1/3$ for (multi, loss) input configurations regardless of their size, whereas larger fusions would see larger drops in success probability when modelled through the composition of two-way error models, due to the accumulation of $2/3$ factors. Since the results in fig. 5.1 are based on a more realistic multiphoton contamination model and on the simulation of a specific linear optical circuit, we expect the trends shown in this plot to be closer to the behaviour of an actual fusion circuit.

Fig. 5.2 shows the marginal Pauli error rates for the same fusion circuits, again directly compared to the plot in fig. 4.8 shown in lighter colours. The observations made in the case of the total Pauli error rate apply here as well, with these marginal error rates being consistently higher than those previously found, due to differences in the error probabilities due to combinations of multiphoton contamination and loss. The discrepancy between these values and the ones found in Chapter 4 is larger for these marginal error rates, which can be explained by the fact that, for $n > 3$, (multi, loss) configurations lead to correlated Pauli error channels which have a significant impact on the height of the orange

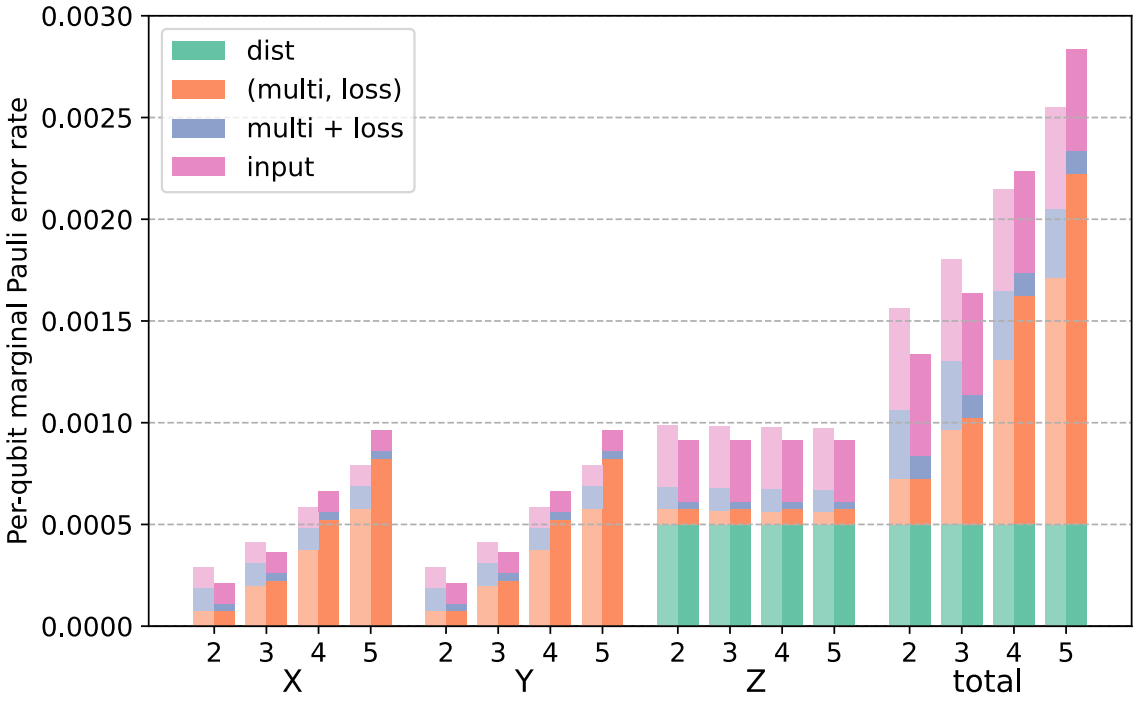

Figure 5.2: Marginal per-qubit Pauli error rate for standard n-GHZ fusions with $2 \leq n \leq 5$. Darker colours correspond to values obtained using the modelling framework described in this chapter, while lighter colours are used for the values from fig. 4.8 in Chapter 4. The breakdown into contributions from different sources, as well as into different single-qubit Pauli errors, is shown. Similarly to the total Pauli error rate, the marginalized total (rightmost) increases with n due to the growth of all contributions.

bars in this plot.

While these plots confirm that larger n-way fusions have a higher probability of suffering Pauli errors, they do not directly rule out their usefulness to improve the performance of RSG schemes. In fact, in that context the error rate of individual fusion circuits has to be considered in conjunction with the number of fusions in the RSG scheme as a whole, which decreases with n. This trade-off is analysed in detail in Chapter 6.

5.3 Pauli errors in 4-GHZ fusion circuits

The modelling methodology described in sec. 5.1 above can be applied to any fully destructive linear optical fusion circuit, lifting the restriction to circuits which can be decomposed into constituent T1 and T2 fusions imposed by the framework described in Chapter 4. This allows us to extend

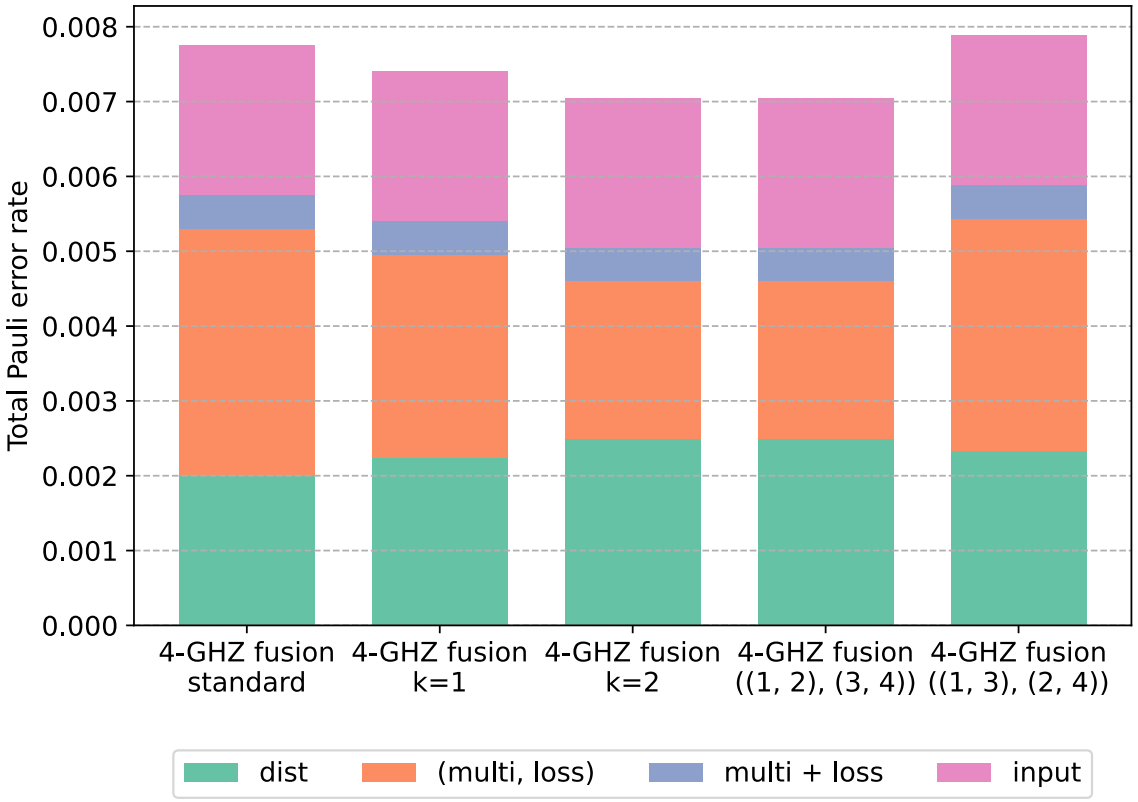

Figure 5.3: Total Pauli error rate for the success outcome of different types of 4-GHZ fusion circuits introduced in Chapter 3. While all fusions perform the same operation in the ideal case, the rate at which the outcome is affected by a Pauli error changes based on the structure of the circuit.

the analysis to the fusion circuits with additional input beamsplitters introduced in Chapter 3, some of which cannot be described simply as concatenations of two-way fusions. In the ideal case, these circuits were found to have favourable properties compared to standard n-GHZ fusions, namely higher success probability or partially entangling failure outcomes. Understanding their behaviour in the presence of errors is important to determine whether they also maintain an advantage in a realistic noisy scenario. $n = 4$ is the smallest fusion size for which representatives of each kind of circuit with interesting properties exist, so we focus on 4-GHZ fusion circuits here. In particular, we compare the performance of five variants with respect to Pauli errors: the standard 4-GHZ fusion, analysed above, the $k = 1$ and $k = 2$ circuits with one pair of additional beamsplitters (see fig. 3.7) and the representatives of the two families of depth-2 $n = 4$ fusions, labelled $((1,2),(3,4))$ and $((1,3),(2,4))$ (see fig. 3.9). Pauli error models are extracted for each of these circuits and used to calculate their total Pauli error rate and per-qubit marginal Pauli error rate.

The total Pauli rate for the success outcome of the five 4-GHZ fusions is shown in fig. 5.3. All five circuits perform the same operation in the ideal case, but their different structures are seen to affect the rate at which the measurement operator suffers a Pauli error in the presence of noise. The total error rates due to input Pauli errors (pink) and multiphoton contamination and loss on the same qubit (blue) are only affected by the circuit size and not by its structure, so their contributions are identical for all five fusions. On the other hand, the probability of a distinguishable input photon causing a Pauli error (green) is lowest for the standard 4-GHZ fusion and has higher values for other circuits, suggesting that the impact of this source of errors is increased by the addition of extra beamsplitters to the 4-GHZ analyzer. At the same time, the rate of Pauli errors due to combinations of multiphoton contamination and loss on different qubits (orange) is reduced by a larger amount for most of the new circuits, with the high-probability depth-2 fusion being the exception. This decrease is due to the fact that the success probability of the fusions is reduced by a larger amount in the presence of these errors, a phenomenon we refer to as *multiphoton error filtering*, described in more detail in sec. 5.3.1 below. When these changes to the various Pauli error contributions are combined, the result shows that some of the 4-GHZ fusions with extra beamsplitters outperform the standard circuit with respect to total Pauli errors. The ((1, 2), (3, 4)) and $k = 2$ fusions have the same total Pauli error rate and the lowest of the analyzed circuits, reduced by $\sim 9\%$ compared to the standard fusion. In the case of the $k = 1$ fusion, the reduction is smaller, $\sim 5\%$, while the high-probability ((1, 3), (2, 4)) fusion has a total Pauli error rate $\sim 2\%$ higher than the standard circuit. The specific values and relative differences obtained here are highly sensitive to the choice of physical error parameters made in eq. 4.41. While we have chosen these to be representative of expected noise rates in an LO-FBQC scheme, it is important to reassess the performance of these fusions under the parameter values which best reflect the architecture under consideration. Other hardware platforms may have significantly different physical error rates, which can result in changes in the relative performance of these fusion circuits and lead to different fusion choices in the design. This is discussed further in Appendix C,

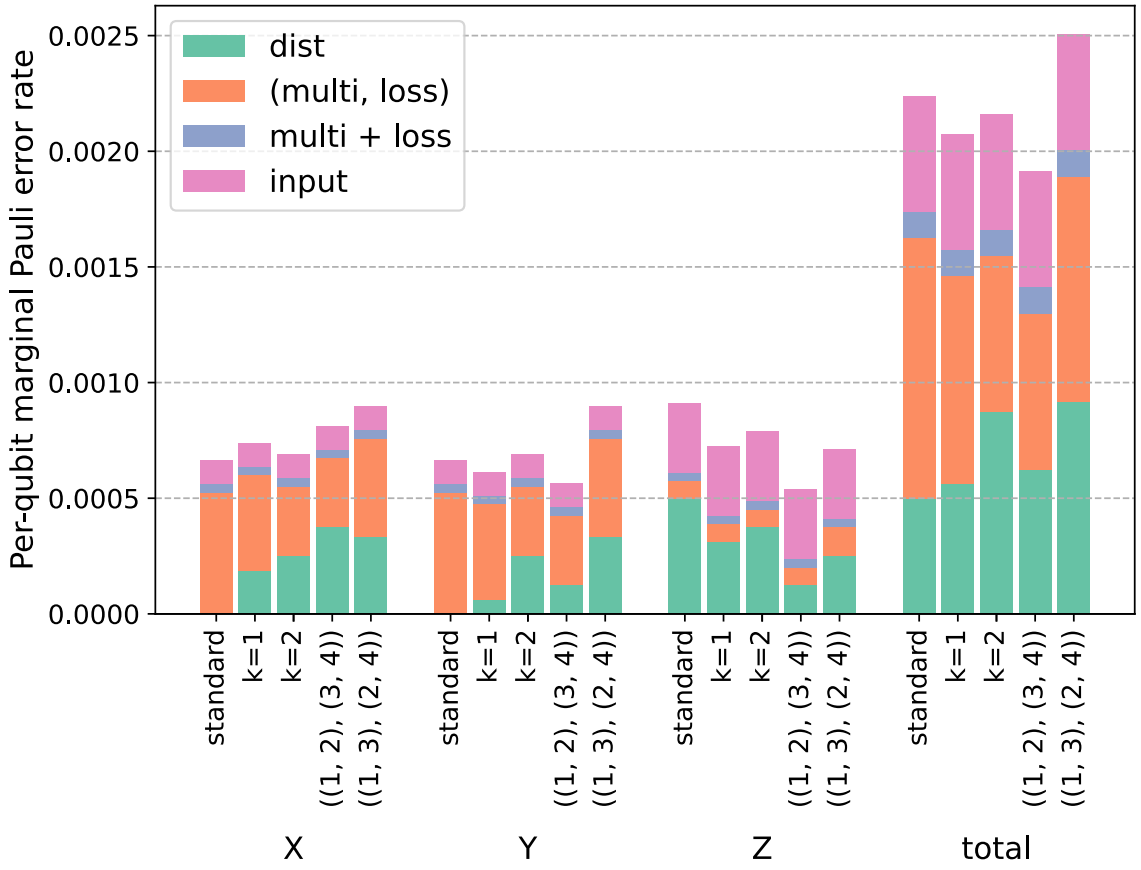

Figure 5.4: Marginal per-qubit Pauli error rate for the success outcome of different types of 4-GHZ fusion circuits. The relative performance of the fusions is different than suggested by the total Pauli error rate, suggesting that the internal structure of the distribution of errors, and correlations in particular, play a significant role.

and one example of a different physical error parameter regime is considered.

The marginal per-qubit Pauli error rate, shown in fig. 5.4, paints a slightly different picture. The $k = 1$, $k = 2$ and $((1,2),(3,4))$ circuits still outperform the standard 4-GHZ fusion, but now the latter has a significantly lower Pauli error rate than the standard circuit ($\sim 15\%$ reduction), followed by the $k = 1$ circuit ($\sim 7\%$ reduction), and the $k = 2$ fusion ($\sim 3\%$ reduction). The substantial difference between the performance of the latter circuit with respect to the two Pauli metrics, in particular the fact that it has the same total Pauli error rate as the ((1, 2), (3, 4)) fusion but higher marginal Pauli error rate, suggests that the internal structure of the error distribution plays an important role in differentiating these circuits. In order to obtain a clearer picture of how these circuits compare, it is useful to inspect the full probability distributions of Pauli errors due to different physical errors.

(a) "dist" Pauli error probability distribution

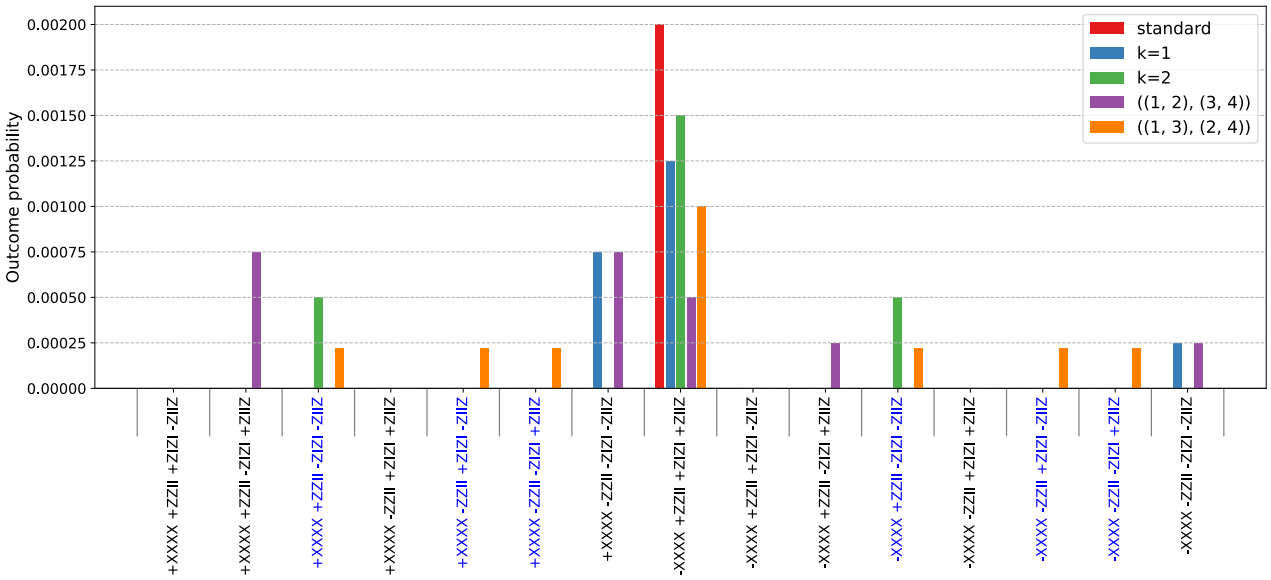

(b) "(multi, loss)" Pauli error probability distribution

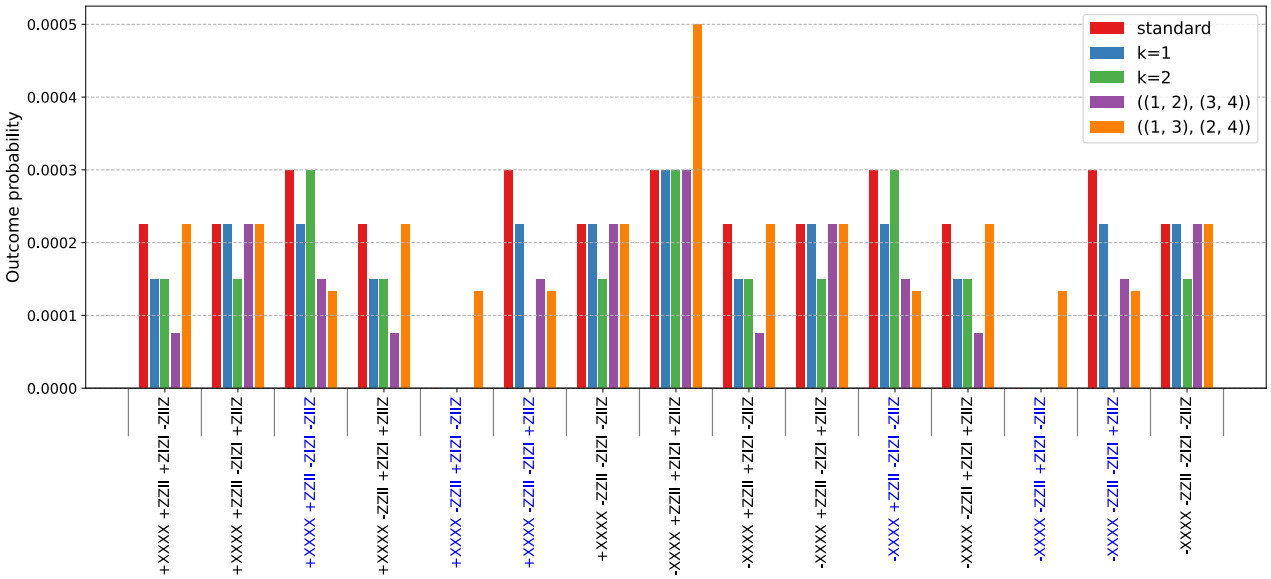

Figure 5.5: Distributions of Pauli errors on the success outcome of different 4-GHZ fusion circuits. Pauli flips are expressed as '-' signs on stabilizer generators of the measurement operator, as described in sec. 5.1. Flips corresponding to correlated errors are highlighted in blue. (a) Pauli errors due to distinguishability at the input of the fusion. (b) Pauli errors due to combinations of multiphoton contamination and loss on different input qubits.

Fig. 5.5 shows the probabilities of each of the 15 possible Pauli flips of the stabilizer generators of the 4-GHZ projection due to distinguishability, in (a), and to multiphoton and loss combinations, in (b). Fig. 5.3 previously showed that the circuits with extra beamsplitters have a higher total Pauli error rate associated with distinguishability. Here it can be seen that the form of the errors is also different. In the standard 4-GHZ fusion, distinguishability can only flip the $XXXX$ stabilizer generator, but in the

new fusions other types of errors can also occur. In particular, distinguishability in the $k = 2$ and ((1, 3), (2, 4)) fusions can lead to correlated XX and XY errors, corresponding to outcomes highlighted in blue on the x axis, which explains the steep increase in the corresponding contribution to the marginal Pauli error rate for these circuits seen in fig. 5.4. Combinations of multiphoton and loss lead to all single-qubit errors in all fusions. One notable difference is that the $((1,3),(2,4))$ circuit is the only one with a non-zero probability of suffering every possible correlated flip, whereas the $k = 2$ circuit can only be affected by two of the six types of correlated error, and the other three fusions by four. In terms of absolute probability of correlated errors, the standard 4-GHZ fusion has the highest rate ($\sim 0.12\%$) and the $k = 2$ and $((1,2),(3,4))$ circuits the lowest ($\sim 0.06\%$).

This analysis shows that the performance of n-GHZ fusions with respect to Pauli errors is dictated not only by their size, but also by the structure of the circuits. Considering that the $k = 1$, $k = 2$ and $((1,2),(3,4))$ 4-GHZ fusions all have the same ideal success probability as the standard 4-GHZ fusion, the results presented above suggest that the new circuits allow us to reduce the Pauli error rate without affecting the resource overhead. The maximum improvement in the total Pauli error rate is $\sim 9\%$ when the standard 4-GHZ fusion is replaced by either the $k = 2$ or the ((1, 2), (3, 4)) circuit. This metric does not distinguish between single-qubit and correlated multi-qubit errors, so it is representative of the gain expected in a scenario where the two can be regarded as equally detrimental. On the other hand, when correlations are considered as damaging as two independent single qubit errors, the ((1, 2), (3, 4)) circuit allows a higher maximum improvement of $\sim 15\%$ with respect to the standard fusion. These circuits also have entangling failure outcomes, which can in principle be used to reduce the resource overhead too. The $((1,3),(2,4))$ 4-GHZ fusion is found to have higher Pauli error rates according to both metrics. Given the significant boost in success probability it provides compared to the standard 4-GHZ fusion ($2.25\times$), however, it is likely that the trade-off between error and footprint performance can still be favourable for this circuit.

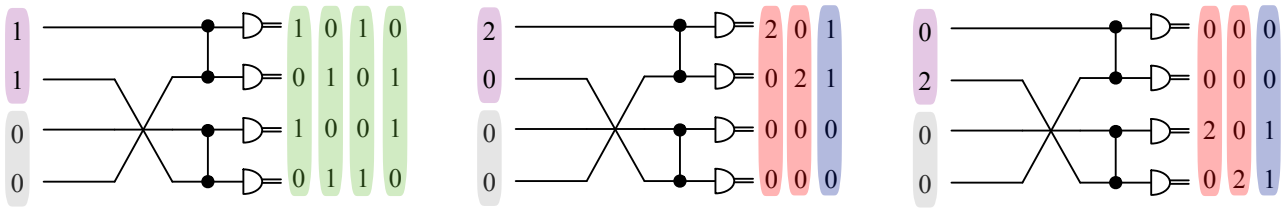

Figure 5.6: Propagation of different two-photon terms (purple) paired with a vacuum qubit (grey) through the T2 fusion circuit. The antibunched term deterministically results in success detection patterns (green), while bunched patterns can only result in failure (red) or nonsense (blue) patterns.

5.3.1 Multiphoton error filtering

This analysis of 4-GHZ fusions shows that the structure of a fusion circuit has a noticeable impact on the rate at which physical errors translate into Pauli errors on its measurement outcomes. In particular, some fusion variants appear to have a significantly reduced error rate caused by combinations of multiphoton and loss on different qubits, as shown clearly in fig. 5.3. It is important to understand what properties a fusion circuit make it more tolerant to these kinds of errors, which may help identify or design circuits which exploit similar mechanisms.

We have seen that in all the n-GHZ fusion circuits analysed so far, including the T2 and larger standard fusions, a combination of multiphoton contamination and loss on two different inputs leads to a reduced success probability. Consider the simplest case of a T2 fusion circuit. The state of a qubit which has suffered multiphoton contamination contains all three two-photon dual rails terms $\{|20\rangle\rangle, |02\rangle\rangle, |11\rangle\rangle\}$ with equal probabilities, while the other qubit is in the vacuum state $|00\rangle\rangle$. When the two bunched terms are propagated through the T2 fusion circuit, they can only result in either a failure detection pattern, with two photons at one detector and none in the others, or in a nonsense pattern, with two antibunched photons after a single beamsplitter. Therefore, the fusion cannot succeed in this case. On the other hand, the $|11\rangle\rangle$ term paired with a vacuum input deterministically produces a detection pattern which heralds success, with one photon after each beamsplitter. The evolution of each of these scenarios through the T2 fusion circuit is shown in fig. 5.6. As a result, the success probability of the fusion in this case is $\frac{2}{3} \times 0 + \frac{1}{3} \times 1 = \frac{1}{3}$, i.e., $\frac{2}{3}\times$ that of the ideal case.

This reduction in success probability is a desirable feature of the fusion, since it converts potential unheralded errors into failures or heralded errors, which can be discarded via multiplexing. For this reason, we refer to this mechanism as *multiphoton error filtering*. The reasoning given above for a T2 fusion, applies to all standard n-GHZ fusion circuits, whose success probability is reduced by a factor of $\frac{1}{3}$ independent of n. However, this ratio can vary for circuits with different structures. The amount of filtering for a given fusion can be quantified by calculating its average success probability given (multi, loss) input errors using the model coefficients in eq. 5.4:

$$\Pr(X_O|(\text{multi, loss})) = \frac{1}{n(n-1)} \sum_{\mathbf{Q}_i \in (\text{multi, loss})} \left(\sum_{\tilde{O}} \Pr(\tilde{O}, X_O|\mathbf{Q}_i) \right), \tag{5.12}$$

where $\tilde{O}$ includes the ideal outcome with no flips, and X_O is the success outcome in our analysis. In general different input configurations $\mathbf{Q}_i$ may lead to different success probabilities, so we consider the average over all possible combinations of inputs of this kind here. The ratio of this quantity with the heralding probability of the same fusion with ideal inputs gives us a sense of its error filtering properties. This calculation can be performed for any E_i, not only (multi, loss), but we focus on this type of error here since it is the one whose input was found to be reduced in the previous analysis.

Fig. 5.7 shows this ratio for the 4-GHZ fusions analysed in sec. 5.3. The higher probability ((1, 3), (2, 4)) circuit has the same behaviour as the standard fusion. The three circuits which were found to have a reduced Pauli error rate associated with this type of input error are all found to perform more filtering than the standard 4-GHZ fusion, with the success probability of the $k = 2$ and $((1,2),(3,4))$ being reduced to $\frac{4}{9}\times$ the ideal one. This suggests that filtering plays an important role in mitigating the impact of these (multi, loss) errors on the Pauli error rate of these fusion circuits.

To understand how higher amounts of filtering are achieved, consider the example of the $k = 1$ 4-GHZ fusion with one multiphoton and one lost input. It is useful to refer to the decomposition of this circuit into T1 and T2 circuits, shown in fig. 5.8. If qubit 2 is in a multiphoton state and qubit 4 is lost, the

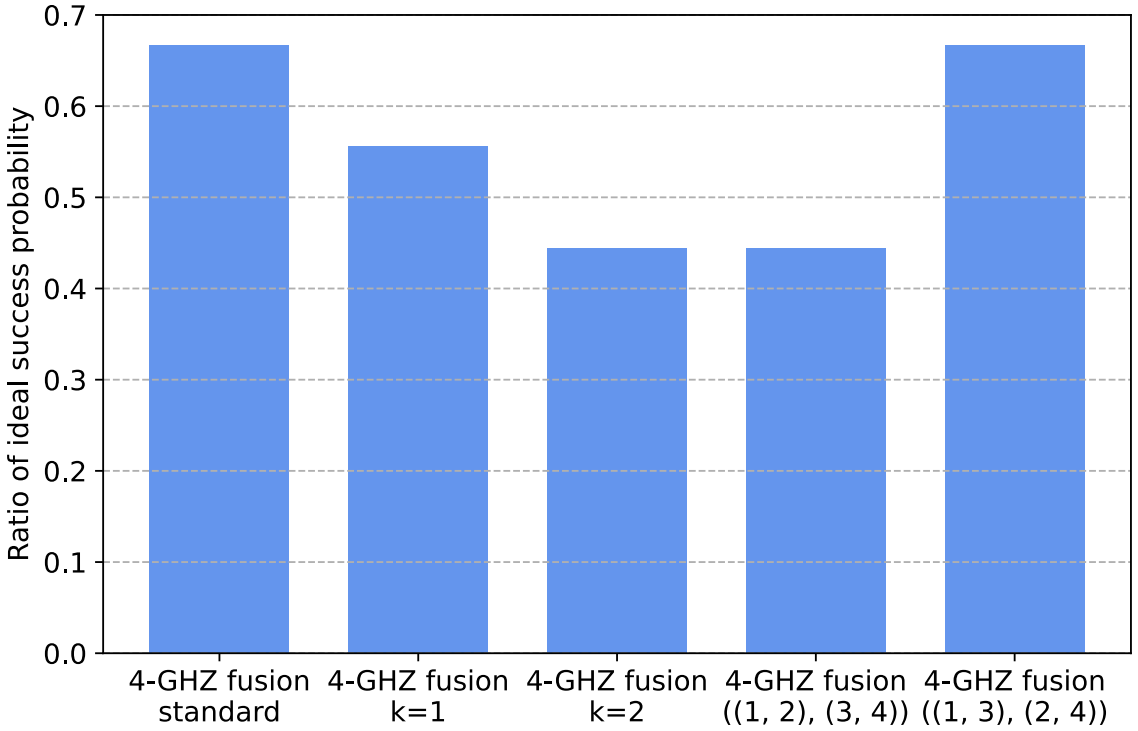

Figure 5.7: Average ratio of the ideal success probability in the presence of multiphoton contamination and loss on different input qubits of various 4-GHZ fusion circuits. The success probability is reduced by a larger amount for the three circuits found to have better tolerance for (multi, loss) errors in the previous section.

only two-photon term which can survive past the second T1 fusion without causing a failure is $|11\rangle\rangle$. When evolved through the extra beamsplitter at the input of the T2 this becomes a superposition of $|20\rangle\rangle$ and $|02\rangle\rangle$, which, combined with the $|00\rangle\rangle$ state of qubit 4, cannot lead to success, as discussed above. So this configuration of input multiphoton and loss errors cannot produce a success outcome in this circuit, i.e., it is fully filtered. A similar mechanism applied to other configurations leads to partial filtering, whereby the T2 outcome can succeed, but at a lower rate. This is the case for all configurations in the case of the $k = 2$ 4-GHZ fusion, for which most of the input combinations lead to a probability reduction factor of $\frac{1}{3}$, resulting in a lower average success probability than the standard fusion. In the case of the ((1, 2), (3, 4)) circuit two input configurations are fully filtered, while the other ones are partially filtered, with some leading to a $\frac{1}{3}\times$ reduction factor and others to $\frac{2}{3}\times$, so that the average reduction is the same as for the $k = 2$ circuit.

The multiphoton filtering described here is an example of a technique which reduces the impact of a

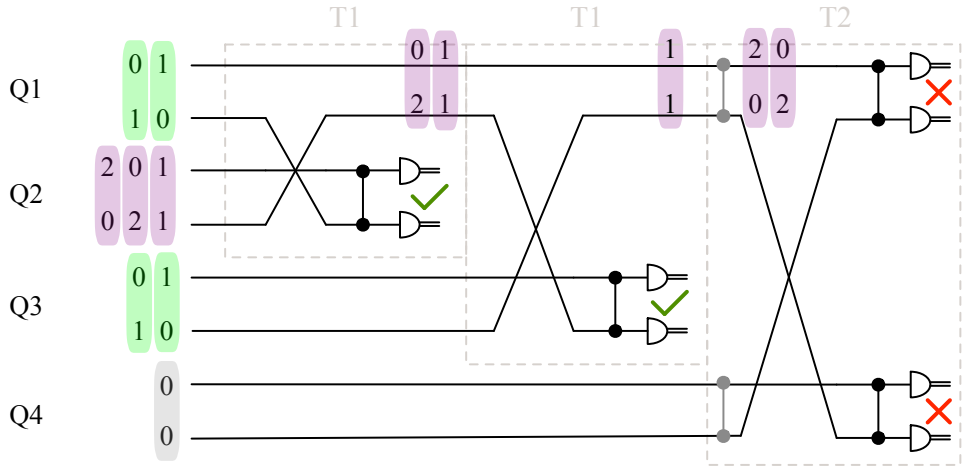

Figure 5.8: Error filtering in the $k = 1$ 4-GHZ fusion. The possible terms for each input qubit are shown to the left of the circuit. If qubit 2 is in a multiphoton state (purple) and qubit 4 is lost (grey), there is no way that all three constituent fusions can succeed, so the whole fusion fails deterministically.

specific kind of unheralded error mechanism by transforming it into a heralded failure or error, which is particularly helpful when multiplexing can be used to discard such outcomes, e.g., in RSG. As mentioned above, filtering can also work for distinguishability errors, but the results in the previous section suggest that any reduction in success probability obtained in that case is not sufficient to outweigh the increase in the Pauli error rate they cause. In the next section we present another filtering method which can help mitigate these kinds of errors as well.

5.3.2 Detection pattern filtering

In the Pauli error analysis of fusions performed so far, a single success outcome has been considered which includes all possible detection patterns heralding a successful n-GHZ projection. The modelling framework allows for a more fine-grained analysis of individual detection patterns, or groups of them, which can provide an insight into differences in the "quality" of the projection performed based on the observed heralding signal. Some detection patterns may be found to correspond to significantly lower error rates than others, for example, or to different distributions of errors, and this information can be used to discard or deprioritise outcomes associated with higher error probabilities.

Consider the example of the ((1, 3), (2, 4)) 4-GHZ circuit, which heralds success when two photons

E_i	X_O^i	Total Pauli	X	Y	Z	XX	XY
dist	anti	1/2	0	0	1/2	0	0
	bunch	3/4	0	0	1/4	1/4	1/4
(multi, loss)	anti	3/8	1/8	1/8	1/8	0	0
	bunch	17/28	5/28	5/28	3/28	1/14	1/14

Table 5.1: Conditional probabilities of different types of errors in the ((1, 3), (2, 4)) 4-GHZ fusion, given the detection of a pattern belonging to group X_O^i and an input error configuration of type E_i.

are detected after each H_4 circuit and at least two of the photons are antibunched. The detection patterns can be split into two groups: fully antibunched patterns, X_O^{anti}, where two single photons are detected at different detectors at the output of each H_4, and partially bunched patterns, X_O^{bunch}, where two photons go to the same detector and the other two are separate. There are 12 patterns of the former kind, each occurring with probability $1/128$, and 48 in the latter group, each registered with probability $1/256$. So the success probability is split between the two groups as:

$$\begin{aligned} p_{\text{succ}}^{\text{anti}} &= 12 \times \frac{1}{128} = \frac{3}{32}, \\ p_{\text{succ}}^{\text{bunch}} &= 48 \times \frac{1}{256} = \frac{3}{16}. \end{aligned} \tag{5.13}$$

Performing the Pauli error analysis separately for the two groups reveals significant differences in the error properties of the measurements performed in the two cases. Table 5.1 summarizes the conditional probabilities of different Pauli errors in the presence of distinguishability or combinations of multiphoton and loss, given that a pattern of one of the two types is observed. Fully antibunched patterns clearly lead to a substantially lower probability of a Pauli error on the resulting 4-GHZ projection, in addition to removing all correlated errors. The probability of distinguishability causing a Pauli error is 66% that of partially bunched patterns, and for combinations of multiphoton contamination and loss the ratio is 62%. This suggests that discarding the partially bunched patterns would result in a significant improvement in the Pauli error rate of the fusion. Fig. 5.9 shows the total Pauli error rate obtained when postselecting only on the fully antibunched patterns compared to that associated with all success patterns. The result of removing the partially bunched patterns is a very steep

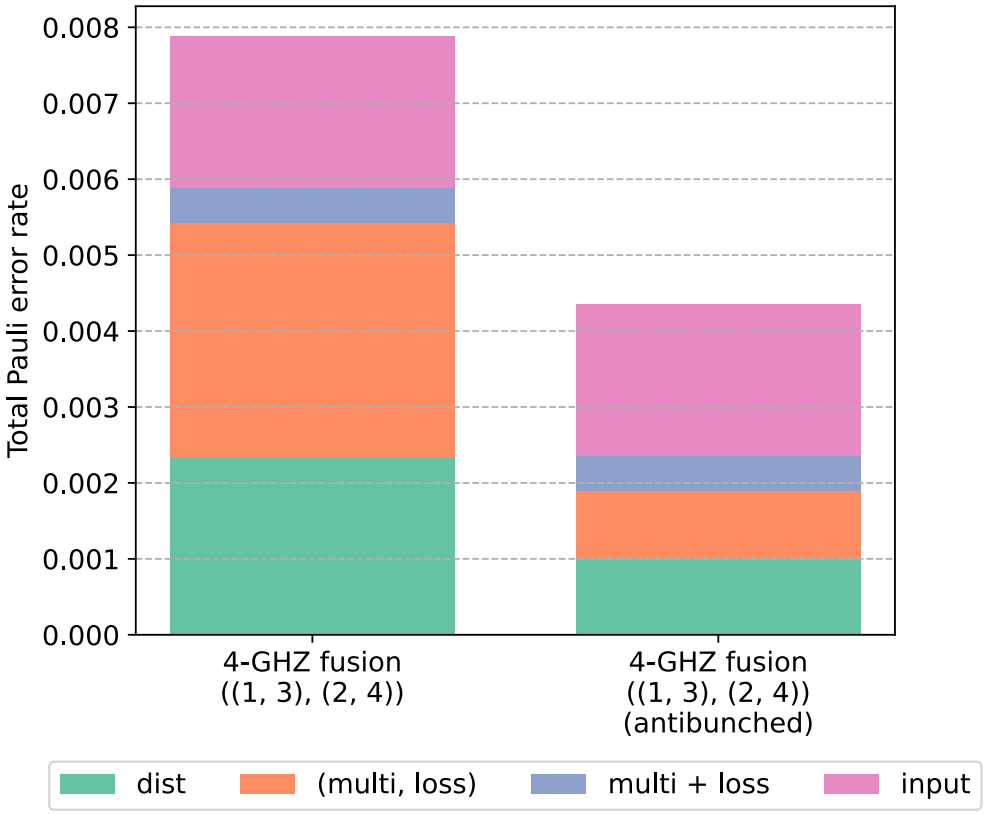

Figure 5.9: Comparison of the total Pauli error rate for the ((1, 3), (2, 4)) 4-GHZ fusion when all possible success patterns are considered (left) and when only the fully antibunched patterns are postselected on (right). Removing the partially bunched patterns reduces the Pauli error probability by $\sim 43\%$.

reduction in the total Pauli error rate by $\sim 45\%$, bringing it to a value close to that of the standard 3-GHZ fusion in fig. 5.1.

We refer to this method of reducing the Pauli error rate of a fusion by selecting specific detection patterns with favourable properties as *detection pattern filtering*. The ((1, 3), (2, 4)) 4-GHZ fusion was considered here, which was found to have a higher total Pauli error rate than the other fusions of the same size, as an example of how the response of a fusion to input errors can be improved just by changing the way its classical heralding information is processed. This applies to any other fusion circuit whose detection patterns are found to have different error properties. Of course, the reduction in Pauli errors comes at the cost of a decrease in success probability, which in the example given here is only $1/3$ of the full success rate upon filtering. This introduces a trade-off between success probability, linked to resource overhead, and Pauli rate, linked to error tolerance. Steep reductions

in success probability are unlikely to be advantageous at every multi-way fusion, since they would negate their advantage in terms of resource consumption. However, it is possible that in certain locations in RSG where Pauli errors are particularly harmful, it is worthwhile to sacrifice the success rate to obtain a projection of higher quality. This is the type of consideration that is essential in the process of RSG design to obtain schemes with good performance.

5.4 Conclusion

In this chapter, a general modelling framework was introduced which allows us to model the impact of physical errors on linear optical fusions as a probability distribution of Pauli error channels on the input qubits. This process is based on Fock state simulations of the fusion circuits, followed by a number of postprocessing steps to convert the extracted quantum channels into a stochastic Pauli error model. Compared to the methodology described in the previous chapter, this procedure can be applied to any of the fully-destructive fusion circuits presented in Chapter 3 and more, it captures the specific structure of the linear optical circuit without relying on abstract subcircuit descriptions, and it uses a more physically motivated multiphoton contamination model.

The results obtained in Chapter 4 on the scaling of Pauli errors due to different physical mechanisms were confirmed by this kind of analysis. In addition, a comparison of different variants of 4-GHZ fusions presented in Chapter 3 was possible, which highlighted the importance of the structure of the linear optical circuit in determining the Pauli error response of a fusion. Some of the new fusions with entangling failure outcomes were found to lead to lower Pauli error rates compared to the standard 4-GHZ fusion, thanks to their ability to filter out combinations of multiphoton contamination and loss. The high-probability circuit does not perform as well with respect to Pauli errors, but a technique was introduced which allows us to trade off success probability against Pauli error rate simply by filtering out some of the detection patterns. Fusions with $n = 4$ were considered here as an illustrative

example, but this kind of analysis and techniques can be easily applied to fusion circuits of any size.

Chapter 6

Resource State Generation with Multi-Way Fusions

This chapter brings together the results of Chapters 3- 5 to answer the question posed at the beginning of this thesis of whether the use of multi-way fusions can improve the performance of Resource State Generation (RSG) schemes, both in terms of the cost and error properties of the target resource state. Methods to estimate the impact of larger fusions on RSG in general are presented and used to obtain a high-level understanding of their potential. Then, techniques to perform calculations for specific RSG schemes are described and applied to the generation of two resource states taken from existing LO-FBQC proposals.

Up to this point, we have focused on understanding the behaviour of multi-way fusion circuits in isolation, both in the ideal case and in the presence of physical errors. In this chapter, we use the modelling framework and results obtained in the previous chapters to assess the effect of using linear optical n-GHZ fusions in the generation of resource states on both their cost and error rates. With this analysis, we test the intuition presented in Chapter 2, Sec. 2.3.3 that multi-way fusions could simultaneously reduce the resource overhead and error rate of linear optical RSG, which was the original motivation behind the work described in this thesis. As a reminder, this starting hypothesis is based on the fact that while the success probability and total error rate of linear optical fusions decrease with their size, n, so does the number of fusions and seed states required to make a fixed target resource

state. The analysis described here aims to determine which of the two effects dominates in the process of RSG, and thus whether multi-way fusions provide an advantage.

To start, it is useful to perform a more abstract comparison of the impact of different fusion circuits on the process of RSG. In this chapter, we introduce metrics which allow to do this kind of mapping between the properties of a fusion circuit and its expected effect on the footprint and error rate of a resource state generated using it, based on a simplified model of RSG. These quantities are helpful to compare several of the fusions introduced in Chapter 3 and to identify promising candidate fusion circuits, allowing to reduce the search space for RSG design.

The optimal choice of fusions to generate a given target state ultimately depends on the specific structure of the state itself. Finding the best protocol to generate a target resource state requires the ability to evaluate different RSG schemes. The performance of different protocols can only be accurately characterised with involved modelling and simulations to take into account all aspects of the RSG scheme, which is not a suitable approach to rapidly sift through the many available options and to identify the optimal design. Here we present simpler heuristic methods which can be used to approximate the cost and error rates of any RSG scheme, providing a path to rapid evaluation and optimisation of designs. At the end of this chapter, this type of calculation is used to compare RSG schemes with only T2 fusions to ones with multi-way fusions for two example resource states.

6.1 Estimating RSG performance

In order to allow a general assessment of RSG schemes with different fusions and fusion sizes, we define a set of metrics based on a simple reference protocol, where n-way fusions are applied between identical states recursively in each stage to generate a large GHZ state. An example of this kind of scheme with 3-GHZ seed states and three stages of 3-way fusions is shown in fig. 6.1 in ZX-diagram language. The notation introduced in Chapter 1 is used, where square nodes correspond to states and

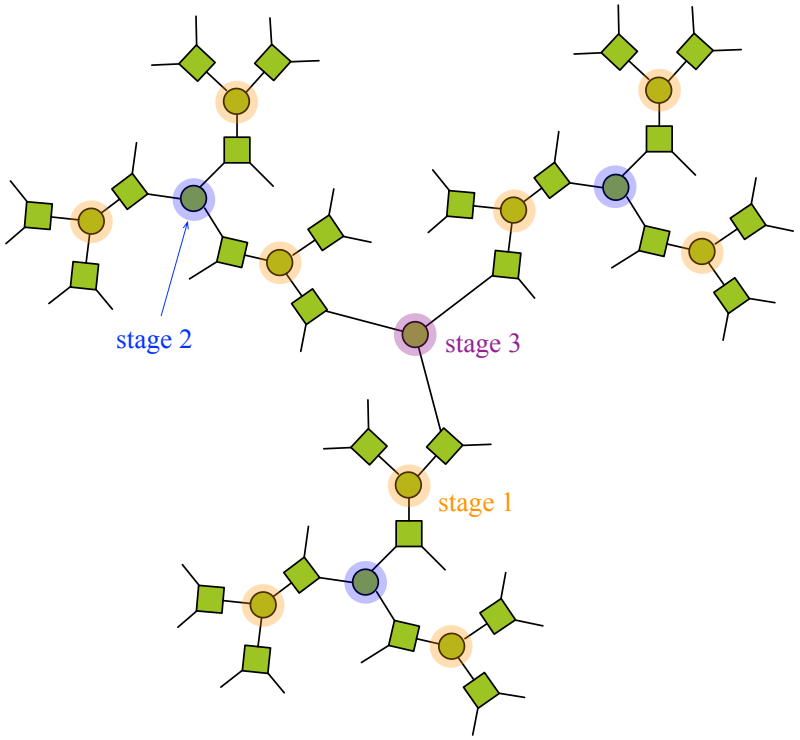

Figure 6.1: Example of the optimal RSG scheme used to assess the impact of different multi-way fusions on the resource state footprint. This scheme uses 3-GHZ seed states and recursively performs 3-GHZ fusions between identical states in three stages, highlighted in blue, orange and purple in order. All seed states and fusions are green spiders.

round ones to fusions.

This reference model is particularly convenient to obtain an insight into how fusions of different sizes for RSG change the high-level properties of RSG schemes for a few reasons. First, by assuming that the state is always generated by combining seed state and fusion spiders of the same colour, effects linked to the specific internal structure of the state can be avoided. The resulting large GHZ state is the most natural to make using the majority of the fusions described so far, which perform an n-GHZ projection upon success. Moreover, the recursive nature of the scheme makes it easy to generalise and to evaluate its properties as parameters such as the size of the target state or fusions change. This will become clear in the description of how the footprint and Pauli error rate are calculated below. The goal of the modelling described in the following sections is to assess the potential usefulness of different fusions for RSG and not to evaluate the performance of specific schemes.

6.1.1 Asymptotic overhead

Our analysis is restricted to fully-destructive n-GHZ fusion circuits, and intentionally omits non-GHZ projections, partially-destructive fusions, e.g., T1 fusions, and boosted fusion circuits. This choice has been made both for simplicity and also because some of these options are expected to have worse performance than the circuits considered here in the presence of noise. In the following we make the assumption that only these kinds of fusions can be used and that the same fusion is used throughout the RSG. In this case, the reference scheme described above minimises the cost of generating a resource state of given size starting from a fixed number of seed states. When using n-way fusions to make an S-qubit target state, this perfectly symmetric scheme uses $\mathcal{O}(\log_n(S))$ fusion stages. At the other end of the spectrum, the most asymmetric construction in which one seed state at a time is added to the resource state requires $\mathcal{O}(S)$ fusion stages. From this it can be inferred that any intermediate schemes which fuse states of different sizes have more fusion stages than the optimal logarithmic scheme.

In Chapter 2, a simple formula was given (eq. 2.26) which can be used recursively to calculate the total cost of a specific RSG consisting of a sequence of fusions on a given set of starting seed states. We repeat it here for reference:

$$C = \frac{\gamma}{p}\sum_i C_i, \tag{2.26}$$

where C_i are the costs of the inputs to an entanglement generation stage (SSG or fusion), $\gamma \geq 1$ is the multiplexing inefficiency factor, and p is the success probability of the fusion. Applying this recursively to each stage in the RSG scheme allows to calculate the resource state cost in terms of number of muxed single photons, which are each assigned a cost of $C_{\text{sp}} = 1$. Erasure, mostly due to fusion failure and photon loss, impacts the output of this calculation by modifying the value of the success probability, p. Accounting for photon loss, for an n-way fusion we define:

$$p_{\text{F}} = p_{\text{succ}} \times \eta^n, \tag{6.1}$$

where p_{succ} is the success probability in the ideal case and η^n is the probability of all n photons being present at the input. Larger fusions are penalised by a decrease in the second factor in this equation, while the first factor is higher in fusion circuits with increased ideal success probability. In the case of Seed State Generation (SSG) circuits, which are not fully destructive, we modify this definition:

$$p_{\text{SSG}} = p_{\text{succ}} \times \eta^{n_{\text{det}}}, \tag{6.2}$$

where n_{det} is the number of photons detected in the circuit upon success.

The size of the state obtained in the reference RSG scheme after each n-way fusion stage, starting with seed states of a qubits, is:

$$\begin{aligned}
&\text{stage 1}: \quad S_1 = n(a-1) \\
&\text{stage 2}: \quad S_2 = n \times [n(a-1)-1] = n^2 a - (n^2+n) \\
&\ldots \\
&\text{stage k}: \quad S_k = n^k a - (n^k + n^{k-1} + \ldots n) = n^k a - n\frac{(n^k-1)}{n-1}.
\end{aligned} \tag{6.3}$$

From this, the number of stages needed to make a state of size S can be calculated:

$$k = \log_n \left(\frac{S - \frac{n}{n-1}}{a - \frac{n}{n-1}} \right). \tag{6.4}$$

It should be noted that not all resource state sizes can be obtained with all seed state and fusion sizes using this type of scheme. Therefore, this expression is an analytic continuation to be used for the purposes of this high level assessment only, and it does not always correspond to an existing RSG scheme. Given that it is derived from a protocol designed to minimise footprint, this should be regarded as a lower bound on the number of fusion stages required.

Each stage s fuses identical states with fusions of the same size n, so its cost can be written as:

$$C_s = nC_{s-1}\frac{\gamma}{p}. \tag{6.5}$$

Applying this recursively to all stages leads to the cost of the full RSG scheme:

$$\begin{aligned} C &= C_{\mathrm{SSG}}\left(n\frac{\gamma}{p}\right)^{k} = \\ &= C_{\mathrm{SSG}}\left(n\frac{\gamma}{p}\right)^{\log_n\left(\frac{S-\frac{n}{n-1}}{a-\frac{n}{n-1}}\right)} = \\ &= C_{\mathrm{SSG}}\left[n^{\log_n\left(n\frac{\gamma}{p}\right)}\right]^{\log_n\left(\frac{S-\frac{n}{n-1}}{a-\frac{n}{n-1}}\right)} = \\ &= C_{\mathrm{SSG}}\left[\frac{S-\frac{n}{n-1}}{a-\frac{n}{n-1}}\right]^{\log_n\left(\frac{\gamma}{p}\right)+1}, \end{aligned} \tag{6.6}$$

where C_{SSG} is the cost of the initial seed state. C provides a lower bound on the overhead of an RSG scheme which uses n-way fusions and a-qubit seed states to produce a resource state of size S, and we thus refer to it as the *asymptotic overhead*.

Fig. 6.2 shows the values of this footprint metric for a resource state of size $S = 64$, generated starting from Bell pairs or 3-GHZ states and using various n-GHZ fusions from Chapter 3. These results include a loss rate of 15% factored into the success probabilities of the SSG circuits and fusions, and a mux inefficiency of $\gamma = 2$, roughly representative of a simple $N \to 1$ mux scheme [94]. It should be noted that these parameter values, used in the rest of this chapter, are chosen for simplicity and illustrative purposes, but they correspond to a conservative model of the system performance, resulting in asymptotic overheads significantly larger than the expected cost of an RSG used in a practical scheme. In particular, high-efficiency multiplexing schemes with $\gamma < 2$ are known, and the loss rate is lower for the photons measured in the SSG than for those measured in fusions, due to the former having a significantly shorter worldline. The size S of the resource state has also been chosen to illustrate the differences between fusions of various sizes, but it is not necessarily representative

of the requirements of a realistic architecture. The difference in the asymptotic overhead can be very large when a different error model and resource state size are considered. For example, the minimum cost in fig. 6.2, obtained using 3-GHZ seed states and the 4-GHZ fusion with success probability $9/32$, is reduced by about three orders of magnitude when targeting a 16-qubit state with the average loss cut by about half and more efficient multiplexing. On the other hand, while the absolute values of C are affected by the choice of these parameters, the trends observed in the plots below and the conclusions extracted from them are expected to hold for a wide range of parameters.

The SSG costs are based on the Bell state generation circuit from Ref. [141], which produces a Bell pair on fixed pairs of rails upon detection of two photons with probability $p_{\text{succ}} = 1/8$, and the 3-GHZ state generator from Ref. [142], which detects three photons and succeeds with probability $p_{\text{succ}} = 1/32$. So the costs are:

$$\begin{aligned} C_{\text{BP}} &= 4 \times \frac{2}{\frac{1}{8} \times 0.85^2} \sim 89 \\ C_{\text{3-GHZ}} &= 6 \times \frac{2}{\frac{1}{32} \times 0.85^3} \sim 625. \end{aligned} \tag{6.7}$$

Ideal success probability and number of inputs are the only properties of a fusion which affect the value of C, so in these plots multiple fusion circuits may be represented by a single data point. For instance, the standard n-GHZ fusions and the new circuits with entangling failure outcomes but same success probability $1/2^{n-1}$ overlap in this plot.

This data clearly suggests that multi-way fusions have the potential to significantly reduce the cost of generating a resource state of fixed size. Considering only fusion circuits with the standard $1/2^{n-1}$ success probability (blue), both plots show that the asymptotic overhead is minimised by fusions with $n > 2$, namely $n = 4$ when starting from Bell pairs and $n = 3$ from 3-GHZ states, and it is higher for both smaller and larger n. The footprint reduction compared to using T2 fusions is $\sim 6.5\times$ (or $\sim 85\%$) for schemes with 3-GHZ seed states. When starting from Bell pairs, two-way fusions cannot

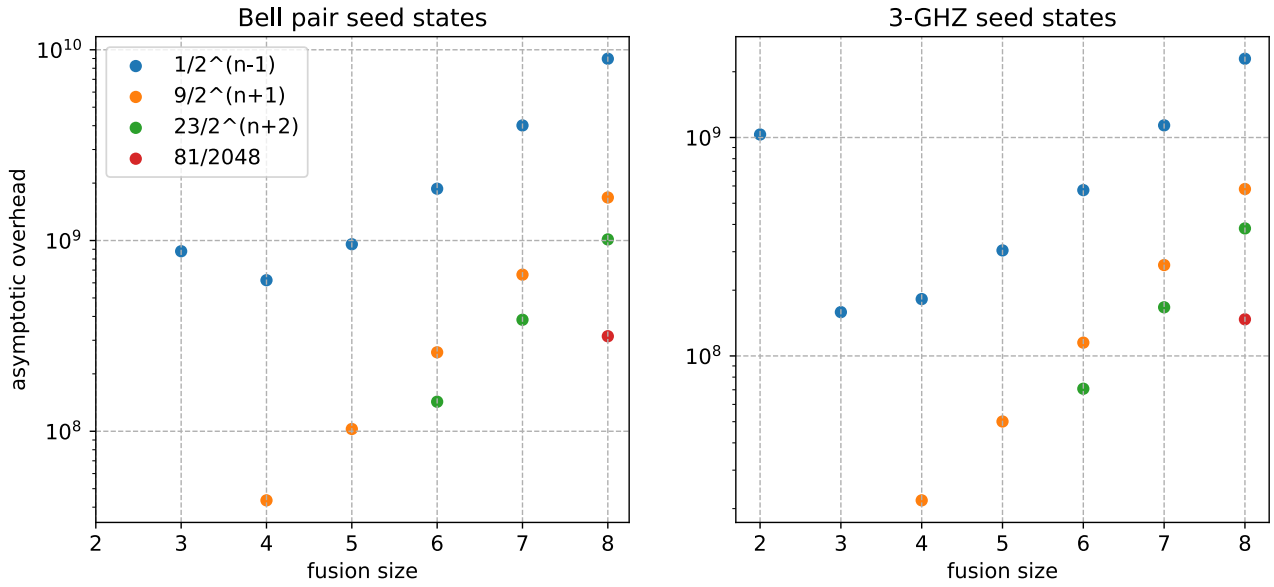

Figure 6.2: Asymptotic overhead for the generation of a 64-qubit resource state using multiway fusions of different sizes and starting from either Bell pairs (left) or 3-GHZ states (right). Fusions are labelled by their success probability. A uniform photon loss rate of 15% for all photons and multiplexing inefficiency $\gamma = 2$ are assumed at each stage. Note that these are conservative assumptions, since photons detected in an SSG circuit have a shorter worldline, and therefore lower loss, than those which reach a fusion, and in addition high-efficiency multiplexing schemes with $\gamma < 2$ are known. The chosen target resource state size should also only be considered as an illustrative example rather than a requirement for a realistic scheme. The asymptotic overhead can vary by several orders of magnitude depending on the details of the error model and chosen resource state size. Under the assumptions used for these plots, the minimum cost is obtained for fusions with $n = 4$, even when considering the standard fusion circuits with probability $1/2^{n-1}$ (blue). The depth-2 circuits with higher success probability introduced in Chapter 3 lead to significant cost reductions.

be used to grow a resource state, so the same comparison cannot be made. However, relative to the scheme using the smallest fusion available in that case, i.e. $n = 3$, the four-way fusion gives a $\sim 1.4\times$ (or $\sim 30\%$) smaller overhead. This confirms that the trade-off between the decrease in the success probability with n and the reduction in the number of fusions and seed states necessary to generate a resource state of fixed size can be favourable for multi-way fusions and lead to footprint gains for certain n.

In addition, the increased success probability of some of the depth-2 fusion circuits introduced in Chapter 3 is seen to translate into very significant decreases in the asymptotic overhead compared to their standard counterparts. In particular, compared to the standard 4-GHZ fusion, the ((1, 3), (2, 4)) circuit (orange, $p_{\text{succ}} = 9/32$) results in a cost reduction of $\sim 14.3\times$ (or $\sim 93\%$) when starting from

Bell pairs and $\sim 8.3\times$ (or $\sim 88\%$) from 3-GHZ states. Due to the underlying RSG scheme, these values reflect the best case scenario with respect to footprint, so the gains obtained in a concrete RSG scheme with a fixed structure and other constraints will be smaller. However, the fact that such big improvements are possible in this model is very promising for more realistic schemes as well.

6.1.2 Adjusted Pauli rate

The total Pauli rate and per-qubit marginal Pauli rate were introduced in Chapter 4 as metrics to quantify the impact of physical noise on the amount of error on the outcomes of a fusion circuit. It was found that both these quantities increase with n for the standard n-GHZ fusions (see figures 5.1 and 5.2). When considering an RSG scheme, the overall probability of Pauli errors on the resource state qubits depends not only on the performance of the individual fusion circuit, but also on the total number of fusions performed and of fused qubits. Given a fixed resource state, these two quantities change with the size of the fusion, so in order to compare the Pauli error rates of RSG schemes which use different fusions, the error metrics need to be adjusted accordingly.

In the reference scheme described above, the total number of fusions performed up to each fusion stage is:

$$\begin{aligned}
&\text{stage 1}: \quad F_1 = 1 \\
&\text{stage 2}: \quad F_2 = n \times 1 + 1 = n + 1 \\
&\text{stage 3}: \quad F_3 = n \times (n+1) + 1 = n^2 + n + 1 \\
&\dots \\
&\text{stage k}: \quad F_k = n^{k-1} + n^{k-2} + \dots 1 = \frac{n^k - 1}{n-1}.
\end{aligned} \tag{6.8}$$

The number of stages needed to generate a state of size S is given in eq. 6.4, which we can use to get

the total number of fusions in the RSG:

$$\begin{aligned} F(S,n,a) &= \frac{n^{\log_n\left(\frac{S-\frac{n}{n-1}}{a-\frac{n}{n-1}}\right)}-1}{n-1} = \\ &= \frac{1}{n-1}\left[\left(\frac{S-\frac{n}{n-1}}{a-\frac{n}{n-1}}\right)-1\right] = \\ &= \frac{S-a}{(n-1)a-n}. \end{aligned} \tag{6.9}$$

Since n-way fusions are used at each stage, the total number of fused qubits is then:

$$Q(S,n,a) = nF(S,n,a) = \frac{n(S-a)}{(n-1)a-n}. \tag{6.10}$$

In fact, these expressions hold independently of the order in which fusions are performed, i.e., not only in the reference scheme in which identical states are fused at each stage. If each fusion stage adds a seed state of size a, after k fusions the size of the state is $S_k = a + k[(n-1)(a-1)]$. Then setting $S_k = S$ and solving for k gives the same result as eq. 6.9.

We are interested in understanding the potential performance improvements given by multi-way fusions with respect to two-way fusions, so we can use eqs. 6.9 and 6.10 to obtain the normalisation factors:

$$f(n,2,a) = \frac{F(n)}{F(2)} = \frac{a-2}{(n-1)a-n} \tag{6.11}$$

$$q(n,2,a) = \frac{Q(n)}{Q(2)} = \frac{n}{2} \times \frac{a-2}{(n-1)a-n}, \tag{6.12}$$

which represent the ratios of the numbers of fusions and qubits compared to a scheme which uses T2 fusions. Both of these are independent of the size of the resource state, S. These ratios are zero for $a = 2$ because Bell pairs cannot be used to generate resource states with two-way fusions, so we consider $a = 3$ in the rest of this section. The total Pauli rate metric captures the overall errors in a fusion, so multiplying it by $f(n,2,3)$ gives a way of rescaling and comparing values for different n to

each other. Similarly, the per-qubit marginal Pauli rate applies to each fused qubit, so the analogous rescaling factor is $q(n,2,3)$. We refer to the values obtained through this renormalisation as the *adjusted total Pauli rate* and *adjusted marginal Pauli rate*.

Fig. 6.3 shows the adjusted total Pauli rates for standard n-GHZ fusions with $2 \leq n \leq 5$, obtained by rescaling the values in fig. 5.1 by $f(n,2,3)$. The trend in this plot matches the expected behaviour described in Chapter 2, with the total amount of error decreasing sharply from $n = 2$ to $n = 3$, reaching a minimum at $n = 4$ and then starting to increase again. The contribution from input distinguishability (green) decreases with n, while errors due to combinations of multiphoton contamination and loss on different qubits (orange) increase. The contribution from Pauli errors already present on the input qubits (pink) and multiphoton + loss errors on the same qubit (blue) is also reduced due to the decrease in the number of fusions as n gets larger. This plot confirms the original intuition, providing evidence that RSG which uses only two-way fusions is far from optimal in terms of error rate on the resulting state, and the use of multi-way fusions can significantly improve it, consequently increasing the threshold of the architecture. Under the simple RSG model considered here, and with this specific choice of parameters, the optimal $n = 4$ gives a $\sim 42\%$ reduction in the adjusted total Pauli error rate compared to $n = 2$.

The adjusted per-qubit marginal Pauli rate is shown in fig. 6.4, obtained by multiplying the values in fig. 5.2 by $q(n,2,3)$. The total error rate and its contributions follow the same trend as in the adjusted total Pauli rate above. The main difference is in the optimal fusion size, $n = 3$ in this case instead of $n = 4$ as found above. This discrepancy can be explained by the presence of correlated Pauli errors in fusions with $n \geq 4$, which increase the marginalised error rate more than the total one, as discussed in Chapter 5. Nevertheless, this data points to the same conclusion as fig. 6.3 that RSG with multi-way fusions can reduce Pauli errors on resource state qubits.

As for the footprint calculations in the previous section, a few caveats apply to these results. They are based on the same simplified RSG scheme, and are therefore not a fully accurate representation

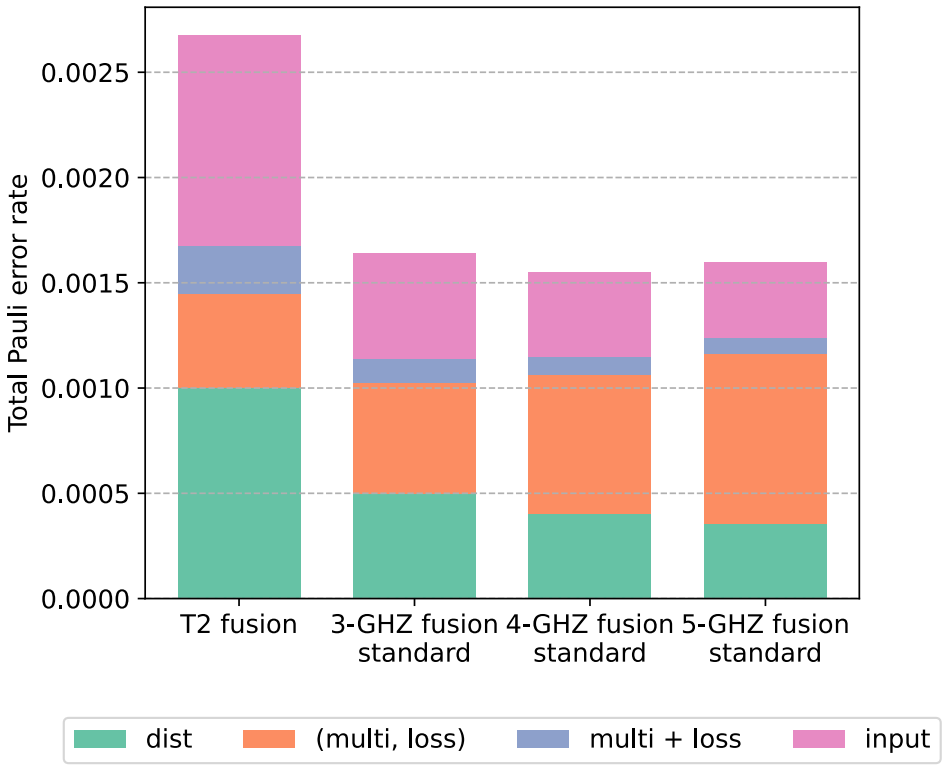

Figure 6.3: Adjusted total Pauli error rate for standard n-GHZ fusions with $2 \leq n \leq 6$. When the total number of fusions in the RSG is taken into account, multi-way fusions are seen to provide a significant advantage in terms of Pauli errors. The optimal size in this specific case is $n = 4$, which reduces the amount of errors by $\sim 37\%$ compared to the T2 fusion.

of the impact of n-way fusions on the error rate of any resource state. The optimal fusion size found here is also dependent on the choice of error parameters. In particular, higher multiphoton and/or loss rates would make the corresponding contribution to Pauli errors increase faster with n, penalising larger fusions more and shifting the optimal size towards lower n. The parameters chosen here are expected to reflect a realistic regime for a physical computing architecture, but it is important to take this variability into account when deciding which fusions to use in an RSG scheme. Overall, however, the conclusion stands that RSG designs which make use of multi-way fusions have the potential to outperform current schemes based on T2 fusions only. In section 6.3 we put this theory to the test by providing a specific RSG scheme which makes use of these larger fusions to achieve better performance than an equivalent scheme to make the same resource state only using two-way fusions. Before that, we review the principles used in the design and optimisation of such a scheme in the next section.

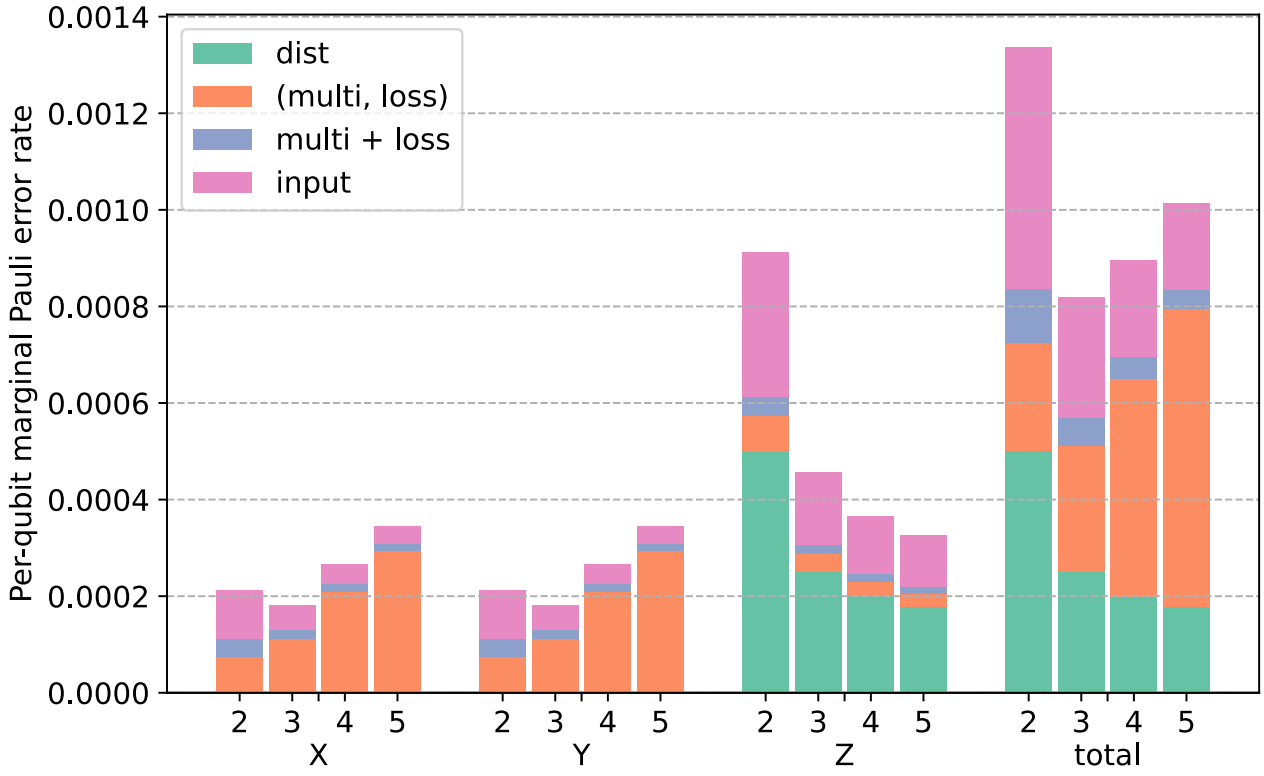

Figure 6.4: Adjusted marginal per-qubit Pauli rate for standard n-GHZ fusions with $2 \leq n \leq 5$.

6.2 Calculating RSG performance

The number of possible RSG schemes for a given resource state is extremely large, especially when the option of using fusions of different sizes is available. The footprint and error rate vary greatly between schemes, so a set of design principles is necessary to guide the exploration of this vast space and identify promising protocols. Given a specific RSG scheme, heuristic methods can be used to calculate its expected cost and error rates and to optimise the design.

6.2.1 Footprint optimisation

As discussed in Chapter 2, any target resource state can be represented by a ZX-diagram and any valid RSG scheme for it is a decomposition of this diagram into seed state and fusion spiders such that each of the former is only connected to the latter and vice versa. Starting from a resource state diagram, an RSG scheme to construct it using T2 fusions can be found by first using the identity rule

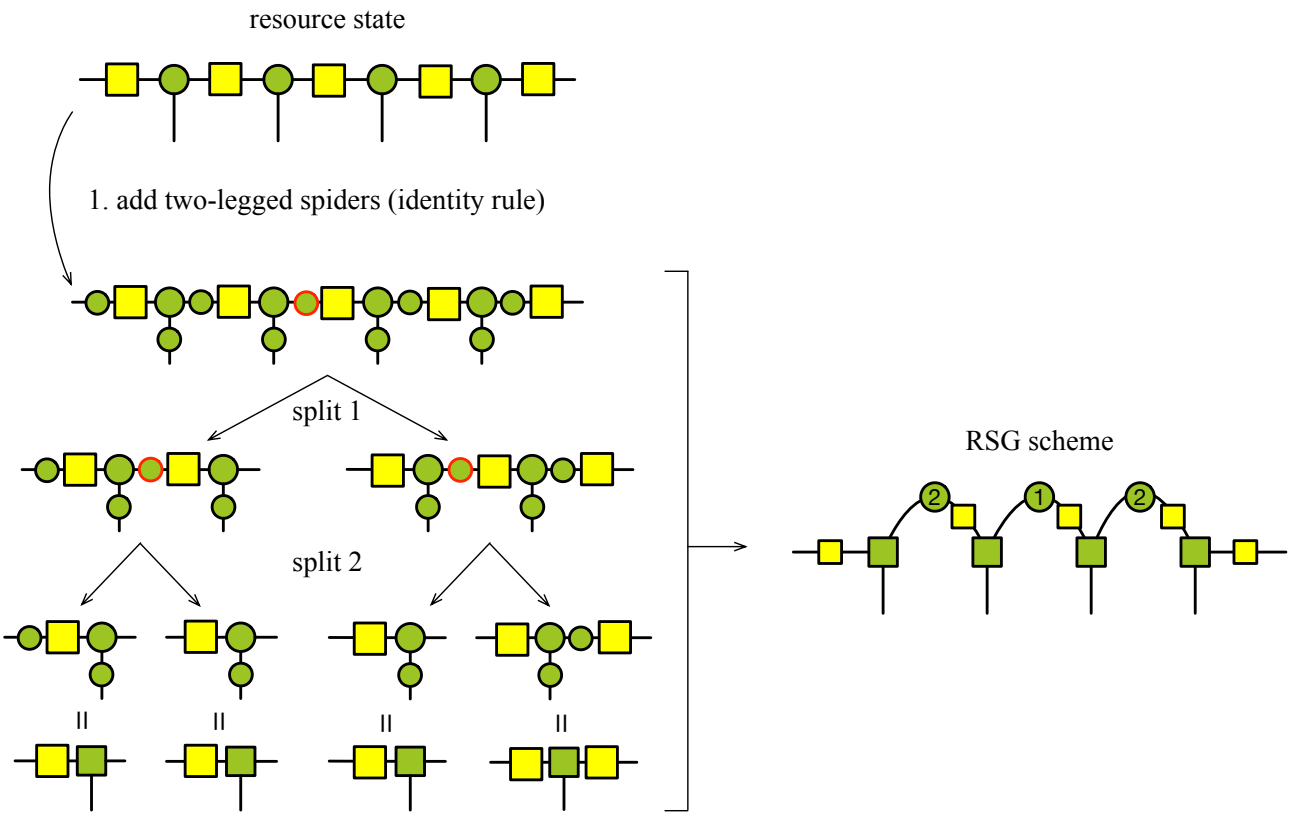

Figure 6.5: Example of the derivation of the T2-based RSG scheme for the six-qubit linear cluster state in fig. 2.9 from the ZX diagram of the state via splitting. First two-legged spiders are added on every leg, then the diagram is recursively split at one of those nodes (indicated by a red outline) until diagrams with a minimum number of legs are left, corresponding to seed states.

of the ZX calculus to insert green nodes on all legs, and then defining a series of recursive splittings at those nodes until a diagram is obtained which consists of disjoint spiders with a small, predetermined number of outer legs, which correspond to seed states. If at any point only spiders with too many ports are left, the inverse of the spider fusion rule can be used to obtain nodes with lower degree and continue splitting the diagram. Once only ZX sub-diagrams with the minimum number of outer legs are left, any remaining identity two-leg spiders can be removed. The RSG scheme then consists of T2 fusions applied in reverse order compared to the splitting, i.e., beginning with fusions between seed states and progressively making larger states. Fig. 6.5 shows an example of how this procedure can be used to obtain the RSG scheme previously introduced in fig. 2.9 for the six-qubit linear cluster state.

This approach can be generalised to construct RSG schemes with multi-way fusions. The size of the fusions is reflected in the choice of which nodes to perform the split at: splitting the diagram

at an n-port spider results in an n-way fusion in the RSG scheme. Using the inverse of the spider fusion rule, any node of degree n in the resource state diagram can be decomposed into a network of connected spiders each with n_i external legs such that $\sum_i n_i = n$. Given a set of fusion and seed state sizes, the first step to obtain an RSG scheme which uses them is to decompose all nodes in the diagram until all spiders have degree corresponding to the chosen fusions and seed states. Then the steps outlined for the two-way fusion case above can be applied, but with the added freedom to choose larger nodes corresponding to the desired multi-way fusions for the splitting. Fig. 6.6 illustrates this with the example of an RSG scheme for a 6-GHZ state which uses 3-GHZ and 4-GHZ fusions and Bell pairs as seed states.

This method can be turned into a systematic search for RSG schemes and gives a way to minimise the footprint of the resource state. Starting with the ZX-diagram representation of the state and iterating over all possible splittings at each step, it is in principle possible to find all the ways of constructing the diagram using the desired fusions and seed states. The resource overhead can be calculated for every scheme using eq. 2.26 with the appropriate parameters corresponding to the fusion and SSG circuits used, and the smallest one can be identified. In practice, this brute force approach to the search is computationally very expensive and its cost increases rapidly with the number of nodes in the starting ZX-diagram, making it impractical for resource states with tens of qubits as typically used in FBQC architectures. To circumvent this problem, the search can be supplemented with a heuristic derived from the structure of the footprint-optimal RSG described in sec. 6.1, which fuses states of identical size at each step. Accordingly, rather than letting the search explore all possible graph splittings, it can be restricted to only allow a split into sub-diagrams with similar numbers of legs by specifying a minimum size ratio. This drastically reduces the complexity of the computation and allows to search for and optimise RSG schemes of interesting size. This footprint optimisation procedure is used to obtain the RSG schemes described in sec. 6.3 below.

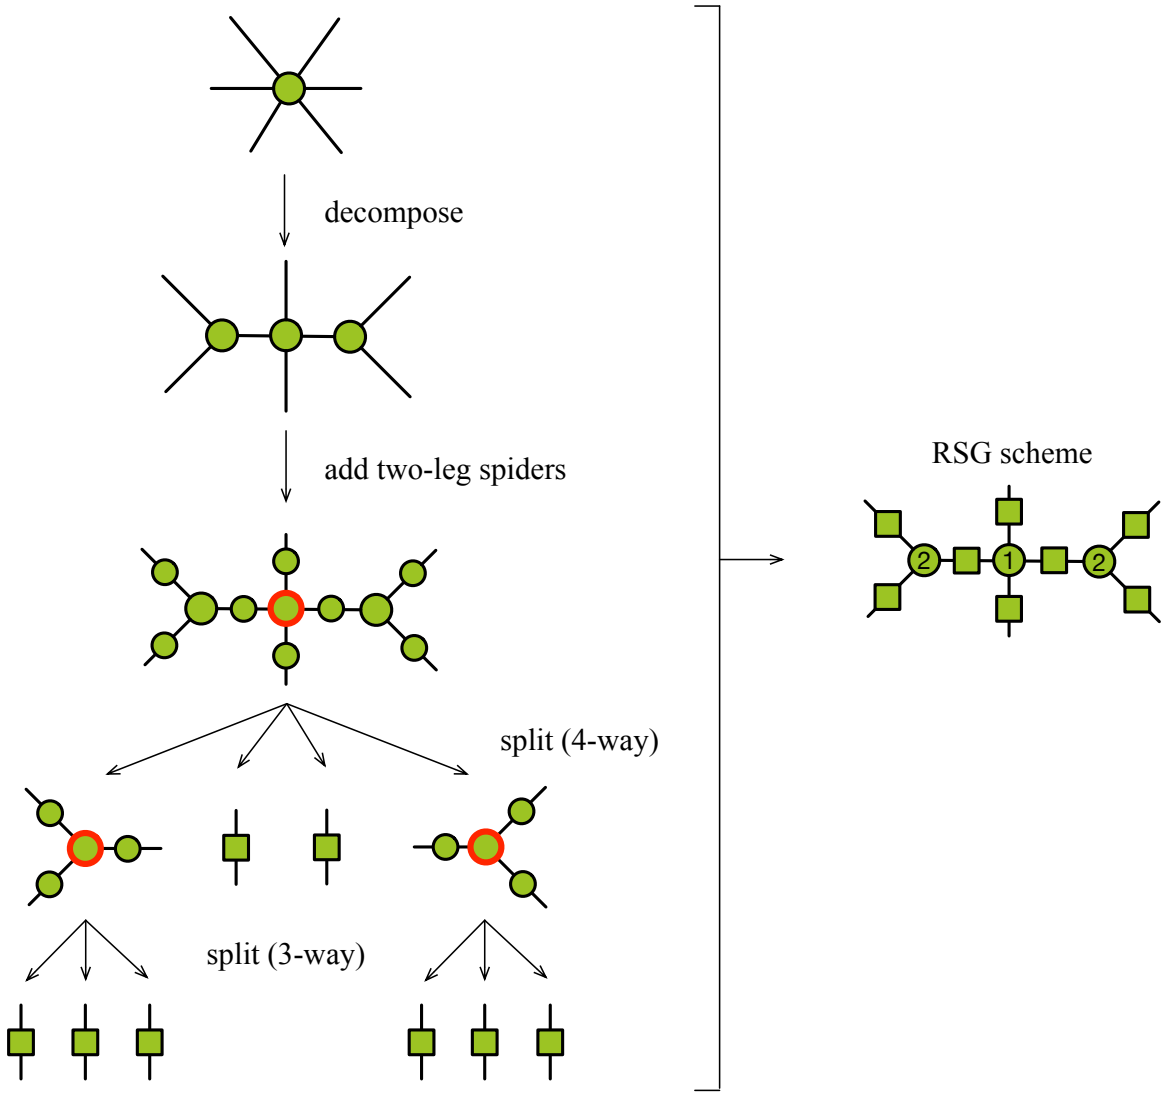

Figure 6.6: Example of the derivation of an RSG scheme for a 6-GHZ state using 3-GHZ and 4-GHZ fusions and Bell pair seed states. The first step consists in decomposing the 6-leg spider into a network of 3- and 4-leg spiders to match the degree of the desired multi-way fusions. The rest of the derivation then proceeds as described above for T2 fusions.

6.2.2 Encoded erasure

As discussed in Chapter 2, local encoding of resource state qubits has been found to be a powerful technique to boost the threshold of FBQC architectures. Given this, we are mostly interested in designing schemes to generate encoded resource states. In particular, we focus here on states whose qubits are replaced with generalised Shor codes (or quantum parity codes), defined in Chapter 2, which offer excellent protection against erasure. Here we describe a heuristic calculation which allows to estimate the rate at which erasure on the local Pauli operators of physical qubits of the code turns

into erasure on its logical operators. This model, also used in Refs. [43, 144], can be seen as describing the behaviour of the local encoding when single qubit measurement in the same basis, X or Z, are performed on all the physical qubits in the encoding, or alternatively, when fusions are applied with a fixed failure basis. Only first-order physical erasure events are accounted for and the erasure rate is assumed to be the same for X and Z and across physical qubits, not capturing any variations due to different SSGs or measurement choices.

Under the assumption that multiplexing removes erasure from RSG fusions, the ability of an encoding to shield the logical operators depends entirely on the structure of the code and it is not influenced by the way in which it is constructed. An (n, m)-Shor code can be represented as a tree ZX-diagram with alternating layers of green and red spiders, as shown in fig. 2.10. The rate at which the encoded qubit suffers a logical erasure can be calculated by propagating the physical erasure through the tree. Consider a spider in the tree with $n+1$ legs. If the value of an operator (X or Z) on one of the bottom n legs is lost with probability p_e, the probability of the same operator being erased on the top leg, P_e, depends on the colour of the spider. A green spider has n ZZ stabilizer operators with support on the top and one of the bottom legs, so the top Z value is only erased if all bottom legs suffer an erasure. On the other hand, there is only one $X^{\otimes(n+1)}$ stabilizer operator, so erasure of any of the bottom X operators turns into an erasure on the top leg. This can be written as:

$$P_e^Z = (p_e^Z)^n, \tag{6.13}$$

$$P_e^X = 1 - (1 - p_e^X)^{n-1}. \tag{6.14}$$

For a red spider, the opposite applies:

$$P_e^X = (p_e^X)^n, \tag{6.15}$$

$$P_e^Z = 1 - (1 - p_e^Z)^{n-1}, \tag{6.16}$$

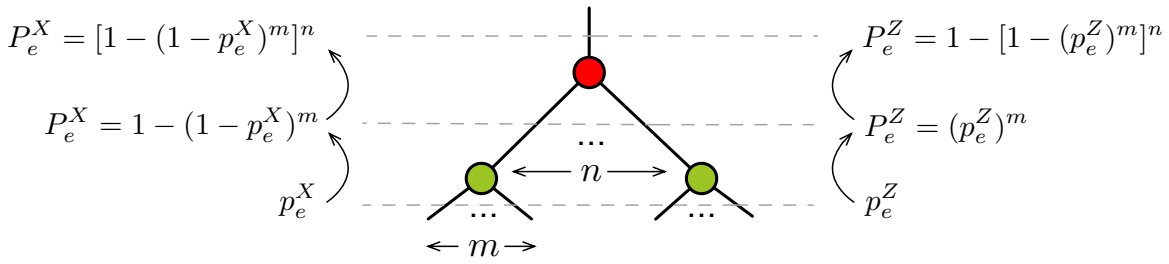

Figure 6.7: First-order calculation of the rate of erasure of the Z and X operators of the encoded qubit by propagating erasure at the physical level through the diagram of the encoding.

or, equivalently, X and Z are swapped across a Hadamard node:

$$P_e^Z = p_e^X, \tag{6.17}$$

$$P_e^X = p_e^Z. \tag{6.18}$$

The rate of erasure on the encoded qubit can be calculated to first order starting from the physical qubits at the bottom of the tree and recursively applying these equations to obtain P_e^Z and P_e^X on the root of the tree. This is shown in fig. 6.7 for the (n,m)-Shor encoded qubit, which in this case we define up to Hadamard on the top leg, which only has the effect of swapping the two logical operators. Consider the example of a (2, 2)-Shor encoding. In this case:

$$P_e^{\overline{Z}} = 2(p_e^Z)^2 - (p_e^Z)^4, \tag{6.19}$$

$$P_e^{\overline{X}} = 4(p_e^X)^2 - 4(p_e^X)^3 + (p_e^X)^4, \tag{6.20}$$

so if $p_e^X = p_e^Z = p_e$, $\overline{Z}$ is better protected than $\overline{X}$. When assessing the performance of a resource state with respect to erasure, it is useful to consider the average logical erasure, $P_e = \frac{P_e^{\overline{X}} + P_e^{\overline{Z}}}{2}$. If $p_e = 15\%$, the loss rate assumed in the error calculations so far, the logical erasure probabilities are significantly suppressed, $P_e^X = 7.7\%$, $P_e^Z = 4.4\%$ and $P_e = 6.1\%$.

While this calculation is based on a simplified model and can therefore only give an approximate value of the expected erasure rate, it can provide useful insight for the comparison of different local

encodings to apply to the resource state.

6.2.3 Encoded Pauli errors

Similarly to erasure, physical Pauli errors in the RSG can also be mapped to flips on the whole encoded qubit at the root of the tree, which need to be corrected by the QEC code implemented by the fusion network. The fundamental difference is that multiplexing does not typically allow to suppress these unheralded errors at individual fusions within the RSG, and as a result the structure of the scheme plays an important role in determining the Pauli error rate at the level of the encoded qubit. While in the case of erasure it is mainly the physical errors on the outer qubits which contribute to the overall error rate on the encoded qubit, Pauli errors also result from every fusion in the RSG. In addition, the primary purpose of local encodings is to reduce the impact of erasure at the physical qubit level on the fusions performed in the fusion network, so they do not necessarily lead to a decrease in the Pauli error rate. In fact, the increasing number of qubits and fusions used to construct the encoding can have the opposite effect of increasing the total amount of Pauli error.

An accurate calculation of the impact of physical Pauli errors in the generation of an encoded qubit requires taking into account the structure of the local encoding as well as the choices of measurements performed on the physical qubits, which determine the probability of reconstructing different representatives of the logical operators, as well as the presence of checks to protect against unheralded Pauli errors. However, a simpler heuristic can be used to obtain an estimate based exclusively on the structure of the tree and its generation scheme.

The probability of the logical operator of an encoded qubit suffering a Pauli error can be estimated to first order by summing the Pauli rates on the fusions along the path between the outer qubits and the root of the tree. Due to the redundancy in the definition of logical operators in a generalised Shor code, the probability of an operator on a spider leg in the tree being in the support of a logical operator can decrease moving away from the root towards the leaves. Ignoring any effect due to

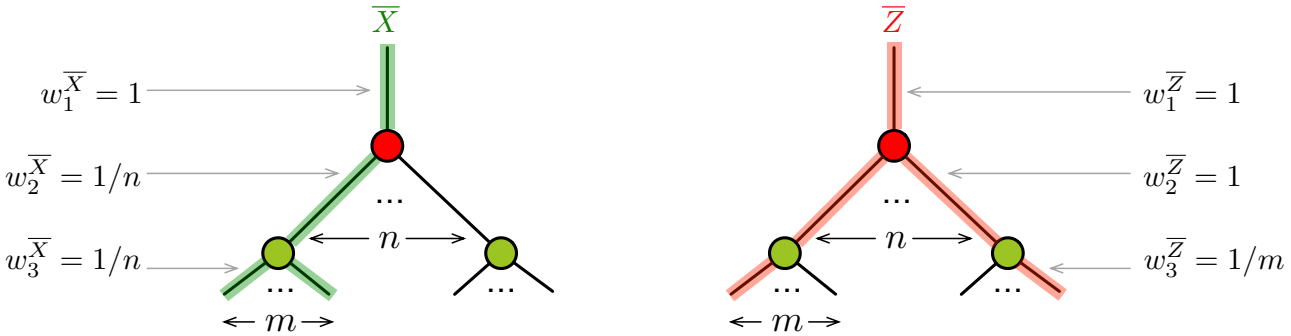

Figure 6.8: Weights associated with the contribution of Pauli operators at different levels of the tree to the corresponding logical operator. Moving away from the root (level 1) towards the leaves (level 3), the weight of Z operators is reduced crossing each green spider and the X and Z operators are swapped across a Hadamard node.

the choice of measurement of the physical qubits in the fusion network, we equate this reduction in probability with the branching ratio of the tree. This can be interpreted as considering the scenario in which only a single representative of each logical operator can be reconstructed, which results in a conservative estimate of the errors. In fact, when multiple representatives of the logical operators are obtained, their values can be compared to perform checks and possibly herald the occurrence of an error, effectively lowering the logical Pauli rate by converting unheralded errors into erasure.

From the Pauli webs for the logical operators of the (n,m) Shor code in fig. 2.11, it can be seen that due to green (red) spiders having multiple ZZ (XX) stabilizers with support on the top port, the probability of a Z (X) operator on one of them being part of the support of the logical operator is reduced. On the other hand, the probability stays the same for the X (Z) operator since there is only one X-type (Z-type) stabilizer operator. So, for a green spider with b lower ports, the weight, i.e., relative importance, of the contribution of the Z operator on each port to a logical operator L is $w_i^L = w_{i-1}^L/b$, where w_{i-1}^L is the weight associated with the upper leg. For the X operator, $w_i^L = w_{i-1}^L$. The inverse is true for a red spider. The root of the tree is assigned $w_1^L = 1$ for both X and Z, since an error on it is guaranteed to flip the logical operator, and all the other weights are calculated relative to it. The weights at various levels of the tree for the (n,m) Shor local encoding are shown in fig. 6.8.

The Pauli error rate on both logical operators can then be obtained as the weighted sum of the Pauli

errors on the inputs of all fusions between the outer qubits and the root. Consider the RSG scheme for a (2, 2) Shor encoded qubit in fig. 6.9. Using the definition of the support of the Pauli operators and the weights associated with each level in fig. 6.8, the total rate of a flip of the two logical operators can be calculated as:

$$\begin{aligned} P_{\overline{Z}} &= \sum_i n_i \times w_i^{\overline{Z}} \times p_i^Z = \\ &= 1 \times 1 \times p_{1,\mathrm{t}}^Z + 2 \times 1 \times p_{1,\mathrm{b}}^Z + 2 \times 1 \times p_{2,\mathrm{t}}^Z + 4 \times \frac{1}{2} \times p_{2,\mathrm{b}}^Z = \\ &= p_{1,\mathrm{t}}^Z + 2p_{1,\mathrm{b}}^Z + 2p_{2,\mathrm{t}}^Z + 2p_{2,\mathrm{b}}^Z, \end{aligned} \tag{6.21}$$

and

$$\begin{aligned} P_{\overline{X}} &= \sum_i n_i \times w_i^{\overline{X}} \times p_i^X = \\ &= 1 \times 1 \times p_{1,\mathrm{t}}^X + 2 \times \frac{1}{2} \times p_{1,\mathrm{b}}^X + 2 \times \frac{1}{2} \times p_{2,\mathrm{t}}^X + 4 \times \frac{1}{2} \times p_{2,\mathrm{b}}^X = \\ &= p_{1,\mathrm{t}}^X + p_{1,\mathrm{b}}^X + p_{2,\mathrm{t}}^X + 2p_{2,\mathrm{b}}^X, \end{aligned} \tag{6.22}$$

where the i are the labels for the fusion ports (t for top and b for bottom), n_i is the number of equivalent ports in the RSG scheme, w_i^L is the weight of the port towards the logical operator L, and p_i^L is the probability of the local Pauli operator in the support of L being flipped. As seen in Chapters 4 and 5, the Pauli error rates can vary significantly between different fusions, especially for different sizes, which significantly influences the values of $P_{\overline{X}}$ and $P_{\overline{Z}}$ and therefore the choice of RSG design. As for erasure, the two Pauli error rates can be combined to obtain an average Pauli rate $P_P = \frac{P_{\overline{Z}} + P_{\overline{X}}}{2}$ as a single metric to quantitatively compare the performance of different local encodings and generation schemes.

As mentioned above, this is expected to be a conservative estimate of the Pauli error rate on the encoded qubit. However, this heuristic is useful for our purpose of comparing the Pauli error rates of different RSG schemes to one another without needing to perform a significantly more involved

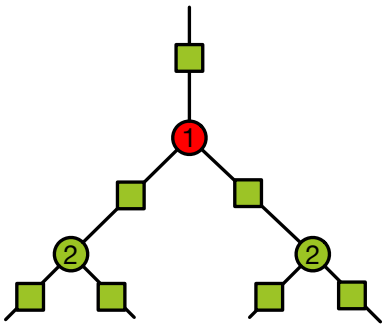

Figure 6.9: Example RSG scheme for the (2, 2) Shor local encoding with Bell pairs as seed states and two stages of 3-GHZ fusions.

calculation based on other aspects of the architecture.

6.3 RSG schemes with multi-way fusions

The analysis in sec. 6.1 suggests that multi-way fusions, in particular 3- and 4-GHZ fusions, have the potential to improve both the cost of generating a resource state, as well as the Pauli error rate on its qubits. The simple model which leads to this conclusion does not capture many of the details which can impact the performance of a realistic RSG scheme, so this should only be considered as a positive indication.

In order to turn this into more solid evidence of the usefulness of multi-way fusions in RSG, we now consider the generation of two specific resource states introduced in the original FBQC proposal [43]: the four-star and six-ring states, and in particular their version with qubits encoded in a (2, 2) Shor code. Two RSG schemes are proposed and analysed for each state, one which only uses T2 fusions and one which includes n-GHZ fusions. The methods described in secton 6.2 are applied to design footprint-optimal RSG schemes and to evaluate the total Pauli errors, and the costs and error rates are compared across the two protocols to generate the same states. Erasure is not factored into this comparison because all states have the same local encoding, and we make the simplified assumption that the scheme used to generate it does not affect the erasure rate. The values of $P_e^{\overline{X}}$ and $P_e^{\overline{Z}}$ for the (2, 2) Shor code considered here are calculated in sec. 6.2.2.

6.3.1 Primitives

The first step in the design of the RSG schemes to be compared is the computational search and footprint optimisation described in sec. 6.2.1. This requires the definition of a set of available primitives for RSG, i.e., SSG and fusion circuit options. These have different success and Pauli error probabilities associated with them, which impact the resource overhead and Pauli error rates of the different schemes. Here we give an overview of the primitives which appear in the RSG schemes described and compared in the rest of this section.

Seed states

The RSG schemes presented below use Bell pairs and 3-GHZ seed states. Both are generated using heralded circuits, the former using four single photons and succeeding with probability $1/8$ [141], and the latter taking in 6 single photons and succeeding with probability $1/32$ [142]. These are the same SSGs considered in sec. 6.1.1, where the cost of each state was calculated as (eq. 6.7):

$$\begin{aligned} C_{\text{BP}} &= 4 \times \frac{2}{\frac{1}{8} \times 0.85^2} \sim 89 \\ C_{\text{3-GHZ}} &= 6 \times \frac{2}{\frac{1}{32} \times 0.85^3} \sim 625. \end{aligned} \tag{6.7}$$

The error probabilities used to calculate the encoded Pauli error rate of the RSG scheme are a combination of various contributions, including Paulis present on the qubits output by the SSG circuits. In this analysis the simplified assumption is made that the Bell state and 3-GHZ state generators produce states with the same per-qubit Pauli error rates, corresponding to the "input" errors in Chapter 4,

Seed state colour	p_i^Z	p_i^X
Green	0.02%	0.04%
Red	0.04%	0.02%

Table 6.1: Probabilities of a local Pauli operator flip for qubits of seed states corresponding to ZX-spiders of different colours.

eq. 4.41:

$$\begin{aligned} p_X &= 0.01\%, \\ p_Y &= 0.01\%, \\ p_Z &= 0.03\%. \end{aligned} \tag{6.23}$$

These values are expected to be representative of SSG circuits which correspond to green ZX-spiders. RSG schemes can contain red seed state spiders too, which are equivalent to green ones with Hadamards on all their legs by the duality rule of the ZX-calculus. The error probabilities for such states are the ones in eq. 6.23 conjugated by a Hadamard, i.e., with p_X and p_Z exchanged. The corresponding probabilities of a local Pauli operator flip, p_i^Z and p_i^X, used to calculate the encoded Pauli error rates in eqs. 6.21 and 6.22, can be obtained by combining the marginal error rates:

$$p_i^Z = p_X + p_Y, \quad p_i^X = p_Z + p_Y. \tag{6.24}$$

Table 6.1 gives the values of these probabilities for the qubits of seed state spiders of different colours.

Fusions

The footprint analysis summarized in fig. 6.2 suggests that some of the depth-2 n-GHZ fusions can lead to significant overhead reductions compared to standard n-GHZ circuits, particularly the ((1, 3), (2, 4)) 4-GHZ fusion with success probability $9/32$. On the other hand, T2 fusions allow more flexibility in the design of specific RSG schemes where finite-size effects play an important role.

From the Pauli error analysis in figs. 6.3 and 6.4 it can be seen that $n = 3$ and $n = 4$ minimise the

impact of Pauli errors on the resource state. While those plots only contain data for standard n-GHZ fusion circuits, the comparison of different 4-GHZ fusion circuits in fig. 5.3 shows a modest increase in total Pauli errors for the ((1, 3), (2, 4)) circuit compared to the standard one, which does not alter the picture significantly. The large footprint savings this fusion allows are very likely to outweigh the higher Pauli error rate.

Overall, these observations suggest that the fusions to be included in the RSG design toolbox are the standard 3-GHZ fusion and the ((1, 3), (2, 4)) 4-GHZ fusion, in addition to the T2 fusion. These are the fusion circuits used in the design and optimisation of the two RSG schemes with multi-way fusions described below.

The marginal per-qubit Pauli error rates for these fusion circuits were obtained in Chapter 5. Those values consist of different contributions which represent Pauli errors affecting the qubit at different points along its worldline. Some are due to physical imperfections on the qubits causing a Pauli flip within the fusion circuit, such as photon distinguishability ("dist") and combinations of multiphoton contamination and loss on different qubits ("(multi, loss)"). Others are caused by these two noise mechanisms acting on the same qubit ("multi + loss"), which can be assumed to mostly happen before it reaches the fusion circuit. Finally, there are Pauli errors which are already present on the qubits output by the SSG circuits ("input"), which we associate with the seed states themselves here, as described in the previous section. The data obtained for these errors assumes fusions which can be represented as green ZX-spiders. A general RSG diagram may contain red fusion spiders, or spiders with Hadamards acting on one or more legs. In this case, like for seed states, the effect of the Hadamard is to swap the p_X and p_Z error rates. However, this only applies to errors which occur within a fusion as a result of physical imperfections on the input qubits. The Paulis which are already present on the inputs, on the other hand, are not affected. Therefore, the Pauli error rates are different between plain inputs of green fusion spiders and inputs with a Hadamard node, which include those of red fusion spiders. Table 6.2 shows these marginal Pauli error probabilities for both cases for all

Fusion circuit	p_{succ}	Input Hadamard	p_i^Z	p_i^X
T2	1/2	No	0.043%	0.113%
		Yes	0.093%	0.062%
3-GHZ, standard	1/4	No	0.073%	0.128%
		Yes	0.108%	0.093%
4-GHZ, ((1, 3), (2, 4))	9/32	No	0.179%	0.161%
		Yes	0.141%	0.199%

Table 6.2: Ideal success probabilities and per-qubit marginal Pauli error rates of the three types of fusions used in the RSG schemes described below. Different error rates are associated with fusion inputs which have a Hadamard applied on them or not.

three fusion circuits used in the RSG schemes described below. The total error rates associated with each leg between a seed state and a fusion node in the RSG diagram are calculated by summing the appropriate combinations of these values and those in table 6.1 to reflect the structure of the scheme.

It is assumed that these probabilities are the same for all inputs of a fusion. In reality, there will be differences due to the input qubits originating from different SSG circuits or to the stage at which the fusion is performed, which determines the size of the multiplexing network seen by a qubit. However, these effects are not easily incorporated in a modular, leading-order analysis of an RSG scheme, and they are not expected to significantly alter the conclusions of the comparison. Therefore we choose to disregard them in the context of this analysis. Another simplification we make is to assume a constant mux inefficiency parameter $\gamma = 2$, independent of the fusion stage and size. This is chosen as representative of the expected performance of the simplest $N \to 1$ mux scheme [94], but in practice both of those aspects would influence the effective value at different locations, and lower inefficiency could be achieved with more advanced multiplexing strategies. Therefore, the resource overheads reported below are likely overestimates.

6.3.2 Four-star resource state

The four-star resource state was part of one of the first FBQC architectures in Ref. [43], and also of subsequent proposals [117]. Schemes based on it have been found to have lower photon loss thresholds than those which use the six-ring resource state, described later, but at the same time, this

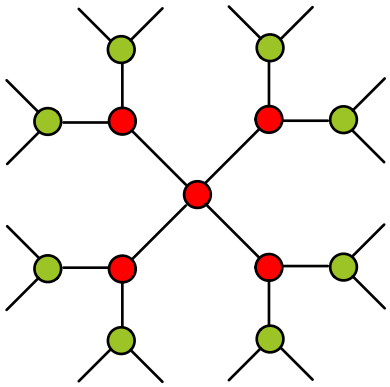

Figure 6.10: ZX-diagram of the 16-qubit four-star resource state with a (2, 2) Shor encoding on each qubit.

state is significantly cheaper to generate, as we will see below. In fact, the four-star is just a 4-GHZ state, which could be simply constructed by performing a T2 fusion between two 3-GHZ seed states, for example. However, we are interested in the version of this state with a (2, 2) Shor encoding applied to its qubits, which is a more complex state of 16 physical qubits. The corresponding ZX-diagram is shown in fig. 6.10.

In the following, we present two footprint-optimised RSG schemes for this encoded four-star state, one using only T2 fusions and the other one including 3- and 4-GHZ fusions as well.

T2 fusions only

Fig. 6.11 shows the RSG scheme given by the cost optimisation routine when only T2 fusions are allowed. Square nodes represent seed states and circular ones are fusions, where the number indicates the stage in which the fusion is performed. The fusion ordering proceeds from the highest to the lowest number, i.e., fusion 1 is the last one to be applied. The central 4-GHZ state (blue shading), which we refer to as the *kernel*, is constructed by fusing two 3-GHZ seed states, and then T2 fusions are used to attach the local encodings to each of the four qubits (grey shading). The total number of multiplexed single photons used by this RSG scheme with a loss rate of 15%, i.e., $\eta = 0.85$, is:

$$C_{\text{4-star, T2}} = 16,958,097 \sim 1.7 \times 10^7 \text{ muxed single photons.} \tag{6.25}$$

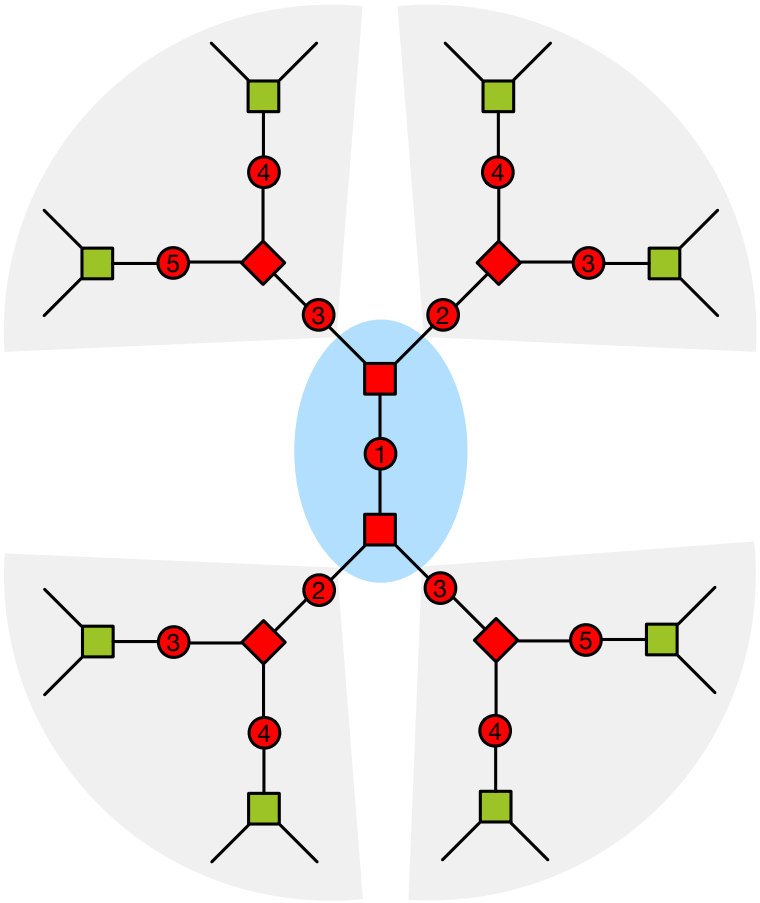

Figure 6.11: RSG scheme for the (2, 2) Shor-encoded four-star resource state using only T2 fusions. Square nodes represent seed states and circular nodes are fusions, where the labels indicate the (reverse) ordering in which they are performed. The grey shading indicates each of the four encoded qubits and the blue one the kernel of the state.

The average Pauli error rate on each encoded qubit is calculated using the method described in sec. 6.2.3. Errors on each fusion in the RSG are mapped to the root of the closest local encoding, i.e., the edge where the kernel and tree shadings meet in fig. 6.11. All four encoded Pauli error rates are then summed and the total is divided by the number of encoded qubits. In this case, due to the symmetry of the scheme, all qubits have the same Pauli error rate, but in general this might not be the case, and this process of averaging gives a way of comparing different schemes.

Errors from fusions in the kernel and trees can be mapped to encoded Pauli errors separately. The

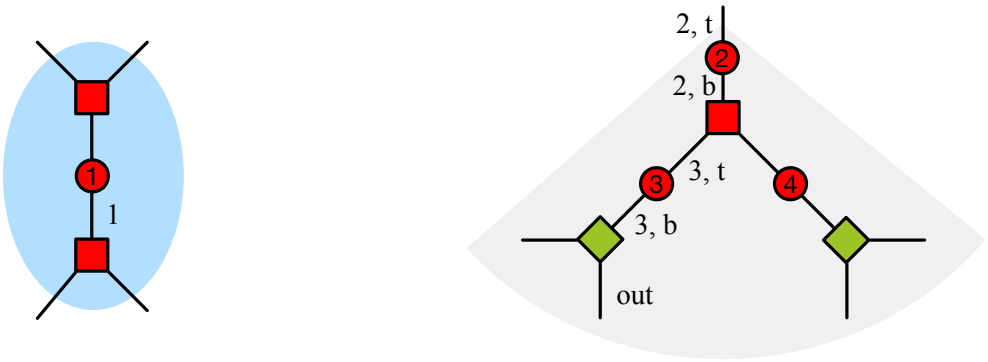

Figure 6.12: Labelling of the edges of the T2-based four-star RSG diagram used to map errors to the encoded qubit level. The calculation of the total encoded Pauli error rate is split into kernel errors (left) and tree errors (right).

diagrams in fig. 6.12 show the labelling of the legs used in the following calculations. We have:

$$P_{\overline{Z}}^{\text{kernel}} = 2 \times 2 \times p_1^Z, \tag{6.26}$$

where one factor of 2 reflects the symmetry of the kernel on the two sides of fusion 1, and the other indicates that a flip of the Z operator on that leg is mapped to flips on two encoded qubits through the 3-GHZ state node. Similarly:

$$P_{\overline{X}}^{\text{kernel}} = 2 \times p_1^X, \tag{6.27}$$

where this time $p_{1,A}^X$ is not multiplied by 2 because a flip of the X operator on the input of fusion 1 only propagates to one encoded qubit. The encoded Pauli error rates from the (2, 2) Shor codes are:

$$P_{\overline{Z}}^{\text{encoding}} = 4 \times \left(p_{2,t}^Z + p_{2,b}^Z + 2 \times p_{3,t}^Z + 2 \times p_{3,b}^Z + 4 \times \frac{1}{2} \times p_{\text{out}}^Z \right), \tag{6.28}$$

$$P_{\overline{X}}^{\text{encoding}} = 4 \times \left(p_{2,t}^X + p_{2,b}^X + 2 \times \frac{1}{2} \times p_{3,t}^X + 2 \times \frac{1}{2} \times p_{3,b}^X + 4 \times \frac{1}{2} \times p_{\text{out}}^X \right) \tag{6.29}$$

where p_{out} indicate error rates on seed state qubits from table 6.1, and the factor of 4 outside the brackets captures the symmetry across all four encoded qubits.

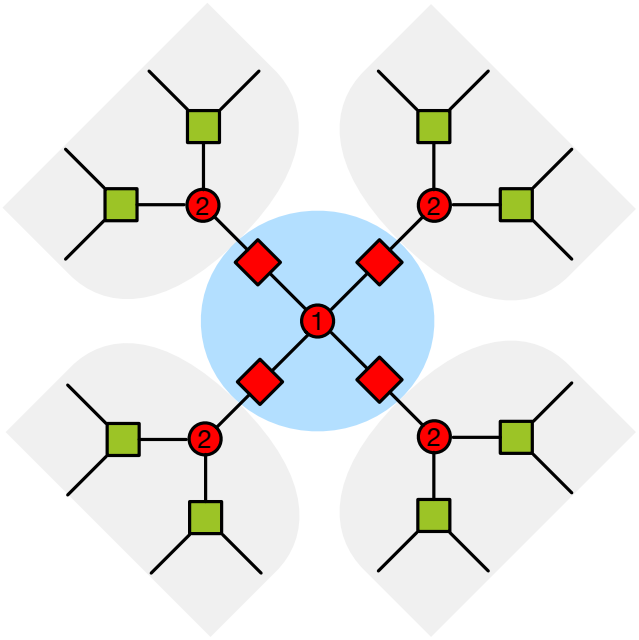

Figure 6.13: RSG scheme for the 24-qubit (2, 2) Shor-encoded four-star resource state using multi-way fusions. Both the standard 3-GHZ fusion circuit and the ((1, 3), (2, 4)) 4-GHZ fusion are used.

The per-qubit encoded Pauli error rates are then:

$$P_{\overline{Z}} = \frac{1}{4} \times \left(P_{\overline{Z}}^{\text{kernel}} + P_{\overline{Z}}^{\text{encoding}} \right) = 0.79\%, \tag{6.30}$$

$$P_{\overline{X}} = \frac{1}{4} \times \left(P_{\overline{X}}^{\text{kernel}} + P_{\overline{X}}^{\text{encoding}} \right) = 0.29\%, \tag{6.31}$$

so the average probability is $P_P = 0.54\%$.

Multi-way fusions

The lowest footprint RSG scheme with multi-way fusions is shown in fig. 6.13. The kernel in this case consists of a single 4-GHZ fusion, and each tree is constructed in a single 3-GHZ fusion step. This is a clear example of how larger fusions reduce the number of seed states and fusion stages required to generate a resource state, compared to the T2-based scheme in fig. 6.11. This results in a very significant reduction in the resource overhead compared to the previous scheme:

$$C_{\text{4-star, MW}} = 950,564 \sim 9.5 \times 10^5 \text{ muxed single photons}, \tag{6.32}$$

which is $\sim 18\times$ smaller than $C_{\text{4-star, T2}}$, or a $\sim 94\%$ saving.

The calculation of the average encoded Pauli error rate proceeds as for the previous example, with the labels used in this case shown in fig. 6.14. In this scheme, each input of the central four-way fusion contributes separately to one encoded qubit and all four local encodings are produced in the same way. Therefore, the average Pauli error rate is the same as the error rate for a single encoding, which we calculated below. We have:

$$P_{\overline{Z}}^{\text{kernel}} = p_1^Z, \tag{6.33}$$

$$P_{\overline{X}}^{\text{kernel}} = p_1^X, \tag{6.34}$$

and

$$P_{\overline{Z}}^{\text{encoding}} = p_{2,t}^Z + 2 \times p_{2,b}^Z + 4 \times \frac{1}{2} \times p_{\text{out}}^Z, \tag{6.35}$$

$$P_{\overline{X}}^{\text{encoding}} = p_{2,t}^X + 2 \times \frac{1}{2} p_{2,b}^X + 4 \times \frac{1}{2} \times p_{\text{out}}^X, \tag{6.36}$$

where we have used $\tilde{p}$ to indicate error rates which only include either the fusion or seed state contributions depending on the spider the corresponding leg belongs to. Using the values from tables 6.1 and 6.2, we then get the following per-qubit error rates:

$$P_{\overline{Z}} = \frac{1}{4}\left(P_{\overline{Z}}^{\text{kernel}} + P_{\overline{Z}}^{\text{encoding}}\right) = 0.54\%, \tag{6.37}$$

$$P_{\overline{X}} = \frac{1}{4}\left(P_{\overline{X}}^{\text{kernel}} + P_{\overline{X}}^{\text{encoding}}\right) = 0.42\%, \tag{6.38}$$

and the average is $P_P = 0.48\%$. This is a relative reduction in the average encoded Pauli error rate of $\sim 10\%$ compared to the T2-based RSG scheme. The Z flip rate is reduced by $\sim 31\%$, but the X flip rate increases by $\sim 46\%$.

So, we find that multi-way fusions allow to dramatically reduce the cost of the (2, 2) Shor encoded

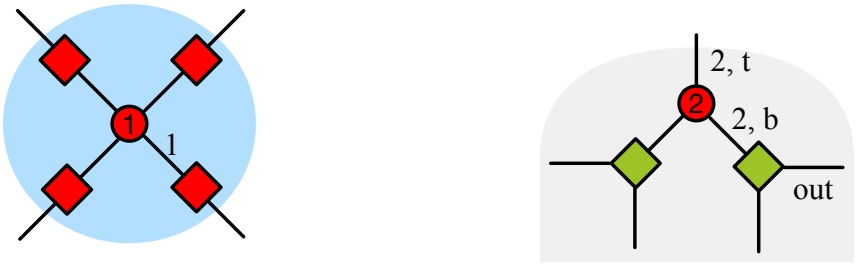

Figure 6.14: Labelling of kernel and tree RSG diagram edges used to map errors to the encoded qubit level in the multi-way four-star RSG scheme.

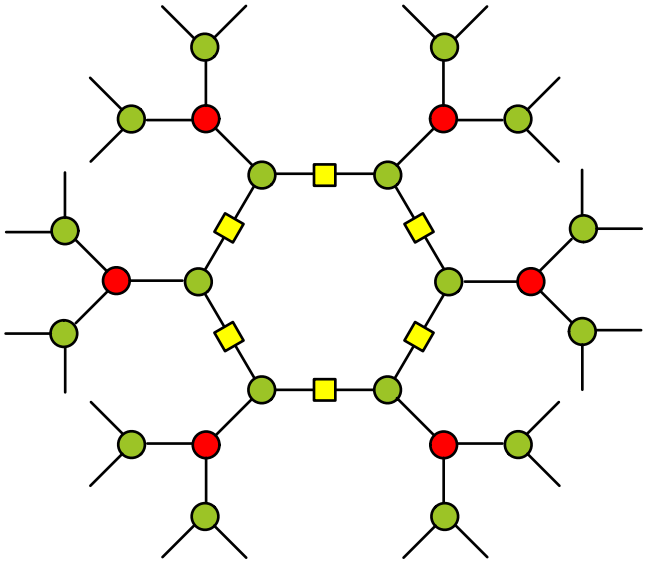

Figure 6.15: ZX-diagram representation of the six-ring resource state with a (2, 2) Shor encoding on each of its qubits.

four-star resource state, while simultaneously decreasing the average rate of Pauli errors on the encoded qubits. It is important to point out, however, that the rate of one type of Pauli flip increases, the logical X flip in this case. In order to assess the impact of this on the performance of the architecture, it is necessary to evaluate this scheme in the context of the full fusion network, which is outside of the scope of this analysis.

6.3.3 Six-ring resource state

The FBQC architecture with the best performance in Ref. [43] has a six-qubit ring resource state with (2, 2) Shor encoding on its qubits. The ZX-diagram of this state is shown in fig. 6.15. Schemes based on this resource state with different local encodings have been studied [112, 115] and found to have good photon loss tolerance properties compared to the original architecture. One key difference with

the four-star state considered in the previous section is that this state contains a loop in its diagram, which requires two qubits from the same input state to be fused to each other in one step to be closed. This makes the state more expensive than one without loops, because the fusion of a state with itself removes the advantage of multiplexing independent inputs. The fusion which closes the loop and the previous one which creates the input state can be seen as happening in parallel since there is no benefit in multiplexing between them, resulting in a single stage with lower success probability $p_1 p_2$, where p_1 and p_2 are the success probabilities of the individual fusions. The generation of this kind of state from linear clusters produced by quantum emitters has been studied and optimised [180].

As in the previous section, we present two RSG schemes for the six-ring resource state, one with T2 fusions only and the other also including multi-way fusions, and a comparison of their costs and encoded Pauli error rates is given.

T2 fusions only

The RSG scheme in fig. 6.16 generated the (2, 2) Shor encoded six-ring resource state using only T2 fusions and 3-GHZ seed states. Every tree encoding is generated in the same way, by fusing three 3-GHZ states together in two steps, and then it is attached to a 3-GHZ in the kernel through a T2 fusion. The fusions in the kernel join pairs of encoded qubits and then close the loop. The cost of this protocol can be calculated as before, but here the last two fusion stages are merged into a single one with success probability $\frac{1}{4}\eta^4$. The result is:

$$C_{\text{6-ring, T2}} = 638,159,943 \sim 6.4 \times 10^8 \text{ muxed single photons.} \tag{6.39}$$

This is over an order of magnitude more expensive than the T2-based scheme for the four-star state ($\sim 38\times$), with a per-encoded qubit cost which is $25\times$ higher. This illustrates the fact that generating tree-like states is significantly easier than making resource states which contain loops.

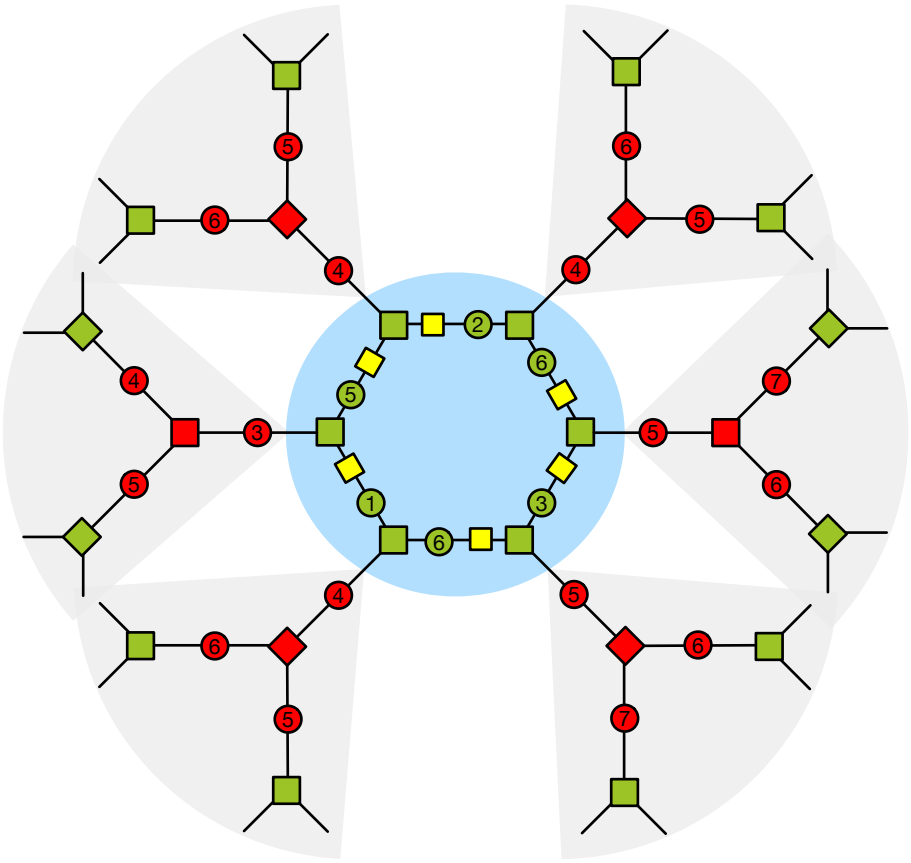

Figure 6.16: RSG scheme to generate the (2, 2) Shor encoded six ring resource state using only T2 fusions and 3-GHZ seed states.

As before, we can calculate the encoded Pauli error rate by splitting the RSG scheme into the kernel and tree parts, shown in fig. 6.17. The structure of the former is such that every T2 fusion has one leg with a Hadamard applied to it and one without. What this means is that every Pauli flip on an edge on the loop can be mapped to a flip on a single encoded qubit, i.e., on only one of the outer legs. A flip of the Z operator on an internal leg of a 3-GHZ spider is equivalent to a flip on its outer leg, whereas a flip of X corresponds to flips on both the outer leg and the other inner one. By pushing through the T2 fusion spiders and using the Hadamard to turn a flip of X into a flip of Z, it is then possible to turn all errors into single flips the Z operator on an outer leg. An example of this is shown in fig. 6.18 for one of the fusions in the kernel RSG diagram.

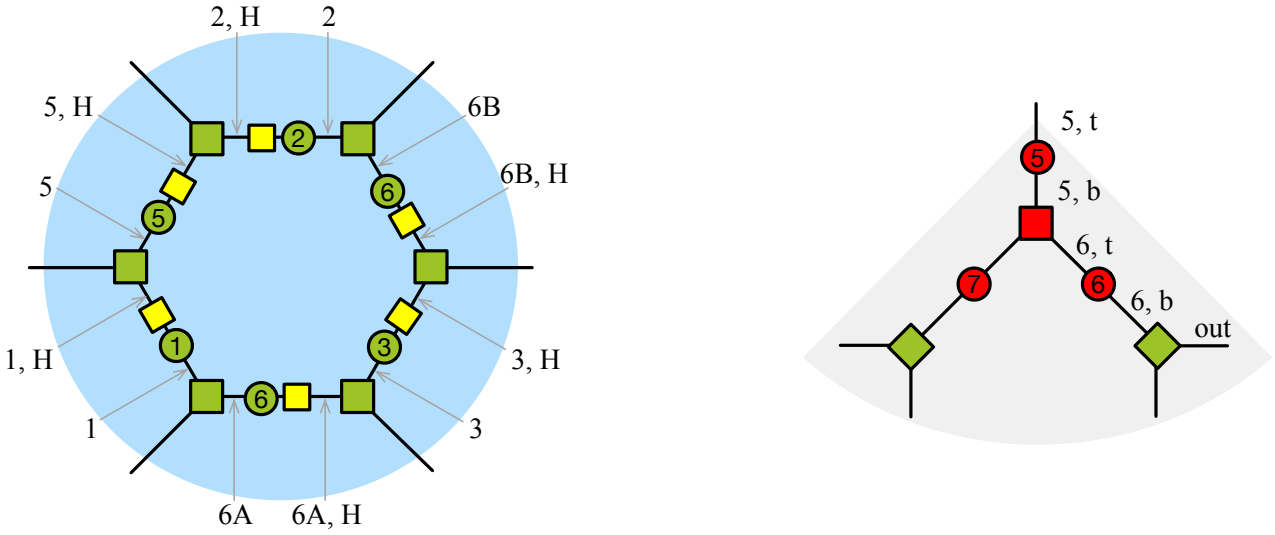

Figure 6.17: Labels associated with each edge in the RSG ZX-diagram for the calculation of the encoded Pauli error rate. Hadamards are considered to be part of the fusion circuits, so the Pauli errors are associated with edges between the Hadamard nodes and the seed state spiders.

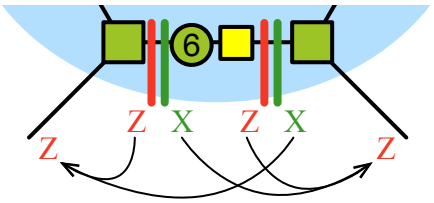

Figure 6.18: Example of the mapping of X and Z flip operators on edges of the loop in the six-qubit RSG scheme to Z operator flips on a single outer edge.

Using this mapping, the total rate of Z operator flips is:

$$\begin{aligned} P_{\overline{Z}}^{\text{kernel}} = {} & p_{6A,H}^{X} + p_{6A}^{Z} + p_{1}^{Z} + p_{1,H}^{X} + \\ & p_{1}^{X} + p_{1,H}^{Z} + p_{5}^{Z} + p_{5,H}^{X} + \\ & p_{5}^{X} + p_{5,H}^{Z} + p_{2,H}^{Z} + p_{2}^{X} + \\ & p_{2,H}^{X} + p_{2}^{Z} + p_{6B}^{Z} + p_{6B,H}^{X} + \\ & p_{6B}^{X} + p_{6B,H}^{Z} + p_{3,H}^{Z} + p_{3}^{X} + \\ & p_{3,H}^{X} + p_{3}^{Z} + p_{6A,H}^{Z} + p_{6A}^{X}, \end{aligned} \tag{6.40}$$

where each line corresponds to one of the outer legs, starting from the bottom left in fig. 6.17 (left) and proceeding clockwise. Since all internal flips are turned into Z operator flips on the outer qubits,

the rate of X flips due to kernel fusions is $P_{\overline{X}}^{\text{kernel}} = 0$. The corresponding flip rates resulting from fusions within the encoding are:

$$P_{\overline{Z}}^{\text{encoding}} = p_{5,t}^{Z} + p_{5,b}^{Z} + 2 \times p_{6,t}^{Z} + 2 \times p_{6,b}^{Z} + 4 \times \frac{1}{2} \times p_{\text{out}}^{Z}, \tag{6.41}$$

$$P_{\overline{X}}^{\text{encoding}} = p_{5,t}^{X} + p_{5,b}^{X} + 2 \times \frac{1}{2} p_{6,t}^{X} + 2 \times \frac{1}{2} p_{6,b}^{X} + 4 \times \frac{1}{2} \times p_{\text{out}}^{X}. \tag{6.42}$$

The per-encoded qubit Pauli error rates are then:

$$P_{\overline{Z}} = \frac{1}{6}\left(P_{\overline{Z}}^{\text{kernel}} + P_{\overline{Z}}^{\text{encoding}}\right) = 0.96\% \tag{6.43}$$

$$P_{\overline{X}} = \frac{1}{6}\left(P_{\overline{X}}^{\text{kernel}} + P_{\overline{X}}^{\text{encoding}}\right) = 0.29\%, \tag{6.44}$$

so the average is $P_P = 0.63\%$.

Multi-way fusions

With the addition of 3-GHZ fusions, the RSG scheme in fig. 6.19 is obtained for the (2, 2) Shor encoded six-ring. Four-way fusions are not useful in this case because the ZX-diagram of the state, fig. 6.15, does not contain any four-valent nodes. In addition to 3-GHZ seed states and T2 fusions, the ring now also contains three-way fusion spiders connected to one half of a Bell pair on the outside. Each encoded qubit is made in a single three-way fusion step, used to connect two 3-GHZs to either another 3-GHZ or to a Bell pair node. This scheme only requires one fusion per encoding and four for the ring, with a total of only four stages, leading to a much smaller resource overhead:

$$C_{\text{6-ring,MW}} = 52,957,620 \sim 5.3 \times 10^{7} \text{ muxed single photons}, \tag{6.45}$$

a 12×, or $\sim 92\%$, reduction compared to the T2-based RSG scheme, a slightly smaller but similar improvement to that obtained for the four-star resource state.

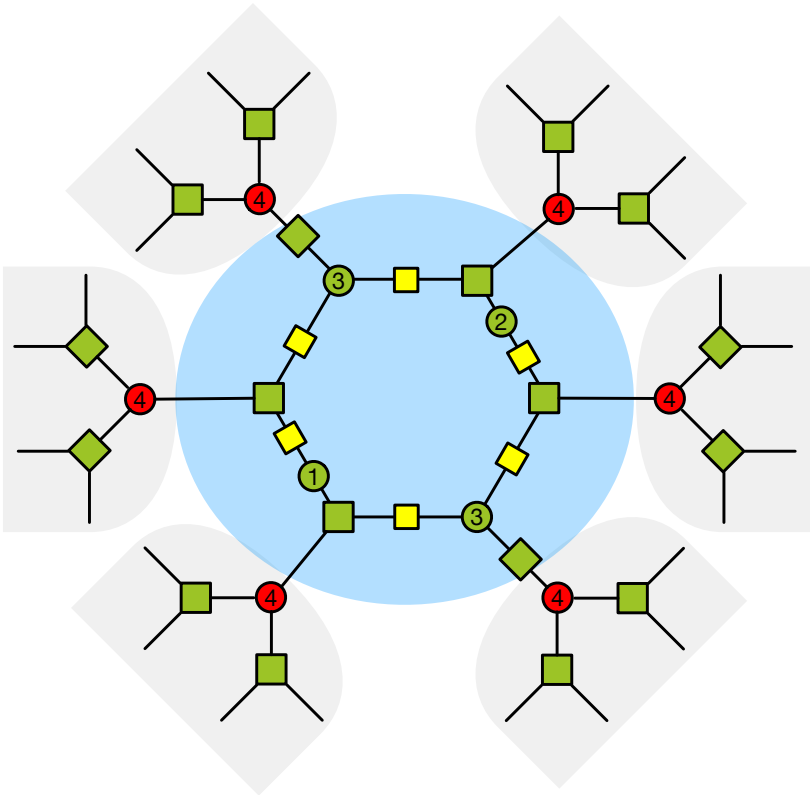

Figure 6.19: RSG scheme for the (2, 2) Shor encoded six-ring resource state with T2 and 3-GHZ fusions.

The labels associated with each edge in the ZX-diagram for the calculation of Pauli errors are shown in fig. 6.20, split into kernel and encoding as before. Similarly to the kernel diagram in fig. 6.17 (left), all spiders with outer legs in this kernel are connected through an edge with a Hadamard node, in some cases with the addition of a two-port green spider. This structure means that every flip on an internal edge can be mapped to a Z flip on a single outer leg here too, using the same reasoning as illustrated in fig. 6.18. Doing this, we calculate the total rate of Z Pauli flips on the outer ports of the

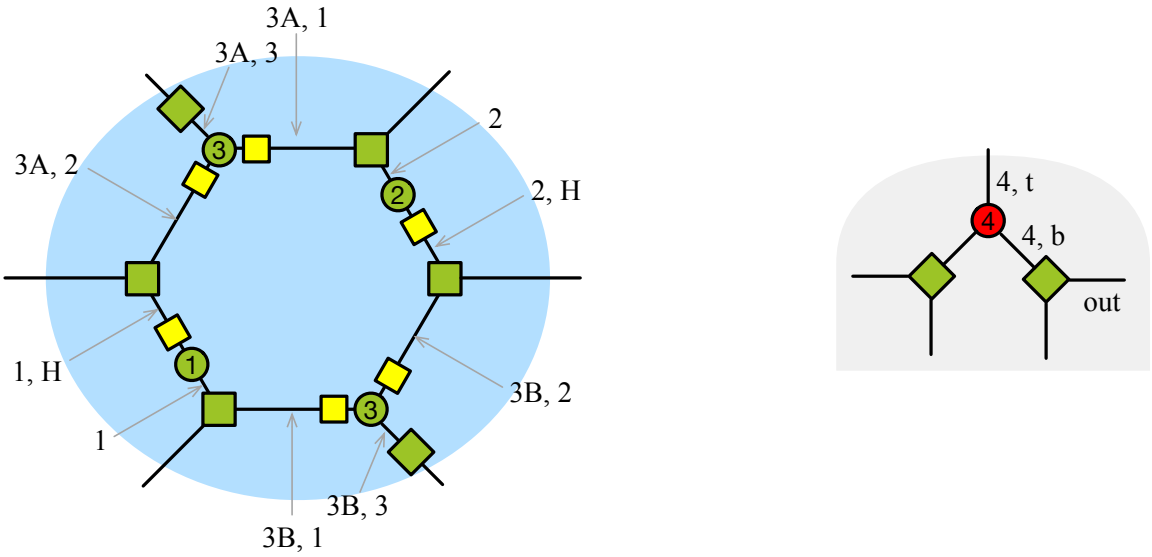

Figure 6.20: Labels used for each edge in the three components of the RSG ZX-diagram for the six-ring resource state.

kernel as follows:

$$\begin{aligned} P_{\overline{Z}}^{\text{kernel}} = {} & p_{3B,1}^{Z} + p_{1}^{Z} + p_{1,H}^{X} + \\ & p_{1}^{X} + p_{1,H}^{Z} + p_{3A,2}^{Z} + \\ & p_{3A,2}^{X} + p_{3A,1}^{X} + p_{3A,3}^{Z} + \\ & p_{3A,1}^{Z} + p_{2}^{Z} + p_{2,H}^{X} + \\ & p_{2}^{X} + p_{2,H}^{Z} + p_{3B,2}^{Z} + \\ & p_{3B,2}^{X} + p_{3B,1}^{X} + p_{3B,3}^{Z}, \end{aligned} \tag{6.46}$$

where each line corresponds to one of the outer ports, starting from the bottom left and proceeding clockwise. Only the two "3" fusions contribute to the X flip rate:

$$P_{\overline{X}}^{\text{kernel}} = p_{3A,3}^{X} + p_{3B,3}^{X}. \tag{6.47}$$

For the encoding we have:

$$P_{\overline{Z}}^{\text{encoding}} = p_{4,t}^{Z} + 2 \times p_{4,b}^{Z} + 4 \times \frac{1}{2} p_{\text{out}}^{Z}, \tag{6.48}$$

$$P_{\overline{X}}^{\text{encoding}} = p_{4,t}^{X} + 2 \times \frac{1}{2} p_{4,b}^{X} + 4 \times \frac{1}{2} p_{\text{out}}^{X}. \tag{6.49}$$

The per-qubit Pauli flip probabilities are then:

$$P_{\overline{Z}} = \frac{1}{6}\left(P_{\overline{Z}}^{\text{kernel}} + P_{\overline{Z}}^{\text{encoding}}\right) = 0.62\%, \tag{6.50}$$

$$P_{\overline{X}} = \frac{1}{6}\left(P_{\overline{X}}^{\text{kernel}} + P_{\overline{X}}^{\text{encoding}}\right) = 0.29\%, \tag{6.51}$$

so an average of $P_P = 0.45\%$. This is a relative $\sim 28\%$ reduction compared to the RSG scheme with only T2 fusions, broken down into a $\sim 35\%$ lowering of Z flips and no change in the X flip rate. So, we see that the use of multi-way fusions in the generation of the (2, 2) Shor encoded six-ring is strictly better than T2 fusions only, with significant improvements to both the cost and the encoded Pauli error rate of the resource state.

The schemes described above provide concrete evidence of the potential of multi-way fusion circuits to improve the resource overhead and Pauli error rates due to the generation of resource states used in existing FBQC architectures. Substantial cost reductions of over an order of magnitude were found in both examples, even larger than expected from the asymptotic overhead calculations in sec. 6.1.1. The difference highlights the importance of the exact structure of the target resource state and finite-size effects in the total cost of the generation protocol, as well as the power of using multiple types of fusion circuits to mirror the required connectivity of the graph state. Mixed results were obtained with respect to Pauli errors, which were found to increase slightly ($\sim 8\%$) on average when multi-way fusions were used to generate the four-star resource state, while they decreased more significantly ($\sim 21\%$) in the case of the six-ring. For both states, the probability of one type of Pauli increased while the other was lowered. These observations indicate that in general the response of the Pauli

error rate to the inclusion of multi-way fusions in the RSG scheme is highly dependent on its specific structure, making it difficult to predict their impact without performing the calculation explicitly. Overall, however, the indication from these two examples is that larger fusions open new paths to improvements in the performance of RSG schemes and therefore of the FBQC architectures they are part of.

It is important to point out that these results are based on simplified seed state and fusion error models and on heuristic calculations of the RSG properties which neglect many of the details of the operations and the overall scheme. An accurate evaluation of the impact of multi-way fusions on the generation of these, and any, states requires threshold simulations of the full FBQC architecture, which take into account not only the properties of the resource state, but also the structure of the fusion network it is used in and the interaction with the QEC code implemented by it. More detailed error models for both SSG and fusion operations are necessary too, to include the effect of many additional physical error mechanisms and correlations in the distribution of errors across different qubits. These kinds of simulations are significantly more complex to perform than the type of analysis presented so far, making them unsuitable for rapid RSG design iteration. In this work we chose to focus on the development and execution of the simpler type of analysis described above as a tool to aid the search for promising RSG schemes to be fully evaluated as part of an architecture.

Further improvements can be made to increase the accuracy of this heuristic analysis framework and to bring its results closer to the actual performance of the RSG scheme. The cost calculation can be refined to account for differences in the value of the inefficiency parameter γ of a mux stage based on the size of the fusion it provides inputs to, and more efficient multiplexing networks could be considered to improve the resource overheads. The value of loss, or efficiency η, can also be made to increase as fusion stages progress to account for the need of a larger mux network before later fusions. The encoded Pauli error rate calculation can also be improved, particularly by removing the assumption of single-qubit measurements being performed on every physical qubit in the encoding.

A different distribution of measurement patterns on the output qubits can be used, determined by the encoded fusion strategy, and check operators of the local code can be accounted for, resulting in a less conservative and more accurate estimate of Pauli errors. This would be reflected in different weights associated with fusion inputs in the calculation. Currently, all inputs to a fusion are also considered to be identical, with the error models being averaged across them. It is possible to incorporate differences between different qubits in the analysis simply by avoiding the averaging step, which could allow to exploit the structure of the errors to orient the fusions in such a way that the overall error rates are reduced.

These are just a few examples of the ways in which the analysis presented in this chapter can be extended and improved to enable a more accurate but still relatively lightweight assessment of the properties of various RSG schemes to facilitate their design. We leave the implementation of these and other upgrades to future work.

6.4 Conclusion

In this chapter, we finally answered the question which served as initial motivation for the research described in this thesis.

First, the general usefulness of multi-way fusions in the context of linear optical RSG was considered by using a simple model of a generation scheme to estimate the impact of larger n-GHZ fusions on both the footprint and the Pauli error rate of the target resource state. The results of this abstract analysis were encouraging, showing that in a realistic noise regime 3- and 4-GHZ fusions could lead to very significant reduction in both cost, up to $\sim 14\times$, and Pauli errors, $\sim 40\%$.

In order to confirm whether these positive results would carry over to the generation of a useful resource state with a more complex structure, we developed a framework to calculate the performance of a specific RSG scheme with respect to three metrics: footprint, encoded qubit erasure, and encoded

qubit Pauli errors. A cost minimisation routine based on the footprint calculation was also described. We then applied this analysis to two example resource states from the original FBQC proposal [43]: the four-star and the six-ring states. The results clearly confirmed our hypothesis that RSG schemes which use multi-way fusions can significantly outperform T2 based protocols: we found cost savings of $\sim 18\times$ and $\sim 12\times$, and average Pauli reductions of $\sim 10\%$ and $\sim 28\%$ for the two states respectively. From this we conclude multi-way fusions are a fundamental tool in RSG design with great potential to result in LO-FBQC architectures with higher thresholds and lower resource requirements than current proposals.

Chapter 7

Computation with Multi-Way Fusions

This chapter considers the use of multi-way fusions between resource states in the fusion network. In particular, the impact of photon loss is considered. This is typically modelled as leading to erasure of all fusion outcome information. Here, a model of photon loss only causing partial erasure is introduced and the question of whether multi-way fusions can lead to improved error rates under this model is investigated.

The main focus of the analysis of errors in multi-way fusion circuits and of their impact on the performance of LO-FBQC architectures has been on their use in the context of linear optical generation of resource states. In the framework of FBQC, regardless of the chosen physical platform, fusions play another fundamental role as primitives of the computation, which is carried out by fusing qubits of different resource states according to the structure of the fusion network. Given the advantages found in the application of linear optical multi-way fusions for RSG, it is natural to ask whether the performance of an LO-FBQC architecture would also benefit from their use in the fusion network, where more than two resource states could be combined with a single operation.

One key difference between RSG and Fusion Network (FN) fusions is that the former are multiplexed, which mitigates the impact of heralded erasure, while the latter are not. Erasure is the dominant type of error in LO-FBQC, since one of its main causes is photon loss, the dominant type of error in the

linear optical platform. In the absence of multiplexing, erasure in the fusion network needs to be handled by the underlying QEC code, so minimising its rate is one of the main considerations in the choice of how FN fusions are performed. As explained in Chapter 6, one way of reducing erasure is to generate resource states whose qubits are locally encoded in small QEC codes. In this chapter, we investigate whether performing physical multi-way fusions in the fusion network can help mitigate this type of error further. In order to reach a conclusive answer on the usefulness of multi-way FN fusions, other aspects would ultimately need to be taken into account, such as their effect on the Pauli error rate in the fusion network. However, erasure occurs at a significantly higher rate, which is why we begin by focusing only on this aspect in the analysis presented here.

This chapter introduces a new model for erasure due to photon loss in fully destructive fusions which takes into account partial stabilizer information which may be recovered in linear optical multi-way fusions even when not all photons reach the detectors. The ability of a few n-GHZ fusion circuits to produce this kind of information is assessed and compared to an analytically-derived minimum amount which would need to be met to match the performance of a T2 fusion under this model. The main result of this analysis is that the T2 fusion appears to minimise overall erasure compared to all the larger fusions considered, even when partial stabilizer information is accounted for. Therefore, we conclude that despite the fact that they project onto larger entangled states upon success, multi-way fusions should not be used in place of T2s in the fusion network, as they would lead to an increase in the erasure rate and therefore a decrease in the error threshold of the architecture.

7.1 Total erasure from photon loss

Any physical error mechanism which alters the number of photons entering a fully destructive fusion causes heralded errors by producing patterns with the wrong number of photons at the detectors. Sources of noise of this kind include photon loss, multiphoton contamination, and detector dark

counts. Since the former is the dominant source of error in linear optics, we focus on the scenario in which it is the only type of error present in the system here. Its effect is commonly described as an erasure of all information about the measurement which was performed, so that none of the stabilizer generators of the operator are recovered [43, 115, 144, 146, 181]. Assuming that in the absence of physical errors fully destructive fusions produce stabilizer outcomes regardless of whether success or failure is obtained, the probability of recovering stabilizer information from the fusion is given by:

$$R_{\text{stab}} = \eta^n, \tag{7.1}$$

where η is the probability that a photon is not lost, i.e., $1-l$ where l is the loss rate. So R_{stab} is the probability that none of the input qubits of an n-way fusion are lost. Larger fusions are exponentially more likely to have their outcomes erased than smaller ones, which makes T2 fusions with $n = 2$ the best choice to minimise erasure in the fusion network.

This model does not allow for the possibility that some information about the fusion measurement may still be available even when not all photons reach the detectors. In the following section, we relax this assumption and introduce a less conservative model which allows to make use of partial information.

7.2 Accounting for partial erasure

In principle, some detection patterns with fewer than the expected number of photons could contain enough information to allow the recovery of stabilizer measurements on a subset of the input qubits. For example, consider the standard 3-GHZ fusion circuit in fig. 7.1. If the detection pattern obtained has two photons bunched in the top mode and no photons in all the other detectors, the two photons must have come from the top and bottom rails of the circuit and the middle qubit must have been lost. So, even though only two photons were detected, the top and bottom qubits are projected onto the pure

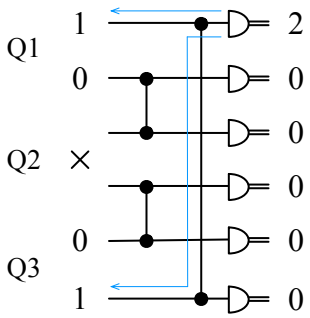

Figure 7.1: Example pattern which allows recovery of partial stabilizer information from a standard 3-GHZ fusion when one photon is lost. Backpropagation of the detected photons shows that the state of qubits 1 and 3 can be inferred from this bunched two-photon pattern, while qubit 2 is assumed to be lost.

states $|0\rangle = |10\rangle\rangle$ and $|1\rangle = |01\rangle\rangle$ respectively and one of the 3-GHZ stabilizers can be reconstructed, namely ZIZ. In this case, the stabilizer information is only partially erased.

As the 3-GHZ fusion example illustrates, the ability of a fusion to provide useful stabilizer information in the presence of loss is highly dependent on the structure of the circuit. In order to quantify the ability of a specific n-way fusion circuit to recover from the loss of m photons, we consider the expected size of the stabilizer group associated with the measurements it performs.

Given a detection pattern x with only $n-m$ photons, we extract the stabilizer group which describes the measurement operator applied to the inputs, $S_{m,x}$. If no photons are lost and a pure stabilizer measurement is performed, this group has $|S_{m,x}| = 2^n$ elements, so $\log_2(|S_{m,x}|) = n$. When $m > 0$ we have $\log_2(|S_{m,x}|) < n$, and in the case that no stabilizer information can be recovered, $\log_2(|S_{m,x}|) = 0$. By considering all the detection patterns with $n-m$ photons, X_{n-m}, we can calculate the expected size of the measured stabilizer group:

$$G_m = \sum_{x \in X_{n-m}} \log_2(|S_{m,x}|) \Pr(x), \tag{7.2}$$

where the $\log_2$ is used to compare the result to the number of generators of the ideal stabilizer group. This expression takes values $0 \leq G_m \leq (n-m)$, with the extremes corresponding to the cases in which none or all of the qubits which are not lost are projected onto stabilizer states by all detection patterns.

It is useful to normalize this as follows:

$$r_m = \frac{G_m}{n-m}, \tag{7.3}$$

such that $0 \leq r_m \leq 1$. From this, a new model can be defined which accounts for the possibility of only partial erasure in case of photon loss:

$$\begin{aligned} R'_{\text{stab}} &= \sum_{m=0}^{n} r_m \frac{n-m}{n} \binom{n}{m} (1-\eta)^m \eta^{n-m} = \\ &= \sum_{m=0}^{n} r_m \binom{n-1}{m} (1-\eta)^m \eta^{n-m}, \end{aligned} \tag{7.4}$$

where the $\frac{n-m}{n}$ factor in the first line renormalizes r_m to the maximum possible number of stabilizer generators in the absence of loss, n, and in the second line we have used the identity $\frac{n-m}{n}\binom{n}{m} = \binom{n-1}{m}$. To first order, we are mostly interested in the case of a single loss event, i.e., we can set $r_m = 0$ for $k > 1$, for which eq. 7.4 becomes

$$R'_{\text{stab},n} = r_0 \eta^n + r_1 (n-1)(1-\eta)\eta^{n-1}. \tag{7.5}$$

In many cases, fusion circuits always perform a pure stabilizer measurement in the absence of loss, i.e., $r_0 = 1$. This is the case for all the standard n-GHZ fusions and the new circuits with one extra pair of input beamsplitters. However, some of the fusions introduced in Chapter 3, such as the depth-2 n-GHZ fusions, have ideal outcomes which project onto other types of states, e.g., W-states, so for those circuits $r_0 < 1$.

In the next section, we analyse the performance of a few different multi-way fusions under this erasure model and compare them to the T2 fusion to determine whether any of them would help mitigate erasure in the fusion network.

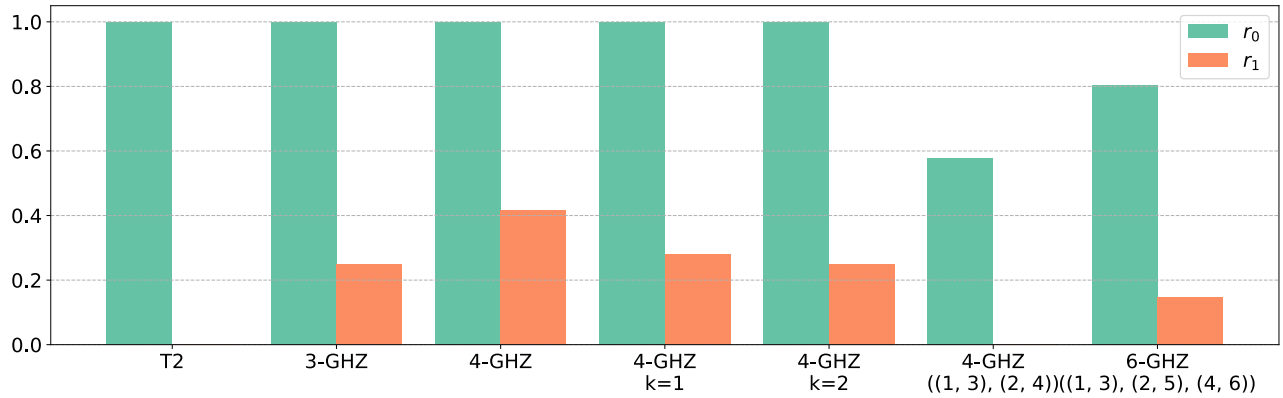

Figure 7.2: Values of the erasure coefficients r_0 (green) and r_1 (orange) for the T2 fusion and a few of the n-GHZ fusion circuits introduced in Chapter 3.

7.3 Erasure in lossy multi-way fusions

In order to evaluate the erasure rate of a circuit under the new model in eq. 7.4, the first step is to calculate their r_m values. This can be done by performing Fock state simulations of the ideal circuits with loss channels applied to the inputs to capture all linear losses, postselecting on the desired number of losses m, and extracting the stabilizer group of the measurements performed using the methods described in Chapter 5. The expected size of the group is then calculated for each m as in eq. 7.2 and used to obtain the r_m coefficients through eq. 7.3.

Focusing on the first-order loss model in eq. 7.5, fig. 7.2 shows the r_0 and r_1 values of a few n-GHZ fusion circuits with $n = 3,4,6$ as well as the standard T2 fusion. As expected, significantly less stabilizer information is available in all fusions when one of the input qubits is lost compared to the ideal case. The trend in the r_1 values (orange), suggests that for standard n-GHZ fusions more stabilizer information can be recovered when one photon is lost as n increases, but when other variants are considered, it is clear that the value of this coefficient is more tied to the structure of the circuit than to its size. Fusions with more additional beamsplitters are found to have lower r_1 values, which agrees with the intuition that the more interference makes it harder to reconstruct which photon was lost and therefore to infer the measurement performed on the rest of the qubits. The fact that for the T2 $r_1 = 0$ and for most multi-way fusion circuits, with the exception of the ((1, 3), (2, 4)) 4-GHZ

fusion circuit, $r_1 > 0$ confirms our expectation that larger fusions allow better recovery of information in the presence of loss. The question we want to answer now is whether the trade-off between the increased chance of photon loss with n and the additional information produced when one photon is missing is favourable for the overall erasure rate.

For a multi-way fusion to outperform a T2 in this context, it has to provide a larger amount of stabilizer information in the presence of loss, i.e., it has to be the case that $R'_{\text{stab},n} > R'_{\text{stab},2}$, where $R'_{\text{stab},n}$ is defined in eq. 7.5. This can be written as:

$$r_0\eta^n + r_1(n-1)(1-\eta)\eta^{n-1} > \eta^2, \tag{7.6}$$

where we have used the fact that $r_1 = 0$ for a two-way fusion. We know from fig. 7.2 that for some n-GHZ fusions $r_0 < 1$, but for now let us assume $r_0 = 1$ for all circuits for simplicity. Then we have:

$$r_1^{\min} = \frac{1}{n-1}\left(\frac{1}{\eta^{n-3}} + \frac{1}{\eta^{n-4}} + \ldots + 1\right), \tag{7.7}$$

where $r_1^{\min}$ is the minimum value of r_1 a circuit has to achieve to provide more stabilizer information than a T2 fusion in the presence of loss under a first-order error model. In the limit of low photon loss this becomes:

$$\lim_{\eta\to 1} r_1^{\min} = \frac{n-2}{n-1}. \tag{7.8}$$

This tends to 1 as n increases, meaning that as the size of the fusion grows, it needs to be able to recover more stabilizer information from the loss of a single photon. This limit also applies to the more general model with higher-order losses in eq. 7.4, i.e., with non-zero r_k for $k > 1$, since the low-loss regime is dominated by the terms capturing the loss of zero or one photon. If $r_0 < 1$, we see from eq. 7.6 that r_1 needs to increase to compensate for the information lost in the ideal case.

Fig. 7.3 compares the values of r_1 from fig. 7.2 (dark bars) to the minimum required value in eq. 7.8

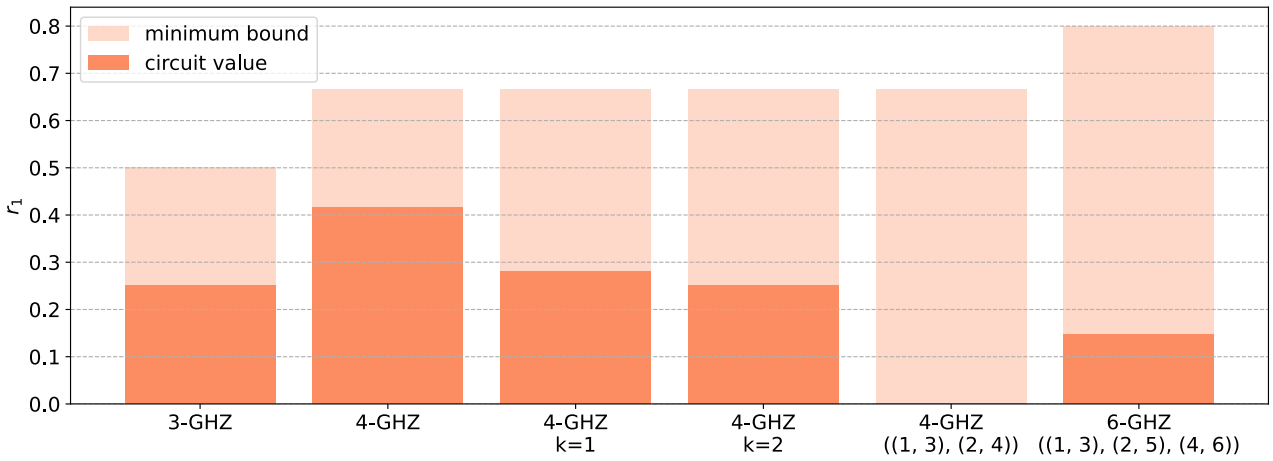

Figure 7.3: r_1 values for several n-GHZ fusions. None of them achieve the minimum $r_1^{\min}$ (lighter bars) needed to recover more stabilizer information than a T2 fusion in the presence of loss.

(lighter bars). The key takeaway from this plot is that none of the fusions achieve the minimum value of r_1 necessary to outperform a T2 fusion in the low loss regime. Moreover, the expression in brackets in eq 7.7 increases monotonically as η decreases, so the value of $r_1^{\min}$ grows with the loss rate, meaning that at higher loss rates the minimum amount of stabilizer information a fusion needs to recover is higher. Given this, the conclusion we draw is that none of the fusions analysed here leads to an improvement in erasure from photon loss compared to a T2 in any loss regime. To confirm this further, we plot the value of $R'_{\text{stab},n}$ for these fusions against the loss rate and compare it to the same ratio for a T2 fusion. This is shown in fig. 7.4. The two-way fusion curve (red, dashed) is consistently above the multi-way fusions for all values of η. Interestingly, the high-probability, depth-2 6-GHZ fusion outperforms the ((1, 3), (2, 4)) 4-GHZ fusion for loss rates below $\sim 25\%$, demonstrating that in this model of erasure it is possible for larger fusions to be preferable to smaller ones.

We have only considered a few examples of multi-way fusion circuits here, chosen as a representative sample of the many variants presented in Chapter 3. Given that less stabilizer information is recovered overall by standard n-GHZ fusions as n increases (see fig. 7.4), and that adding beamsplitters to the fusion inputs appears to reduce the value of r_1, we do not expect any of the other circuits previously described to be able to outperform T2 fusions. So we take this result as indication that it is preferable to perform two-way fusions between resource states in the fusion network to minimise the amount of

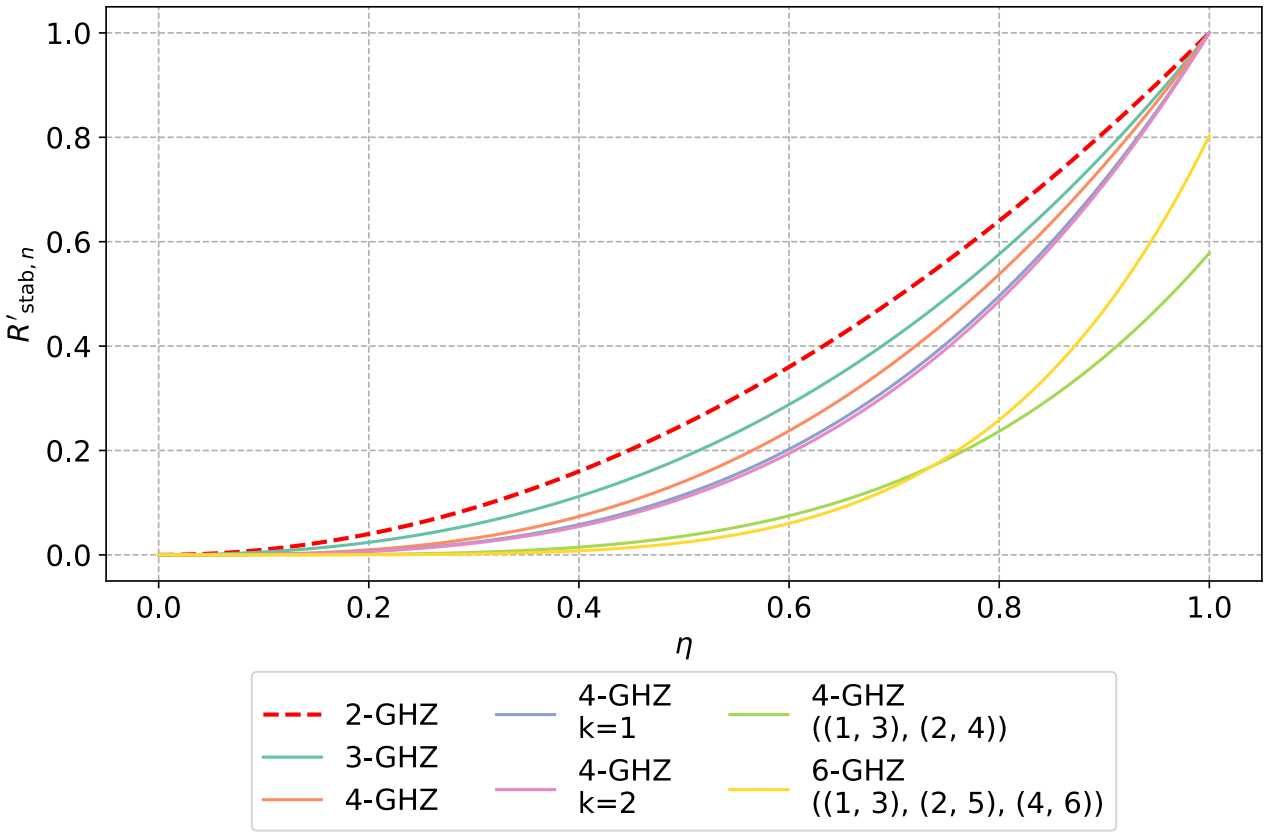

Figure 7.4: Expected ratio of input qubits undergoing a stabilizer measurement, $R'_{\text{stab},n}$, as a function of η. The T2 fusion (red, dashed) outperforms all the multi-way fusions considered here in all loss regimes. Interestingly, the depth-2 6-way fusion is better than the depth-2 4-way fusion at recovering stabilizer information for loss rates lower than $\sim 25\%$.

erasure in an LO-FBQC architecture.

7.4 Conclusion

After analysing the impact of using multi-way fusions in the generation of resource states for FBQC, here we consider whether they could also be useful as primitive operations between resource states in the fusion network. In this context, where multiplexing cannot be used to filter out heralded errors, these play a much more important role than in RSG. In fact, since they occur as a result of high probability events, such as fusion failure or photon loss, fusion outcome erasures become the dominant type of error to consider in this case.

We have seen that the success probability of multi-way fusions decreases with size, making them

less appealing in the absence of multiplexing. Moreover, their susceptibility to photon loss is also typically modelled as increasing exponentially with size, under the assumption that erasure occurs if any of the input photons are lost. In this chapter, we considered a less conservative model for erasure due to loss, which admits the recovery of some outcome information even if not all photons are found at the detectors. The ability of a fusion to recover this partial information entirely depends on the structure of its linear optical circuit.

We applied this model to several of the multi-way fusions introduced in Chapter 3 and compared their rate of erasure due to loss to that of a T2 fusion. Unfortunately, what we obtained was a negative result, which suggested that even when the availability of partial outcome information is accounted for, multi-way fusions suffer from a higher erasure rate compared to two-way fusions, which therefore remain the best option as fusion network fusions between resource states.

Conclusions and Outlook

A universal fault-tolerant quantum computer which can execute useful and ground-breaking computations currently not accessible to classical devices is widely expected to be, at least in its first incarnations, a large system of millions of physical qubits repeatedly undergoing rounds of quantum error correction. Such a complex machine can only be built in practice if the amount of physical noise it can tolerate can be made to meet the performance of the hardware platform used to realise it, while also keeping the resource requirements within realistic bounds. This is a significant challenge which many research groups across academia and industry have been tackling for decades using a variety of physical systems and architectural designs.

The framework of Fusion-Based Quantum Computation (FBQC) has allowed to make significant progress towards a fault-tolerant, universal quantum computing architecture over the last few years by directly integrating the capabilities and limitations of the underlying physical hardware into the design process. Any platform in which the two primitives of FBQC, entangled resource state and multi-qubit projective fusion measurements, are readily available is well suited for this paradigm. Linear optics is one of these. Linear Optical Quantum Computing (LOQC) benefits from a number of properties of photonic qubits which make this approach uniquely scalable, among them the ability to rely on high-volume semiconductor fabrication for all the components of the quantum computer.

In Linear Optical FBQC (LO-FBQC), fusions are used as primitives of the computation as well as in the generation of resource states. While most proposals presented so far have made use of fusions between two qubits, multi-way fusions between more qubits are just as easy to implement using linear

optics. The work described in this thesis investigated the usefulness of larger fusions to improve the performance of LO-FBQC architectures in terms of both resource overhead and error tolerance. In particular, we focused on the potential impact of multi-way fusions on the generation of resource states, based on the intuition that they would provide a path to reduce both their cost and the rate of Pauli errors on their qubits. The results described in Chapter 6 confirmed this starting hypothesis: it was found that multi-way fusions allow to generate the two resource states from the original FBQC proposal in Ref. [43] with a $12-18\times$ lower resource overhead and $\sim 10-28\%$ lower Pauli error rate compared to two-way fusions alone. These significant performance improvements stem from the combination of results described throughout this thesis.

In Chapter 3 we introduced new families of linear optical multi-way fusion circuits with improved entanglement generation capability, having either higher success probability compared to previously known circuits, or failure outcomes which entangle subsets of the input qubits, or in some cases both. The increased success rate of some of these circuits plays a significant role in reducing the cost of generating a resource state of fixed size. While our work focuses on the application of these fusions to resource state generation, with their increased entanglement generation rates, they can be advantageous in many other contexts in which the probabilistic nature of linear optical entangling measurements poses limitations.

In Chapters 4 and 5 we analysed the behaviour of several types of multi-way fusion circuits in the presence of noise, particularly focusing on the description of unheralded error mechanisms as Pauli error channels, since these are the most detrimental types of errors in resource state generation. We developed and applied analysis frameworks based on leading-order error physical models, which allow to gain insight into the expected behaviour of linear optical fusions without the added complexity of a very detailed analysis. In addition to providing an understanding of fusion circuits in isolation, these models are a fundamental component of the assessment of resource state generation protocols described in Chapter 6. Ultimately it is essential to evaluate any final architectural design using a

much more complex error model which accounts for many more possible physical noise mechanisms and which brings together all layers of the computing stack, from the linear optical circuits to the fault-tolerant protocol. This kind of modelling is a fundamental step towards the realisation of a working fault-tolerant quantum computer, but it also requires a very significant theoretical and computational effort, which makes it unsuitable for the kind of rapid iteration required as part of the design process. The analysis methods and models presented in these chapters provide the ability to perform a light-weight yet informative evaluation of schemes to guide the design towards quantum architectures with improved performance.

The analysis in Chapter 7 looked at another potential use of multiway fusions, as primitives of the computation in the fusion network, finding that two-way fusions provide better protection against erasure. However, this negative result obtained here for multi-way fusions should not be seen as a definitive evidence that these circuits can only lead to improvements in resource state generation in LO-FBQC. It is possible that taking into account the entanglement generated by different outcomes might favour the fusion circuits with entangling failures, or that other types of information are available to use other than the size of the stabilizer group. We leave further investigation of this question to future work.

Overall, the key conclusion of this work is that multi-way fusions hold significant promise for the design of LO-FBQC architectures with improved resource overhead and tolerance to physical errors, helping move current proposals towards the design of a realistic fault-tolerant quantum computing architecture. Along the way towards this result, novel linear optical circuits, error models, and analysis frameworks were introduced which can be used in other contexts and beyond the specific examples considered in this thesis to investigate further improvements to quantum architecture designs.

Bibliography

[1] R. P. Feynman. “Simulating Physics with Computers”. *International Journal of Theoretical Physics* **21** (6/7), (1982).

[2] L. K. Grover. “A fast quantum mechanical algorithm for database search”. In: *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pp. 212–19.

[3] P. W. Shor. “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer”. *SIAM Review* **41** (2), pp. 303–32 (1999).

[4] M. T. Nguyen, Y.-L. Lee, D. Alfonso, Q. Shao, and Y. Duan. “Description of reaction and vibrational energetics of CO2–NH3 interaction using quantum computing algorithms”. *AVS Quantum Science* **5** (1), (2023).

[5] I. H. Kim, Y.-H. Liu, S. Pallister, W. Pol, S. Roberts, and E. Lee. “Fault-tolerant resource estimate for quantum chemical simulations: Case study on Li-ion battery electrolyte molecules”. *Physical Review Research* **4** (2), p. 023019 (2022).

[6] Y.-h. Lam, Y. Abramov, R. S. Ananthula, J. M. Elward, L. R. Hilden, S. O. Nilsson Lill, P.-O. Norrby, A. Ramirez, E. C. Sherer, and J. Mustakis. “Applications of quantum chemistry in pharmaceutical process development: Current state and opportunities”. *Organic Process Research & Development* **24** (8), pp. 1496–507 (2020).

[7] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, and C. Ottaviani. “Advances in quantum cryptography”. *Advances in optics and photonics* **12** (4), pp. 1012–236 (2020).

[8] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd. “Quantum machine learning”. *Nature* **549** (7671), pp. 195–202 (2017).

[9] M. Steudtner, S. Morley-Short, W. Pol, S. Sim, C. L. Cortes, M. Loipersberger, R. M. Parrish, M. Degroote, N. Moll, and R. Santagati. “Fault-tolerant quantum computation of molecular observables”. *Quantum* **7**, p. 1164 (2023).

[10] A. Kan and B. Symons. “Resource-optimized fault-tolerant simulation of the Fermi-Hubbard model and high-temperature superconductor models”. arXiv:2411.02160 (2024).

[11] M. A. Nielsen and I. L. Chuang. “Quantum computation and quantum information”. Cambridge institution press, 2010.

[12] P. O. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan. “A new universal and fault-tolerant quantum basis”. *Information Processing Letters* **75** (3), pp. 101–07 (2000).

[13] R. Jozsa and N. Linden. “On the role of entanglement in quantum-computational speed-up”. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **459** (2036), pp. 2011–32 (2003).

[14] S. Ding and Z. Jin. “Review on the study of entanglement in quantum computation speedup”. *Chinese Science Bulletin* **52** (16), pp. 2161–66 (2007).

[15] C. Löschnauer, J. M. Toba, A. Hughes, S. King, M. Weber, R. Srinivas, R. Matt, R. Nourshargh, D. Allcock, and C. Ballance. “Scalable, high-fidelity all-electronic control of trapped-ion qubits”. arXiv:2407.07694 (2024).

[16] W. K. Wootters and W. H. Zurek. “A single quantum cannot be cloned”. *Nature* **299** (5886), pp. 802–03 (1982).

[17] D. Gottesman. “Class of quantum error-correcting codes saturating the quantum Hamming bound”. *Physical Review A* **54** (3), p. 1862 (1996).

[18] D. Gottesman. “The Heisenberg representation of quantum computers”. arXiv:quant-ph/9807006 (1998).

[19] D. Gottesman. “Stabilizer codes and quantum error correction”. PhD thesis. California Institute of Technology, 1997.

[20] N. Delfosse, A. Paz, A. Vaschillo, and K. M. Svore. “How to choose a decoder for a fault-tolerant quantum computer? The speed vs accuracy trade-off”. arXiv:2310.15313 (2023).

[21] P. W. Shor. “Scheme for reducing decoherence in quantum computer memory”. *Physical Review A* **52** (4), R2493 (1995).

[22] M. Grassl, T. Beth, and T. Pellizzari. “Codes for the quantum erasure channel”. *Physical Review A* **56** (1), p. 33 (1997).

[23] T. C. Ralph, A. Hayes, and A. Gilchrist. “Loss-tolerant optical qubits”. *Physical Review Letters* **95** (10), p. 100501 (2005).

[24] A. R. Calderbank and P. W. Shor. “Good quantum error-correcting codes exist”. *Physical Review A* **54** (2), p. 1098 (1996).

[25] A. Steane. “Multiple-particle interference and quantum error correction”. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **452** (1954), pp. 2551–77 (1996).

[26] A. Kitaev. “Fault-tolerant quantum computation by anyons”. *Annals of physics* **303** (1), pp. 2–30 (2003).

[27] S. B. Bravyi and A. Y. Kitaev. “Quantum codes on a lattice with boundary”. arXiv:quant-ph/9811052 (1998).

[28] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill. “Topological quantum memory”. *Journal of Mathematical Physics* **43** (9), pp. 4452–505 (2002).

[29] G. Q. AI. “Suppressing quantum errors by scaling a surface code logical qubit”. *Nature* **614** (7949), pp. 676–81 (2023).

[30] B. Eastin and E. Knill. “Restrictions on transversal encoded quantum gate sets”. *Physical Review Letters* **102** (11), p. 110502 (2009).

[31] S. Bravyi and A. Kitaev. “Universal quantum computation with ideal Clifford gates and noisy ancillas”. *Physical Review A* **71** (2), p. 022316 (2005).

[32] D. Aharonov and M. Ben-Or. “Fault-tolerant quantum computation with constant error”. In: *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pp. 176–88.

[33] C. Wang, J. Harrington, and J. Preskill. “Confinement-Higgs transition in a disordered gauge theory and the accuracy threshold for quantum memory”. *Annals of Physics* **303** (1), pp. 31–58 (2003).

[34] T. M. Stace, S. D. Barrett, and A. C. Doherty. “Thresholds for topological codes in the presence of loss”. *Physical Review Letters* **102** (20), p. 200501 (2009).

[35] D. P. DiVincenzo. “The physical implementation of quantum computation”. *Fortschritte der Physik: Progress of Physics* **48** (9-11), pp. 771–83 (2000).

[36] C. A. Pérez-Delgado and P. Kok. “Quantum computers: Definition and implementations”. *Physical Review A* **83** (1), p. 012303 (2011).

[37] J. M. Gambetta, J. M. Chow, and M. Steffen. “Building logical qubits in a superconducting quantum computing system”. *npj quantum information* **3** (1), p. 2 (2017).

[38] H. Häffner, C. F. Roos, and R. Blatt. “Quantum computing with trapped ions”. *Physics reports* **469** (4), pp. 155–203 (2008).

[39] D. S. Weiss and M. Saffman. “Quantum computing with neutral atoms”. *Physics Today* **70** (7), pp. 44–50 (2017).

[40] D. Loss and D. P. DiVincenzo. “Quantum computation with quantum dots”. *Physical Review A* **57** (1), p. 120 (1998).

[41] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn. “Linear optical quantum computing with photonic qubits”. *Reviews of modern physics* **79** (1), pp. 135–74 (2007).

[42] R. Raussendorf and H. J. Briegel. “A one-way quantum computer”. *Physical Review Letters* **86** (22), p. 5188 (2001).

[43] S. Bartolucci, P. Birchall, H. Bombín, H. Cable, C. Dawson, M. Gimeno-Segovia, E. Johnston, K. Kieling, N. Nickerson, M. Pant, F. Pastawski, T. Rudolph, and C. Sparrow. “Fusion-based quantum computation”. *Nature Communications* **14** (1), p. 912 (2023).

[44] T. Rudolph. “Quantum computing at the speed of light”. YouTube - Caltech PMA. Apr. 18, 2025.

[45] R. Raussendorf, D. E. Browne, and H. J. Briegel. “Measurement-based quantum computation on cluster states”. *Physical Review A* **68** (2), p. 022312 (2003).

[46] M. Van den Nest, A. Miyake, W. Dür, and H. J. Briegel. “Universal resources for measurement-based quantum computation”. *Physical Review Letters* **97** (15), p. 150504 (2006).

[47] D. Gross and J. Eisert. “Novel schemes for measurement-based quantum computation”. *Physical Review Letters* **98** (22), p. 220503 (2007).

[48] R. Raussendorf and J. Harrington. “Fault-tolerant quantum computation with high threshold in two dimensions”. *Physical Review Letters* **98** (19), p. 190504 (2007).

[49] R. Raussendorf, J. Harrington, and K. Goyal. “Topological fault-tolerance in cluster state quantum computation”. *New Journal of Physics* **9** (6), p. 199 (2007).

[50] M. Newman, L. A. de Castro, and K. R. Brown. “Generating fault-tolerant cluster states from crystal structures”. *Quantum* **4**, p. 295 (2020).

[51] B. J. Brown and S. Roberts. “Universal fault-tolerant measurement-based quantum computation”. *Physical Review Research* **2** (3), p. 033305 (2020).

[52] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters. “Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels”. *Physical Review Letters* **70** (13), p. 1895 (1993).

[53] H. Bombín, C. Dawson, R. V. Mishmash, N. Nickerson, F. Pastawski, and S. Roberts. “Logical blocks for fault-tolerant topological quantum computation”. *PRX Quantum* **4** (2), p. 020303 (2023).

[54] H. Bombín, C. Dawson, T. Farrelly, Y. Liu, N. Nickerson, M. Pant, F. Pastawski, and S. Roberts. “Fault-tolerant complexes”. arXiv:2308.07844 (2023).

[55] H. Bombín, D. Litinski, N. Nickerson, F. Pastawski, and S. Roberts. “Unifying flavors of fault tolerance with the ZX calculus”. *Quantum* **8**, p. 1379 (2024).

[56] B. Coecke and R. Duncan. “Interacting quantum observables”. In: *International Colloquium on Automata, Languages, and Programming*. Springer, 2008, pp. 298–310.

[57] B. Coecke and R. Duncan. “Interacting quantum observables: categorical algebra and diagrammatics”. *New Journal of Physics* **13** (4), p. 043016 (2011).

[58] M. Backens, H. Miller-Bakewell, G. de Felice, L. Lobski, and J. van de Wetering. “There and back again: A circuit extraction tale”. *Quantum* **5**, p. 421 (2021).

[59] R. Duncan, A. Kissinger, S. Perdrix, and J. Van De Wetering. “Graph-theoretic Simplification of Quantum Circuits with the ZX-calculus”. *Quantum* **4**, p. 279 (2020).

[60] N. Chancellor, A. Kissinger, S. Zohren, J. Roffe, and D. Horsman. “Graphical structures for design and verification of quantum error correction”. *Quantum science and technology* **8** (4), p. 045028 (2023).

[61] E. Jeandel, S. Perdrix, and R. Vilmart. “A complete axiomatisation of the ZX-calculus for Clifford+ T quantum mechanics”. In: *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*, pp. 559–68.

[62] A. Hadzihasanovic, K. F. Ng, and Q. Wang. “Two complete axiomatisations of pure-state qubit quantum computing”. In: *Proceedings of the 33rd annual ACM/IEEE symposium on logic in computer science*, pp. 502–11.

[63] M. Backens. “The ZX-calculus is complete for stabilizer quantum mechanics”. *New Journal of Physics* **16** (9), p. 093021 (2014).

[64] J. van de Wetering. “ZX-calculus for the working quantum computer scientist”. arXiv:2012.13966 (2020).

[65] B. Coecke and A. Kissinger. “Picturing quantum processes: A first course on quantum theory and diagrammatic reasoning”. In: *Diagrammatic Representation and Inference: 10th International Conference, Diagrams 2018, Edinburgh, UK, June 18-22, 2018, Proceedings 10*. Springer. 2018, pp. 28–31.

[66] B. Coecke. “Basic ZX-calculus for students and professionals”. arXiv:2303.03163 (2023).

[67] M. McEwen, D. Bacon, and C. Gidney. “Relaxing hardware requirements for surface code circuits using time-dynamics”. *Quantum* **7**, p. 1172 (2023).

[68] N. C. Harris, Y. Ma, J. Mower, T. Baehr-Jones, D. Englund, M. Hochberg, and C. Galland. “Efficient, compact and low loss thermo-optic phase shifter in silicon”. *Optics express* **22** (9), pp. 10487–93 (2014).

[69] J. W. Silverstone, D. Bonneau, J. L. O'Brien, and M. G. Thompson. "Silicon quantum photonics". *IEEE Journal of Selected Topics in Quantum Electronics* **22** (6), pp. 390–402 (2016).

[70] T. Rudolph. "Why I am optimistic about the silicon-photonic route to quantum computing". *APL photonics* **2** (3), (2017).

[71] L. Feng, M. Zhang, J. Wang, X. Zhou, X. Qiang, G. Guo, and X. Ren. "Silicon photonic devices for scalable quantum information applications". *Photonics Research* **10** (10), A135–A53 (2022).

[72] J. L. O'Brien. "Optical quantum computing". *Science* **318** (5856), pp. 1567–70 (2007).

[73] J. Carolan, C. Harrold, C. Sparrow, E. Martín-López, N. J. Russell, J. W. Silverstone, P. J. Shadbolt, N. Matsuda, M. Oguma, and M. Itoh. "Universal linear optics". *Science* **349** (6249), pp. 711–16 (2015).

[74] X. Qiang, X. Zhou, J. Wang, C. M. Wilkes, T. Loke, S. O'Gara, L. Kling, G. D. Marshall, R. Santagati, and T. C. Ralph. "Large-scale silicon quantum photonics implementing arbitrary two-qubit processing". *Nature photonics* **12** (9), pp. 534–39 (2018).

[75] J. M. Arrazola, V. Bergholm, K. Brádler, T. R. Bromley, M. J. Collins, I. Dhand, A. Fumagalli, T. Gerrits, A. Goussev, and L. G. Helt. "Quantum circuits with many photons on a programmable nanophotonic chip". *Nature* **591** (7848), pp. 54–60 (2021).

[76] C. Taballione, M. C. Anguita, M. de Goede, P. Venderbosch, B. Kassenberg, H. Snijders, N. Kannan, W. L. Vleeshouwers, D. Smith, and J. P. Epping. "20-mode universal quantum photonic processor". *Quantum* **7**, p. 1071 (2023).

[77] PsiQuantum Team. "A manufacturable platform for photonic quantum computing". *Nature* (2025).

[78] M. A. Foster, A. C. Turner, J. E. Sharping, B. S. Schmidt, M. Lipson, and A. L. Gaeta. "Broadband optical parametric gain on a silicon photonic chip". *Nature* **441** (7096), pp. 960–63 (2006).

[79] A. Christ and C. Silberhorn. "Limits on the deterministic creation of pure single-photon states using parametric down-conversion". *Physical Review A* **85** (2), p. 023829 (2012).

[80] A. L. Migdall, D. Branning, and S. Castelletto. "Tailoring single-photon and multiphoton probabilities of a single-photon on-demand source". *Physical Review A* **66** (5), p. 053805 (2002).

[81] C.-K. Hong, Z.-Y. Ou, and L. Mandel. "Measurement of subpicosecond time intervals between two photons by interference". *Physical Review Letters* **59** (18), p. 2044 (1987).

[82] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani. "Experimental realization of any discrete unitary operator". *Physical Review Letters* **73** (1), p. 58 (1994).

[83] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walmsley. "Optimal design for universal multiport interferometers". *Optica* **3** (12), pp. 1460–65 (2016).

[84] T. Rudolph. "Photonic quantum computing with probabilistic single photon sources but without coherent switches". arXiv:2303.03454 (2023).

[85] P. Imany, Z. Wang, R. A. DeCrescent, R. C. Boutelle, C. A. McDonald, T. Autry, S. Berweger, P. Kabos, S. W. Nam, and R. P. Mirin. "Quantum phase modulation with acoustic cavities and quantum dots". *Optica* **9** (5), pp. 501–04 (2022).

[86] T. Pittman, B. Jacobs, and J. Franson. “Single photons on pseudodemand from stored parametric down-conversion”. *Physical Review A* **66** (4), p. 042303 (2002).

[87] K. T. McCusker, N. A. Peters, A. P. VanDevender, and P. G. Kwiat. “A deterministic single-photon source”. In: *Conference on Lasers and Electro-Optics*. Optica Publishing Group, 2008, JTuA117.

[88] T. Jennewein, M. Barbieri, and A. G. White. “Single-photon device requirements for operating linear optics quantum computing outside the post-selection basis”. *Journal of Modern Optics* **58** (3-4), pp. 276–87 (2011).

[89] X.-s. Ma, S. Zotter, J. Kofler, T. Jennewein, and A. Zeilinger. “Experimental generation of single photons via active multiplexing”. *Physical Review A* **83** (4), p. 043814 (2011).

[90] D. Bonneau, G. J. Mendoza, J. L. O’Brien, and M. G. Thompson. “Effect of loss on multiplexed single-photon sources”. *New Journal of Physics* **17** (4), p. 043057 (2015).

[91] G. J. Mendoza, R. Santagati, J. Munns, E. Hemsley, M. Piekarek, E. Martín-López, G. D. Marshall, D. Bonneau, M. G. Thompson, and J. L. O’Brien. “Active temporal and spatial multiplexing of photons”. *Optica* **3** (2), pp. 127–32 (2016).

[92] F. Kaneda and P. G. Kwiat. “High-efficiency single-photon generation via large-scale active time multiplexing”. *Science Advances* **5** (10), eaaw8586 (2019).

[93] M. Gimeno-Segovia, H. Cable, G. J. Mendoza, P. Shadbolt, J. W. Silverstone, J. Carolan, M. G. Thompson, J. L. O’Brien, and T. Rudolph. “Relative multiplexing for minimising switching in linear-optical quantum computing”. *New Journal of Physics* **19** (6), p. 063013 (2017).

[94] S. Bartolucci, P. Birchall, D. Bonneau, H. Cable, M. Gimeno-Segovia, K. Kieling, N. Nickerson, T. Rudolph, and C. Sparrow. “Switch networks for photonic fusion-based quantum computing”. arXiv:2109.13760 (2021).

[95] R. H. Hadfield. “Single-photon detectors for optical quantum information applications”. *Nature photonics* **3** (12), pp. 696–705 (2009).

[96] I. Esmaeil Zadeh, J. Chang, J. W. Los, S. Gyger, A. W. Elshaari, S. Steinhauer, S. N. Dorenbos, and V. Zwiller. “Superconducting nanowire single-photon detectors: A perspective on evolution, state-of-the-art, future developments, and applications”. *Applied Physics Letters* **118** (19), (2021).

[97] S. Bartolucci, P. M. Birchall, M. Gimeno-Segovia, E. Johnston, K. Kieling, M. Pant, T. Rudolph, J. Smith, C. Sparrow, and M. D. Vidrighin. “Creation of entangled photonic states using linear optics”. arXiv:2106.13825 (2021).

[98] P. Kok and S. L. Braunstein. “Limitations on the creation of maximal entanglement”. *Physical Review A* **62** (6), p. 064301 (2000).

[99] E. Knill, R. Laflamme, and G. J. Milburn. “A scheme for efficient quantum computation with linear optics”. *Nature* **409** (6816), pp. 46–52 (2001).

[100] S. Scheel, K. Nemoto, W. J. Munro, and P. L. Knight. “Measurement-induced nonlinearity in linear optics”. *Physical Review A* **68** (3), p. 032310 (2003).

[101] M. Gimeno-Segovia. “Towards practical linear optical quantum computing”. PhD thesis. Imperial College London, 2016.

[102] J. D. Franson, M. Donegan, M. Fitch, B. Jacobs, and T. Pittman. "High-fidelity quantum logic operations using linear optical elements". *Physical Review Letters* **89** (13), p. 137901 (2002).

[103] F. M. Spedalieri, H. Lee, and J. P. Dowling. "High-fidelity linear optical quantum computing with polarization encoding". *Physical Review A* **73** (1), p. 012334 (2006).

[104] N. Yoran and B. Reznik. "Deterministic linear optics quantum computation with single photon qubits". *Physical Review Letters* **91** (3), p. 037903 (2003).

[105] M. A. Nielsen. "Optical quantum computation using cluster states". *Physical Review Letters* **93** (4), p. 040503 (2004).

[106] D. E. Browne and T. Rudolph. "Resource-efficient linear optical quantum computation". *Physical Review Letters* **95** (1), p. 010501 (2005).

[107] C. M. Dawson, H. L. Haselgrove, and M. A. Nielsen. "Noise thresholds for optical cluster-state quantum computation". *Physical Review A* **73** (5), p. 052306 (2006).

[108] K. Kieling, T. Rudolph, and J. Eisert. "Percolation, renormalization, and quantum computing with nondeterministic gates". *Physical Review Letters* **99** (13), p. 130501 (2007).

[109] K. Kieling, D. Gross, and J. Eisert. "Minimal resources for linear optical one-way computing". *Journal of the Optical Society of America B* **24** (2), pp. 184–88 (2007).

[110] M. Gimeno-Segovia, P. Shadbolt, D. E. Browne, and T. Rudolph. "From three-photon Greenberger-Horne-Zeilinger states to ballistic universal quantum computation". *Physical Review Letters* **115** (2), p. 020502 (2015).

[111] H. A. Zaidi, C. Dawson, P. van Loock, and T. Rudolph. "Near-deterministic creation of universal cluster states with probabilistic Bell measurements and three-qubit resource states". *Physical Review A* **91** (4), p. 042301 (2015).

[112] T. J. Bell, L. A. Pettersson, and S. Paesani. "Optimizing graph codes for measurement-based loss tolerance". *PRX Quantum* **4** (2), p. 020328 (2023).

[113] S. Paesani and B. J. Brown. "High-Threshold Quantum Computing by Fusing One-Dimensional Cluster States". *Physical Review Letters* **131** (12), p. 120603 (2023).

[114] K. Sahay, J. Claes, and S. Puri. "Tailoring Fusion-Based Error Correction for High Thresholds to Biased Fusion Failures". *Physical Review Letters* **131** (12), p. 120604 (2023).

[115] H. Bombín, C. Dawson, N. Nickerson, M. Pant, and J. Sullivan. "Increasing error tolerance in quantum computers with dynamic bias arrangement". arXiv:2303.16122 (2023).

[116] B. Pankovich, A. Kan, K. H. Wan, M. Ostmann, A. Neville, S. Omkar, A. Sohbi, and K. Brádler. "High-photon-loss threshold quantum computing using ghz-state measurements". *Physical Review Letters* **133** (5), p. 050604 (2024).

[117] W. Song, N. Kang, Y.-S. Kim, and S.-W. Lee. "Encoded-Fusion-Based Quantum Computation for High Thresholds with Linear Optics". *Physical Review Letters* **133** (5), p. 050605 (2024).

[118] H. Weinfurter. "Experimental Bell-state analysis". *Europhysics Letters* **25** (8), p. 559 (1994).

[119] S. L. Braunstein and A. Mann. "Measurement of the Bell operator and quantum teleportation". *Physical Review A* **51** (3), R1727 (1995).

[120] M. Michler, K. Mattle, H. Weinfurter, and A. Zeilinger. "Interferometric Bell-state analysis". *Physical Review A* **53** (3), R1209 (1996).

[121] L. Vaidman and N. Yoran. "Methods for reliable teleportation". *Physical Review A* **59** (1), p. 116 (1999).

[122] N. Lütkenhaus, J. Calsamiglia, and K.-A. Suominen. “Bell measurements for teleportation”. *Physical Review A* **59** (5), p. 3295 (1999).

[123] J. Calsamiglia and N. Lütkenhaus. “Maximum efficiency of a linear-optical Bell-state analyzer”. *Applied Physics B* **72**, pp. 67–71 (2001).

[124] Y.-H. Kim, S. P. Kulik, and Y. Shih. “Quantum teleportation of a polarization state with a complete Bell state measurement”. *Physical Review Letters* **86** (7), p. 1370 (2001).

[125] S. Walborn, S. Pádua, and C. Monken. “Hyperentanglement-assisted Bell-state analysis”. *Physical Review A* **68** (4), p. 042313 (2003).

[126] H. A. Zaidi and P. van Loock. “Beating the one-half limit of ancilla-free linear optics Bell measurements”. *Physical Review Letters* **110** (26), p. 260501 (2013).

[127] W. P. Grice. “Arbitrarily complete Bell-state measurement using only linear optical elements”. *Physical Review A* **84** (4), p. 042331 (2011).

[128] F. Ewert and P. van Loock. “3/4-efficient bell measurement with passive linear optics and unentangled ancillae”. *Physical Review Letters* **113** (14), p. 140403 (2014).

[129] A. Olivo and F. Grosshans. “Ancilla-assisted linear optical bell measurements and their optimality”. *Physical Review A* **98** (4), p. 042323 (2018).

[130] J. A. Smith and L. Kaplan. “Approaching near-perfect state discrimination of photonic Bell states through the use of unentangled ancilla photons”. arXiv:1802.10527 (2018).

[131] F. Schmidt and P. van Loock. “Generalized fusions of photonic quantum states using static linear optics”. arXiv:2410.20261 (2024).

[132] M. Hein, W. Dür, J. Eisert, R. Raussendorf, M. Van den Nest, and H.-J. Briegel. “Entanglement in graph states and its applications”. In: *Quantum computers, algorithms and chaos*. IOS Press, 2006, pp. 115–218.

[133] S. Stanisic, N. Linden, A. Montanaro, and P. S. Turner. “Generating entanglement with linear optics”. *Physical Review A* **96** (4), p. 043861 (2017).

[134] I. Forbes, F. Ghafari, E. C. Deacon, S. P. Singh, E. Lavie, P. Yard, R. D. Shaw, A. Laing, and N. Tischler. “Heralded generation of entanglement with photons”. arXiv:2502.009823 (2025).

[135] N. H. Lindner and T. Rudolph. “Proposal for pulsed on-demand sources of photonic cluster state strings”. *Physical Review Letters* **103** (11), p. 113602 (2009).

[136] D. Huber, M. Reindl, J. Aberl, A. Rastelli, and R. Trotta. “Semiconductor quantum dots as an ideal source of polarization-entangled photon pairs on-demand: a review”. *Journal of Optics* **20** (7), p. 073002 (2018).

[137] M. Gimeno-Segovia, T. Rudolph, and S. E. Economou. “Deterministic generation of large-scale entangled photonic cluster state from interacting solid state emitters”. *Physical Review Letters* **123** (7), p. 070501 (2019).

[138] D. Istrati, Y. Pilnyak, J. Loredo, C. Antón, N. Somaschi, P. Hilaire, H. Ollivier, M. Esmann, L. Cohen, and L. Vidro. “Sequential generation of linear cluster states from a single photon emitter”. *Nature communications* **11** (1), p. 5501 (2020).

[139] S. Paesani. “Photonic quantum computing with quantum emitters”. In: *Photonics for Quantum 2023*. SPIE, PC126330W.

[140] D. M. Greenberger, M. A. Horne, and A. Zeilinger. “Going beyond Bell’s theorem”. In: *Bell’s theorem, quantum theory and conceptions of the universe*. Springer, 1989, pp. 69–72.

[141] Q. Zhang, X.-H. Bao, C.-Y. Lu, X.-Q. Zhou, T. Yang, T. Rudolph, and J.-W. Pan. “Demonstration of a scheme for the generation of “event-ready” entangled photon pairs from a single-photon source”. *Physical Review A* **77** (6), p. 062316 (2008).

[142] M. Varnava, D. E. Browne, and T. Rudolph. “How Good Must Single Photon Sources and Detectors Be for Efficient Linear Optical Quantum Computation?” *Physical Review Letters* **100** (6), p. 060502 (2008).

[143] B. Pankovich, A. Neville, A. Kan, S. Omkar, K. H. Wan, and K. Brádler. “Flexible entangled-state generation in linear optics”. *Physical Review A* **110** (3), p. 032402 (2024).

[144] P. Hilaire, Y. Castor, E. Barnes, S. E. Economou, and F. Grosshans. “Linear optical logical bell state measurements with optimal loss-tolerance threshold”. *PRX Quantum* **4** (4), p. 040322 (2023).

[145] M. Bergmann and P. van Loock. “Quantum error correction against photon loss using NOON states”. *Physical Review A* **94** (1), p. 012311 (2016).

[146] S.-W. Lee, T. C. Ralph, and H. Jeong. “Fundamental building block for all-optical scalable quantum networks”. *Physical Review A* **100** (5), p. 052303 (2019).

[147] H. Bombín, I. H. Kim, D. Litinski, N. Nickerson, M. Pant, F. Pastawski, S. Roberts, and T. Rudolph. “Interleaving: Modular architectures for fault-tolerant photonic quantum computing”. arXiv:2103.08612 (2021).

[148] Y. Li, P. C. Humphreys, G. J. Mendoza, and S. C. Benjamin. “Resource Costs for Fault-Tolerant Linear Optical Quantum Computing”. *Physical Review X* **5** (4), p. 041007 (2015).

[149] H. L. Haselgrove and P. P. Rohde. “Trade-off between the tolerance of located and unlocated errors in nondegenerate quantum error-correcting codes”. arXiv:quant-ph/0605183 (2006).

[150] P. P. Rohde, T. C. Ralph, and W. J. Munro. “Error tolerance and tradeoffs in loss-and failure-tolerant quantum computing schemes”. *Physical Review A* **75** (1), p. 010302 (2007).

[151] J.-w. Pan and A. Zeilinger. “Greenberger-horne-zeilinger-state analyzer”. *Physical Review A* **57** (3), p. 2208 (1998).

[152] R. H. Dicke. “Coherence in Spontaneous Radiation Processes”. *Physical Review* **93** (1), pp. 99–110 (1954).

[153] W. Dür, G. Vidal, and J. I. Cirac. “Three qubits can be entangled in two inequivalent ways”. *Physical Review A* **62** (6), p. 062314 (2000).

[154] J. J. Sylvester. “LX. Thoughts on inverse orthogonal matrices, simultaneous signsuccessions, and tessellated pavements in two or more colours, with applications to Newton’s rule, ornamental tile-work, and the theory of numbers”. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **34** (232), pp. 461–75 (1867).

[155] J. Hadamard. “Résolution d’une question relative aux déterminants”. *Bull. des sciences math.* **2**, pp. 240–46 (1893).

[156] W. Tadej and K. Życzkowski. “A concise guide to complex Hadamard matrices”. *Open Systems & Information Dynamics* **13** (2), pp. 133–77 (2006).

[157] F. Szöllősi. “Construction, classification and parametrization of complex Hadamard matrices”. arXiv:1110.5590 (2011).

[158] S. Scheel and N. Lütkenhaus. “Upper bounds on success probabilities in linear optics”. *New Journal of Physics* **6** (1), p. 51 (2004).

[159] J. Eisert. “Optimizing Linear Optics Quantum Gates”. *Physical Review Letters* **95** (4), p. 040502 (2005).

[160] D. B. Uskov, L. Kaplan, A. M. Smith, S. D. Huver, and J. P. Dowling. “Maximal success probabilities of linear-optical quantum gates”. *Physical Review A* **79** (4), p. 042326 (2009).

[161] P. P. Rohde and T. C. Ralph. “Error models for mode mismatch in linear optics quantum computing”. *Physical Review A* **73** (6), p. 062312 (2006).

[162] P. Birchall. “Fundamental Advantages and Practicalities of Quantum-Photonic Metrology and Computing”. PhD thesis. institution of Bristol, 2018.

[163] C. Sparrow. “Quantum interference in universal linear optical devices for quantum computation and simulation”. PhD thesis. Imperial College London, 2018.

[164] R. D. Shaw, A. E. Jones, P. Yard, and A. Laing. “Errors in heralded circuits for linear optical entanglement generation”. arXiv:2305.08452 (2023).

[165] J. Saied, J. Marshall, N. Anand, S. Grabbe, and E. G. Rieffel. “Advancing quantum networking: some tools and protocols for ideal and noisy photonic systems”. In: *Quantum Computing, Communication, and Simulation IV*. Vol. 12911. SPIE, pp. 37–65.

[166] S. M. Barnett, J. Jeffers, A. Gatti, and R. Loudon. “Quantum optics of lossy beam splitters”. *Physical Review A* **57** (3), p. 2134 (1998).

[167] M. Oszmaniec and D. J. Brod. “Classical simulation of photonic linear optics with lost particles”. *New Journal of Physics* **20** (9), p. 092002 (2018).

[168] P. Kok and S. L. Braunstein. “Detection devices in entanglement-based optical state preparation”. *Phys. Rev. A* **63** (3), p. 033812 (2001).

[169] Z. Y. Ou and L. Mandel. “Observation of Spatial Quantum Beating with Separated Photodetectors”. *Physical Review Letters* **61** (1), pp. 54–57 (1988).

[170] V. Tamma and S. Laibacher. “Multiboson Correlation Interferometry with Arbitrary Single-Photon Pure States”. *Physical Review Letters* **114** (24), p. 243601 (2015).

[171] C. M. Wilkes, X. Qiang, J. Wang, R. Santagati, S. Paesani, X. Zhou, D. A. B. Miller, G. D. Marshall, M. G. Thompson, and J. L. O’Brien. “60dB high-extinction auto-configured Mach–Zehnder interferometer”. *Optics Letters* **41** (22), pp. 5318–21 (2016).

[172] E. Magesan, D. Puzzuoli, C. E. Granade, and D. G. Cory. “Modeling quantum noise for efficient testing of fault-tolerant circuits”. *Physical Review A* **87** (1), p. 012324 (2013).

[173] W. Dür, M. Hein, J. I. Cirac, and H.-J. Briegel. “Standard forms of noisy quantum operations via depolarization”. *Physical Review A* **72** (5), p. 052326 (2005).

[174] Z. Aqua and B. Dayan. “Temporal quantum eraser: Fusion gates with distinguishable photons”. *Physical Review A* **110** (4), p. 043709 (2024).

[175] L. Stasi, G. Gras, R. Berrazouane, M. Perrenoud, H. Zbinden, and F. Bussières. “Fast high-efficiency photon-number-resolving parallel superconducting nanowire single-photon detector”. *Physical Review Applied* **19** (6), p. 064041 (2023).

[176] M.-D. Choi. “Completely positive linear maps on complex matrices”. *Linear algebra and its applications* **10** (3), pp. 285–90 (1975).

[177] A. Jamiołkowski. “Linear transformations which preserve trace and positive semidefiniteness of operators”. *Reports on mathematical physics* **3** (4), pp. 275–78 (1972).

[178] D. Schlingemann. “Cluster states, algorithms and graphs”. arXiv:quant-ph/0305170 (2003).

[179] M. Van den Nest, J. Dehaene, and B. De Moor. “Efficient algorithm to recognize the local Clifford equivalence of graph states”. *Physical Review A* **70** (3), p. 034302 (2004).

[180] P. Hilaire, L. Vidro, H. S. Eisenberg, and S. E. Economou. “Near-deterministic hybrid generation of arbitrary photonic graph states using a single quantum emitter and linear optics”. *Quantum* **7**, p. 992 (2023).

[181] F. Ewert, M. Bergmann, and P. Van Loock. “Ultrafast long-distance quantum communication with static linear optics”. *Physical Review Letters* **117** (21), p. 210501 (2016).

[182] S. C. Wein, T. G. de Brugière, L. Music, P. Senellart, B. Bourdoncle, and S. Mansfield. “Minimizing resource overhead in fusion-based quantum computation using hybrid spin-photon devices”. arXiv:2412.08611 (2024).

[183] M. L. Chan, T. J. Bell, L. A. Pettersson, S. X. Chen, P. Yard, A. S. Sørensen, and S. Paesani. “Tailoring Fusion-Based Photonic Quantum Computing Schemes to Quantum Emitters”. *PRX Quantum* **6** (2), p. 020304 (2025).

[184] M. L. Chan, A. A. Capatos, P. Lodahl, A. S. Sørensen, and S. Paesani. “Practical blueprint for low-depth photonic quantum computing with quantum dots”. arXiv:2507.16152 (2025).

[185] T. Dessertaine, B. Bourdoncle, A. Denys, G. de Gliniasty, P. C. d’Istria, G. Valentí-Rojas, S. Mansfield, and P. Hilaire. “Enhanced Fault-tolerance in Photonic Quantum Computing: Comparing the Honeycomb Floquet Code and the Surface Code in Tailored Architecture”. arXiv:2410.07065 (2025).

[186] P. Senellart, G. Solomon, and A. White. “High-performance semiconductor quantum-dot single-photon sources”. *Nature Nanotechnology* **12** (11), pp. 1026–39 (2017).

[187] N. Maring, A. Fyrillas, M. Pont, E. Ivanov, P. Stepanov, N. Margaria, W. Hease, A. Pishchagin, A. Lemaître, I. Sagnes, T. H. Au, S. Boissier, E. Bertasi, A. Baert, M. Valdivia, M. Billard, O. Acar, A. Brieussel, R. Mezher, S. C. Wein, A. Salavrakos, P. Sinnott, D. A. Fioretto, P.-E. Emeriau, N. Belabas, S. Mansfield, P. Senellart, J. Senellart, and N. Somaschi. “A versatile single-photon-based quantum computing platform”. *Nature Photonics* **18** (6), pp. 603–09 (2024).

[188] S. Bartolucci, T. Bell, H. Bombin, P. Birchall, J. Bulmer, C. Dawson, T. Farrelly, S. Gartenstein, M. Gimeno-Segovia, D. Litinski, Y. Liu, R. Knegjens, N. Nickerson, A. Olivo, M. Pant, A. Patil, S. Roberts, T. Rudolph, C. Sparrow, D. Tuckett, and A. Veitia. “Comparison of schemes for highly loss tolerant photonic fusion based quantum computing”. 2025. arXiv:2506.11975.

Appendix A

Outcomes of high-probability depth-2 n-GHZ fusions

In Chapter 3, it was shown that depth-2 n-GHZ fusions with the standard success probability of $1/2^{n-1}$ have entangling failure outcomes, listed in table 3.3. Other families of depth-2 n-GHZ fusions were introduced which have higher success probabilities for $n \geq 4$, but at the cost of some failure outcomes which project onto non-stabilizer states, whose total probabilities are shown in table 3.2. These fusions also have failure outcomes which project onto stabilizer states. For $n = 4$ and $n = 5$ these are only single-qubit measurements in the X or Z basis of each input qubit, but for $n \geq 6$ some of the failure outcomes are entangling projections on subsets of input qubits, like in the case of lower probability fusions. All stabilizer outcomes of the high-probabilty depth-2 n-GHZ fusions with $4 \leq n \leq 8$ are shown in table A.1, listed by success probability and family representative, and with outcomes of the same type grouped into a single entry, even if the projection may occur on different subsets of qubits in different cases. For example, if a 6-GHZ circuit can project qubits (1, 2, 3, 4, 5) or qubits (1, 2, 4, 5, 6) onto a GHZ state and measure the remaining one in Z, the total probability of these two outcomes is combined into a single entry "5-GHZ + $1 \times Z$". This is the same notation used in table 3.3.

n	Success probability	Family representative	Failure outcomes	Probability
4	9/32	((1, 3), (2, 4))	$4 \times Z$	1/8
5	9/64	((1, 3), (2, 4))	$4 \times Z$	$1/8$
6	9/128	((1, 2), (3, 5), (4, 6))	5-GHZ + X	9/128
			$1 \times X + 5 \times Z$	$11/32$
			$6 \times Z$	$1/32$
	23/256	((1, 3), (2, 5), (4, 6))	4-GHZ + BP	3/64
			4-GHZ + $2 \times Z$	3/16
			BP + $4 \times Z$	9/32
			$6 \times Z$	$3/32$
7	9/256	((1, 3), (2, 4), (5, 7))	5-GHZ + BP	9/256
			5-GHZ + $2 \times Z$	$9/128$
			BP + $5 \times Z$	$11/64$
			$7 \times Z$	$13/64$
		((1, 2), (3, 5), (4, 6))	6-GHZ + $1 \times X$	9/256
			$1 \times X + 6 \times Z$	49/128

			$7 \times Z$	3/32
		((1, 2), (3, 5), (4, 7))	6-GHZ + $1 \times X$	9/256
			BP + $1 \times X + 4 \times Z$	5/64
			$1 \times X + 6 \times Z$	27/64
			$7 \times Z$	3/32
	23/512	((1, 3), (2, 5), (4, 6))	5-GHZ + BP	1/64
			5-GHZ + $2 \times Z$	1/16
			4-GHZ + 3-GHZ	1/128
			4-GHZ + $3 \times Z$	1/16
			3-GHZ + $4 \times Z$	3/64
			BP + $5 \times Z$	15/64
			$7 \times Z$	33/128
8	9/512	((1, 2), (3, 8), (4, 6), (5, 7))	7-GHZ + $1 \times X$	11/128
			5-GHZ + 3-GHZ	9/512
			5-GHZ + BP + $1 \times X$	9/512
			3-GHZ + $5 \times Z$	11/128
			BP + $1 \times X + 5 \times Z$	11/128
			$1 \times X + 7 \times Z$	25/128
		((1, 2), (3, 4), (5, 7), (6, 8))	7-GHZ + $1 \times X$	9/256
			6-GHZ + $2 \times X$	9/512
			$2 \times X + 6 \times Z$	19/64
			$1 \times X + 7 \times Z$	11/64
			$8 \times Z$	1/128
		((1, 2), (3, 5), (4, 8), (6, 7))	7-GHZ + $1 \times X$	9/256
			6-GHZ + $2 \times X$	9/512
			BP + $2 \times X + 4 \times Z$	5/64
			$2 \times X + 6 \times Z$	43/128
			$1 \times X + 7 \times Z$	11/64
			$8 \times Z$	1/128
	23/1024	((1, 2), (3, 5), (4, 7), (6, 8))	7-GHZ + $1 \times X$	23/1024
			6-GHZ + BP	1/128
			6-GHZ + $2 \times Z$	1/32
			5-GHZ + BP + $1 \times X$	1/128
			5-GHZ + $1 \times X + 2 \times Z$	1/32
			$2\times$ 4-GHZ	1/256
			4-GHZ + 3-GHZ + $1 \times X$	1/256
			4-GHZ + $1 \times X + 3 \times Z$	3/64
			4-GHZ + $4 \times Z$	5/128
			3-GHZ + $1 \times X + 4 \times Z$	3/128
			BP + $1 \times X + 5 \times Z$	3/16
			BP + $6 \times Z$	3/64
			$1 \times X + 7 \times Z$	15/64

			$8 \times Z$	3/128
	55/2048	((1, 3), (2, 6), (4, 8), (5, 7))	6-GHZ + BP	1/64
			$2\times$ 4-GHZ	1/128
			4-GHZ + $2\times$ BP	1/64
			4-GHZ + BP + $2 \times Z$	1/8
			4-GHZ + $4 \times Z$	1/8
			$2\times$ BP + $4 \times Z$	5/32
			BP + $6 \times Z$	15/64
			$8 \times Z$	7/128
	19/512	((1, 3), (2, 5), (4, 7), (6, 8))	6-GHZ + BP	5/512
			6-GHZ + $2 \times Z$	5/128
			$2\times$ BP + $4 \times Z$	1/32
			BP + $6 \times Z$	11/128
			$8 \times Z$	9/32
	81/2048	((1, 3), (2, 4), (5, 7), (6, 8))	$8 \times Z$	65/256

Table A.1: Stabilizer measurement outcomes of depth-2 n-GHZ fusion families with success probability higher than $1/2^{n-1}$ for $4 \leq n \leq 8$. For $n \geq 6$ some of the failure outcomes project subsets of qubits onto entangled states.

Appendix B

Pauli error channels of standard 4- and 5-GHZ fusions

(a) 4-GHZ fusion **(b)** 5-GHZ fusion

Figure B.1: ZX-diagrams of the standard 4- and 5-GHZ fusion circuits used to extract Pauli error models. These are the most balanced of the possible topologically-inequivalent choices for these fusion sizes.

The error propagation rules described in Chapter 4 can be applied to standard n-GHZ fusions of any size. Here we explicitly show the calculation of the Pauli error channels derived in this way for the 4- and 5-GHZ fusion circuits whose error rates are included in the plots Chapter 4. Both these fusions admit topologically inequivalent ZX-diagram representations, and we choose to consider their most balanced, tree-like versions, which correspond to the linear optical circuit with the lowest average number of crossings per mode. For an n-GHZ fusion, there are n topologically equivalent diagrams of this kind, obtained by applying cyclic permutations to the input qubits according to the symmetry of the underlying linear optical circuit. A specific representative of each size, shown in fig. B.1, is used to extract the Pauli error channels associated with different input error configurations, and then these are symmetrised to obtain an average channel across all n diagrams.

As explained in Chapter 4, the effect of distinguishability on any standard n-GHZ fusion can be described as a single Z dephasing error channel on the faulty input qubit, so that considering all input error configurations, the total Pauli error channel is given by eq. 4.39, repeated here for convenience:

$$\rho_{\text{in}} \to p_d \sum_{i=1}^{n} \Delta_{Z_i}(\rho_{\text{in}}) = p_d \left(\frac{n}{2}\rho_{\text{in}} + \frac{1}{2}\sum_{i=1}^{n} Z_i \rho_{\text{in}} Z_i \right). \tag{B.1}$$

Similarly, the combination of multiphoton contamination and loss on the two inputs of one T1 fusion spider gives rise to the same Pauli error channel regardless of the size of the overall fusion, as in

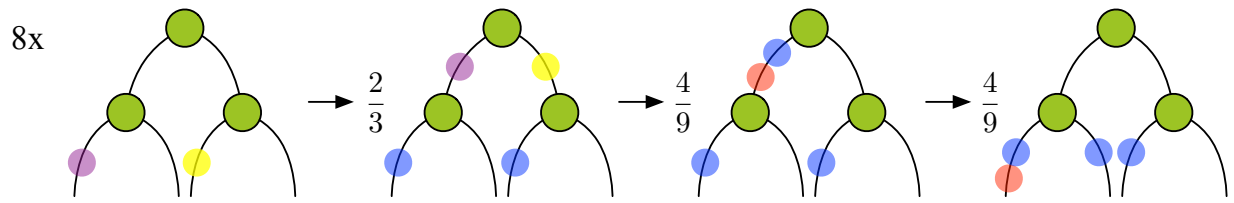

Figure B.2: Derivation of the Pauli error channel associated with multiphoton contamination and loss errors on qubits entering different constituent T1 fusions through the standard 4-GHZ fusion ZX-diagram. There are eight configurations of this kind, all of which lead to equivalent channels up to permutations of qubits.

eq. 4.40, repeated here for convenience:

$$\begin{aligned}\frac{2}{3}p_m p_l\left(\Delta_{X_i}\circ\Delta_{X_{i+1}}\circ\Delta_Z\right) =& \frac{1}{3}p_m p_l\left(\Delta_{X_i}+\Delta_{X_{i+1}}+\Delta_{X_iX_{i+1}}-I\right)\circ\Delta_Z = \\ =& \frac{1}{6}p_m p_l\left(\Delta_{X_i}+\Delta_{Y_i}+\Delta_{X_{i+1}}+\Delta_{Y_{i+1}}+\Delta_{X_iX_{i+1}}+\frac{1}{2}\left(\Delta_{Y_iX_{i+1}}+\Delta_{X_iY_{i+1}}\right)+\Delta_Z-5I\right).\end{aligned} \tag{B.2}$$

From the structure of the diagrams in fig. B.1, it can be seen that the 4-GHZ and the 5-GHZ fusion each admit four configurations of this kind, consisting of (multi, loss) or (loss, multi) on the first two or last two qubits. The only input error configurations which need to be propagated through the diagrams are then combinations of multiphoton contamination and loss on qubits which are not input into the same T1 fusion. The derivation of the channels associated with these configurations in the 4-GHZ and 5-GHZ fusions is described below.

B.1 4-GHZ fusion

In the case of the 4-GHZ fusion, there are eight equivalent configurations of multiphoton contamination and loss entering different T1 fusions, all of which can be associated with identical Pauli error channels. The graphical calculation is shown in fig. B.2 for one such configuration, (multi, ideal, loss, ideal), from which the Pauli error channel is found to be $\frac{4}{9}p_m p_l\left(\Delta_{X_1}\circ\Delta_{X_2}\circ\Delta_{X_3}\circ\Delta_Z\right)$. Symmetrising over all cyclic permutations of input qubits to reflect the four topologically-equivalent ZX-diagrams, and accounting for all eight configurations gives the total channel:

$$\begin{aligned}\rho_{\text{in}} \to& \frac{4}{9}p_m p_l \times 8 \times \frac{1}{4}\sum_{i=1}^{4}\left(\Delta_{X_i}\circ\Delta_{X_{i+1}}\circ\Delta_{X_{i+2}}\circ\Delta_Z\right)(\rho_{\text{in}}) = \\ =& \frac{2}{9}p_m p_l\rho_{\text{in}} + \frac{2}{9}p_m p_l\sum_{i=1}^{4}\left(X_i\rho_{\text{in}}X_i + Y_i\rho_{\text{in}}Y_i + \frac{1}{4}Z_i\rho_{\text{in}}Z_i\right) + \\ & \frac{1}{9}p_m p_l\sum_{i=1}^{4}\sum_{j>i}\left[X_iX_j\rho_{\text{in}}X_iX_j + \frac{1}{2}\left(X_iY_j\rho_{\text{in}}X_iY_j + Y_iX_j\rho_{\text{in}}Y_iX_j\right)\right],\end{aligned} \tag{B.3}$$

where the subscripts $i+k$ are used to indicate $(i+k)$ mod 4 for simplicity, and care has been taken to reduce the weight of Pauli operators and to symmetrise across all four qubits when composing the dephasing channels in line 1.

The Pauli channel due to the four configurations with multiphoton contamination and loss on the two inputs of a T1 spider is given by eq. B.2. When averaging over the four equivalent ZX-diagrams this becomes:

$$\begin{aligned}\rho_{\text{in}} \to & \frac{2}{3} p_m p_l \times 4 \times \frac{1}{4} \sum_{i=1}^{4} \left(\Delta_{X_i} \circ \Delta_{X_{i+1}} \circ \Delta_Z\right)(\rho_{\text{in}}) = \\ = & \frac{1}{3} p_m p_l \rho_{\text{in}} + \frac{1}{6} p_m p_l \sum_{i=1}^{4} \left(X_i \rho_{\text{in}} X_i + Y_i \rho_{\text{in}} Y_i + \frac{1}{2} Z_i \rho_{\text{in}} Z_i\right) + \\ & \frac{1}{12} p_m p_l \sum_{i=1}^{4} \left[X_i X_{i+1} \rho_{\text{in}} X_i X_{i+1} + \frac{1}{2}\left(X_i Y_{i+1} \rho_{\text{in}} X_i Y_{i+1} + Y_i X_{i+1} \rho_{\text{in}} Y_i X_{i+1}\right)\right].\end{aligned} \tag{B.4}$$

The total Pauli error channel for the 4-GHZ fusion is obtained by combining the two channels in eqs. B.3 and B.4 with the channel due to distinguishability, given by eq. B.1 with $n = 4$. This can then be turned into a probability distribution of flips of the stabilizer generators of the success outcome of the fusion, as done in Chapter 4 for the T2 and 3-GHZ fusion, resulting in the following:

$$\begin{aligned}
\Pr(-XXXX, +ZZII, +ZIZI, +ZIIZ|\text{success}) &= 2p_d + \frac{5}{9} p_m p_l + p_m p_l + 4p_Z, \\
\Pr(+XXXX, -ZZII, +ZIZI, +ZIIZ|\text{success}) &= \frac{7}{18} p_m p_l + \frac{1}{4} p_m p_l + p_X, \\
\Pr(+XXXX, +ZZII, -ZIZI, +ZIIZ|\text{success}) &= \frac{7}{18} p_m p_l + \frac{1}{4} p_m p_l + p_X, \\
\Pr(+XXXX, +ZZII, +ZIZI, -ZIIZ|\text{success}) &= \frac{7}{18} p_m p_l + \frac{1}{4} p_m p_l + p_X, \\
\Pr(+XXXX, -ZZII, -ZIZI, +ZIIZ|\text{success}) &= \frac{2}{9} p_m p_l, \\
\Pr(+XXXX, -ZZII, +ZIZI, -ZIIZ|\text{success}) &= \frac{2}{9} p_m p_l, \\
\Pr(+XXXX, +ZZII, -ZIZI, -ZIIZ|\text{success}) &= \frac{5}{9} p_m p_l, \\
\Pr(+XXXX, -ZZII, -ZIZI, -ZIIZ|\text{success}) &= \frac{7}{18} p_m p_l + \frac{1}{4} p_m p_l + p_X, \\
\Pr(-XXXX, -ZZII, +ZIZI, +ZIIZ|\text{success}) &= \frac{7}{18} p_m p_l + \frac{1}{4} p_m p_l + p_Y, \\
\Pr(-XXXX, +ZZII, -ZIZI, +ZIIZ|\text{success}) &= \frac{7}{18} p_m p_l + \frac{1}{4} p_m p_l + p_Y, \\
\Pr(-XXXX, +ZZII, +ZIZI, -ZIIZ|\text{success}) &= \frac{7}{18} p_m p_l + \frac{1}{4} p_m p_l + p_Y, \\
\Pr(-XXXX, -ZZII, -ZIZI, +ZIIZ|\text{success}) &= \frac{2}{9} p_m p_l, \\
\Pr(-XXXX, -ZZII, +ZIZI, -ZIIZ|\text{success}) &= \frac{2}{9} p_m p_l, \\
\Pr(-XXXX, +ZZII, -ZIZI, -ZIIZ|\text{success}) &= \frac{5}{9} p_m p_l, \\
\Pr(-XXXX, -ZZII, -ZIZI, -ZIIZ|\text{success}) &= \frac{7}{18} p_m p_l + \frac{1}{4} p_m p_l + p_Y.
\end{aligned} \tag{B.5}$$

This is the probability distribution used to evaluate the Pauli error metrics for the 4-GHZ fusion shown

in the plots in Chapter 4.

B.2 5-GHZ fusion

Due to its less symmetric structure, the 5-GHZ fusion has a larger number of inequivalent multiphoton contamination and loss input configurations, which give rise to different Pauli error channels. Five groups of configurations can be identified:

$$\begin{aligned}
&\text{(a)} \quad \{(\text{multi}_1, \text{loss}_3), (\text{multi}_2, \text{loss}_3)\} \\
&\text{(b)} \quad \{(\text{loss}_1, \text{multi}_3), (\text{loss}_2, \text{multi}_3)\} \\
&\text{(c)} \quad \{(\text{multi}_1, \text{loss}_4), (\text{multi}_2, \text{loss}_4), (\text{multi}_1, \text{loss}_5), (\text{multi}_2, \text{loss}_5)\} \\
&\text{(d)} \quad \{(\text{loss}_1, \text{multi}_4), (\text{loss}_2, \text{multi}_4), (\text{loss}_1, \text{multi}_5), (\text{loss}_2, \text{multi}_5)\} \\
&\text{(e)} \quad \{(\text{loss}_3, \text{multi}_4), (\text{loss}_3, \text{multi}_5), (\text{multi}_3, \text{loss}_4), (\text{multi}_3, \text{loss}_5)\},
\end{aligned} \tag{B.6}$$

where in the more concise notation $(\text{multi}_i, \text{loss}_j)$ qubits i and j are affected by the respective errors and all other qubits are assumed to be ideal. The propagation of one representative from each of these groups through the 5-GHZ fusion ZX-diagram is shown in fig. B.3.

From these diagrams, it can be seen that different groups of configurations lead to similar Pauli error channels, up to different success probability reduction factors. In particular, groups (a), (b) and (e) are all associated with channels of the form $\Delta_{X_i} \circ \Delta_{X_{i+1}} \circ \Delta_{X_{i+2}} \circ \Delta_Z$, and similarly, for groups (c) and (d) $\Delta_{X_i} \circ \Delta_{X_{i+1}} \circ \Delta_{X_{i+2}} \circ \Delta_{X_{i+3}} \circ \Delta_Z$. Combining these with the appropriate factors and averaging over the five equivalent ZX diagrams gives the total channel:

$$\begin{aligned}
\rho_{\text{in}} \to &\left(6 \times \frac{4}{9} + 2 \times \frac{2}{3}\right) p_m p_l \times \frac{1}{5} \sum_{i=1}^{5} \left(\Delta_{X_i} \circ \Delta_{X_{i+1}} \circ \Delta_{X_{i+2}} \circ \Delta_Z\right)(\rho_{\text{in}}) + \\
&\left(4 \times \frac{8}{27} + 4 \times \frac{4}{9}\right) p_m p_l \times \frac{1}{5} \sum_{i=1}^{5} \left(\Delta_{X_i} \circ \Delta_{X_{i+1}} \circ \Delta_{X_{i+2}} \circ \Delta_{X_{i+3}} \circ \Delta_Z\right)(\rho_{\text{in}}) = \\
= &\frac{37}{108} p_m p_l \rho_{\text{in}} + \frac{131}{540} p_m p_l \sum_{i=1}^{5} \left(X_i \rho_{\text{in}} X_i + Y_i \rho_{\text{in}} Y_i + \frac{37}{131} Z_i \rho_{\text{in}} Z_i\right) + \\
&p_m p_l \sum_{i=1}^{5} \left[\frac{131}{540} \left(X_i X_{i+1} \rho_{\text{in}} X_i X_{i+1} + \frac{1}{2}\left(X_i Y_{i+1} \rho_{\text{in}} X_i Y_{i+1} + Y_i X_{i+1} \rho_{\text{in}} Y_i X_{i+1}\right)\right) + \right. \\
&\left. \frac{77}{1080} \sum_{j=2}^{3} \left(X_i X_{i+j} \rho_{\text{in}} X_i X_{i+j} + \frac{1}{2}\left(X_i Y_{i+j} \rho_{\text{in}} X_i Y_{i+j} + Y_i X_{i+j} \rho_{\text{in}} Y_i X_{i+j}\right)\right)\right],
\end{aligned} \tag{B.7}$$

where the subscripts $i+k$ stand for $(i+k)$ mod 5, and the weight of the Pauli error operators has been reduced to the minimum and they have been symmetrised across all five qubits using the $X^{\otimes 5}$ stabilizer operator. The remaining four configurations with multiphoton contamination and loss at the

input of the same T1 spider give the channel:

$$\begin{aligned}\rho_{\text{in}} \to & \frac{2}{3}p_m p_l \times 4 \times \frac{1}{5}\sum_{i=1}^{4}\left(\Delta_{X_i}\circ\Delta_{X_{i+1}}\circ\Delta_Z\right)(\rho_{\text{in}}) = \\ = & \frac{1}{3}p_m p_l \rho_{\text{in}} + \frac{2}{15}p_m p_l \sum_{i=1}^{4}\left(X_i\rho_{\text{in}}X_i + Y_i\rho_{\text{in}}Y_i + \frac{1}{2}Z_i\rho_{\text{in}}Z_i\right) + \\ & \frac{1}{15}p_m p_l \sum_{i=1}^{4}\left[X_iX_{i+1}\rho_{\text{in}}X_iX_{i+1} + \frac{1}{2}\left(X_iY_{i+1}\rho_{\text{in}}X_iY_{i+1} + Y_iX_{i+1}\rho_{\text{in}}Y_iX_{i+1}\right)\right].\end{aligned} \tag{B.8}$$

The total channel consists of the sum of eqs. B.6 and B.8 with the Pauli error channel due to input distinguishability, given by eq. B.1 with $n = 5$. This can be expressed as the following probability distribution of flips of the stabilizers of the success outcome of the fusion:

$$\begin{aligned}
\Pr(-XXXXX,+ZZIII,+ZIZII,+ZIIZI,+ZIIIZ|\text{success}) &= \frac{5}{2}p_d + \frac{73}{108}p_m p_l + \frac{5}{4}p_m p_l + 5p_Z,\\
\Pr(+XXXXX,-ZZIII,+ZIZII,+ZIIZI,+ZIIIZ|\text{success}) &= \frac{203}{540}p_m p_l + \frac{1}{4}p_m p_l + p_X,\\
\Pr(+XXXXX,+ZZIII,-ZIZII,+ZIIZI,+ZIIIZ|\text{success}) &= \frac{203}{540}p_m p_l + \frac{1}{4}p_m p_l + p_X,\\
\Pr(+XXXXX,+ZZIII,+ZIZII,-ZIIZI,+ZIIIZ|\text{success}) &= \frac{203}{540}p_m p_l + \frac{1}{4}p_m p_l + p_X,\\
\Pr(+XXXXX,+ZZIII,+ZIZII,+ZIIZI,-ZIIIZ|\text{success}) &= \frac{203}{540}p_m p_l + \frac{1}{4}p_m p_l + p_X,\\
\Pr(+XXXXX,-ZZIII,-ZIZII,+ZIIZI,+ZIIIZ|\text{success}) &= \frac{167}{540}p_m p_l,\\
\Pr(+XXXXX,-ZZIII,+ZIZII,-ZIIZI,+ZIIIZ|\text{success}) &= \frac{77}{540}p_m p_l,\\
\Pr(+XXXXX,-ZZIII,+ZIZII,+ZIIZI,-ZIIIZ|\text{success}) &= \frac{77}{540}p_m p_l,\\
\Pr(+XXXXX,+ZZIII,-ZIZII,-ZIIZI,+ZIIIZ|\text{success}) &= \frac{167}{540}p_m p_l,\\
\Pr(+XXXXX,+ZZIII,-ZIZII,+ZIIZI,-ZIIIZ|\text{success}) &= \frac{77}{540}p_m p_l,\\
\Pr(+XXXXX,+ZZIII,+ZIZII,-ZIIZI,-ZIIIZ|\text{success}) &= \frac{167}{540}p_m p_l,\\
\Pr(+XXXXX,-ZZIII,-ZIZII,-ZIIZI,+ZIIIZ|\text{success}) &= \frac{77}{540}p_m p_l,\\
\Pr(+XXXXX,-ZZIII,+ZIZII,-ZIIZI,-ZIIIZ|\text{success}) &= \frac{77}{540}p_m p_l,\\
\Pr(+XXXXX,+ZZIII,-ZIZII,-ZIIZI,-ZIIIZ|\text{success}) &= \frac{167}{540}p_m p_l,\\
\Pr(+XXXXX,-ZZIII,-ZIZII,-ZIIZI,-ZIIIZ|\text{success}) &= \frac{203}{540}p_m p_l + \frac{1}{4}p_m p_l + p_X,\\
\Pr(-XXXXX,+ZZIII,+ZIZII,+ZIIZI,+ZIIIZ|\text{success}) &= \frac{203}{540}p_m p_l + \frac{1}{4}p_m p_l + p_Y,\\
\Pr(-XXXXX,+ZZIII,-ZIZII,+ZIIZI,+ZIIIZ|\text{success}) &= \frac{203}{540}p_m p_l + \frac{1}{4}p_m p_l + p_Y,
\end{aligned} \tag{B.9}$$

$$\Pr(-XXXXX,+ZZIII,+ZIZII,-ZIIZI,+ZIIIZ|\text{success}) = \frac{203}{540}p_m p_l + \frac{1}{4}p_m p_l + p_Y,$$
$$\Pr(-XXXXX,+ZZIII,+ZIZII,+ZIIZI,-ZIIIZ|\text{success}) = \frac{203}{540}p_m p_l + \frac{1}{4}p_m p_l + p_Y,$$
$$\Pr(-XXXXX,-ZZIII,-ZIZII,+ZIIZI,+ZIIIZ|\text{success}) = \frac{167}{540}p_m p_l,$$
$$\Pr(-XXXXX,-ZZIII,+ZIZII,-ZIIZI,+ZIIIZ|\text{success}) = \frac{77}{540}p_m p_l,$$
$$\Pr(-XXXXX,-ZZIII,+ZIZII,+ZIIZI,-ZIIIZ|\text{success}) = \frac{77}{540}p_m p_l,$$
$$\Pr(-XXXXX,+ZZIII,-ZIZII,-ZIIZI,+ZIIIZ|\text{success}) = \frac{167}{540}p_m p_l,$$
$$\Pr(-XXXXX,+ZZIII,-ZIZII,+ZIIZI,+ZIIIZ|\text{success}) = \frac{77}{540}p_m p_l,$$
$$\Pr(-XXXXX,+ZZIII,+ZIZII,+ZIIZI,-ZIIIZ|\text{success}) = \frac{167}{540}p_m p_l,$$
$$\Pr(-XXXXX,+ZZIII,+ZIZII,-ZIIZI,-ZIIIZ|\text{success}) = \frac{167}{540}p_m p_l,$$
$$\Pr(-XXXXX,-ZZIII,-ZIZII,-ZIIZI,+ZIIIZ|\text{success}) = \frac{77}{540}p_m p_l,$$
$$\Pr(-XXXXX,-ZZIII,+ZIZII,-ZIIZI,-ZIIIZ|\text{success}) = \frac{77}{540}p_m p_l,$$
$$\Pr(-XXXXX,+ZZIII,-ZIZII,-ZIIZI,-ZIIIZ|\text{success}) = \frac{167}{540}p_m p_l,$$
$$\Pr(-XXXXX,-ZZIII,-ZIZII,-ZIIZI,-ZIIIZ|\text{success}) = \frac{203}{540}p_m p_l + \frac{1}{4}p_m p_l + p_Y.$$

This is the probability distribution used to evaluate the Pauli error metrics for the 5-GHZ fusion shown in the plots in Chapter 4.

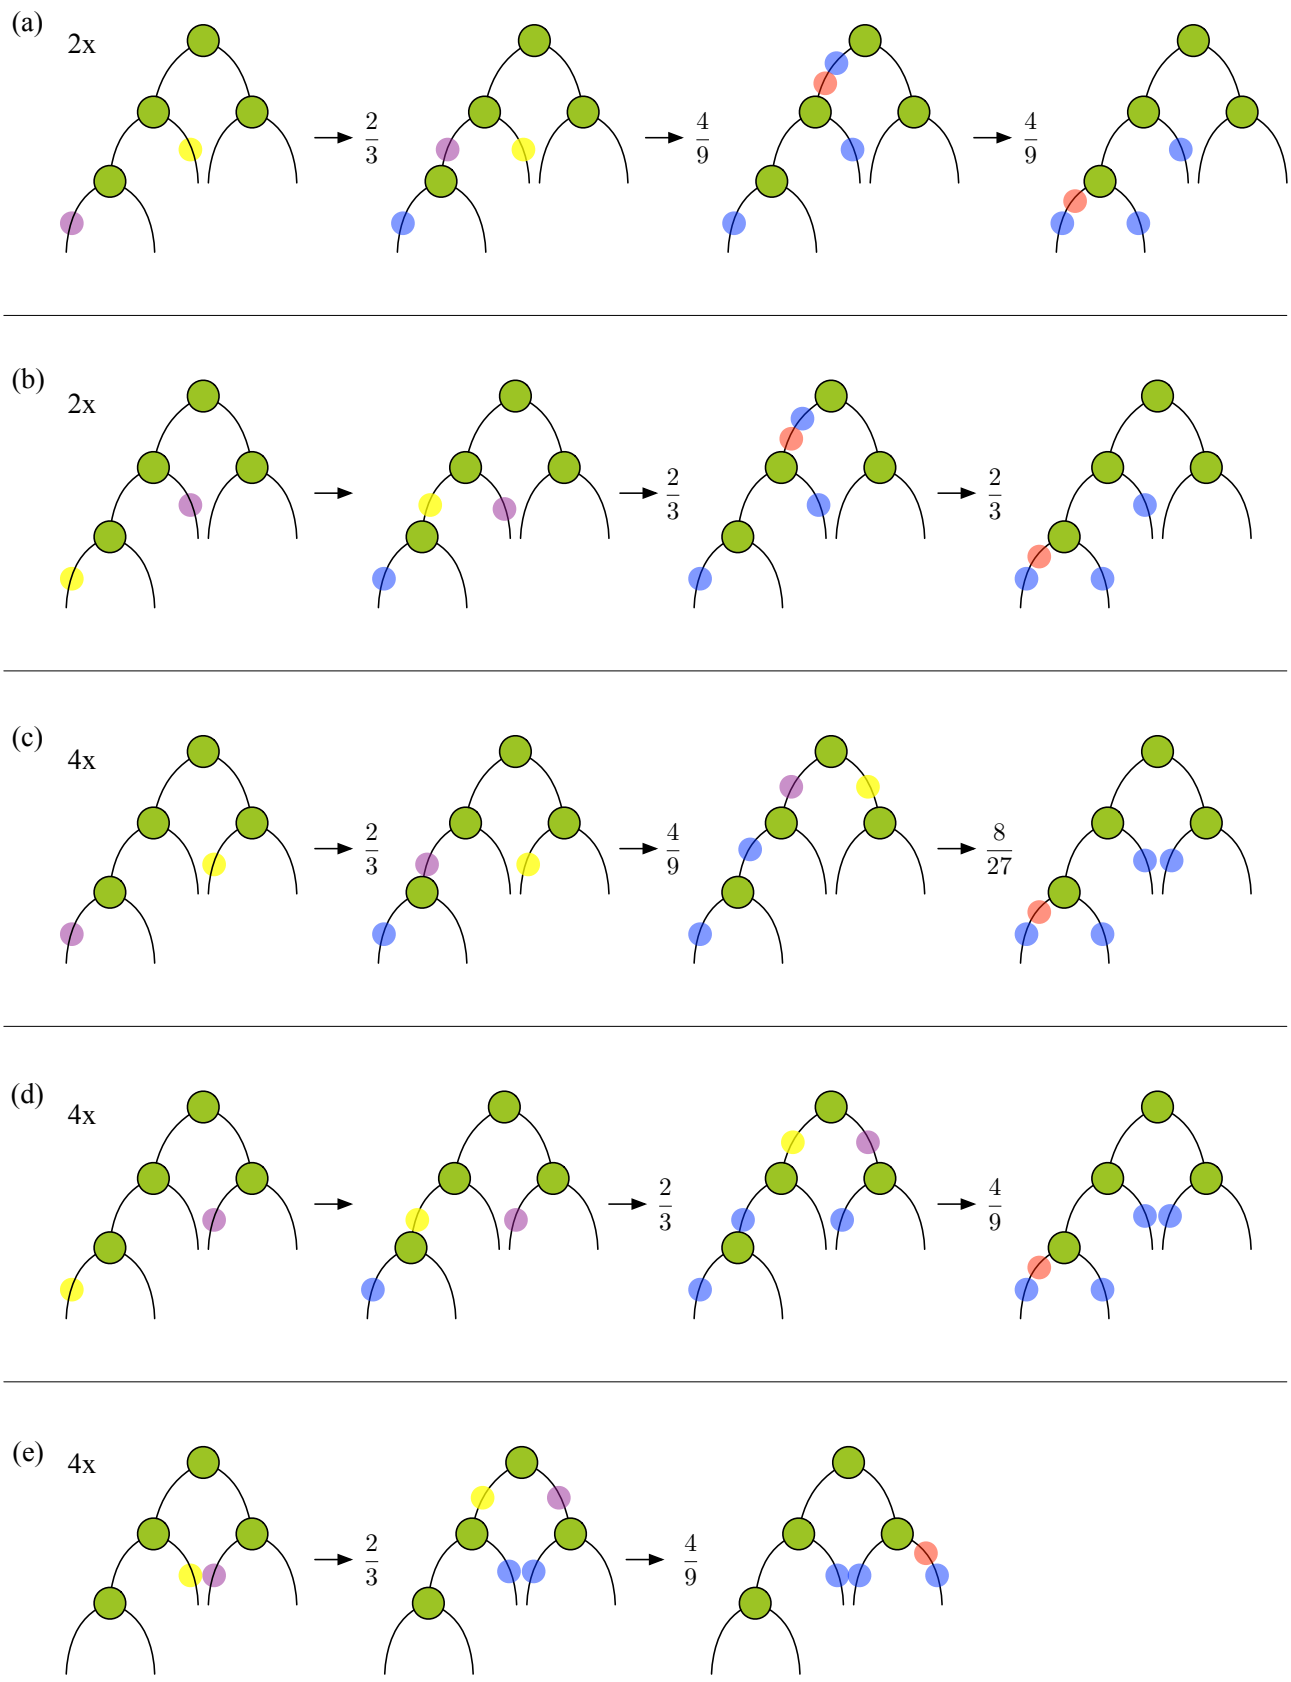

Figure B.3: Graphical Pauli error channel extraction for one representative input multiphoton contamination and loss configuration for each of the five inequivalent groups for the standard 5-GHZ fusion. The four configurations in which these two errors occur on the inputs of the same T1 fusion are not captured here. Up to differences in the success probability reduction factors, groups (a), (b) and (e) share channels of the same form, and similarly for groups (c) and (d).

Appendix C

Pauli error rates in different physical parameter regimes

The error modelling frameworks described in Chapters 4 and 5 can be used to analyse the performance of multi-way fusion circuits in different physical error parameter regimes. The results presented in the body of this thesis focus on one specific choice of parameter values, repeated below, considered to be representative of the noise profile of an FBQC architecture with heralded single photon sources, as described in Chapter 2:

$$\begin{aligned} p_d &= 0.1\%, \\ p_m &= 0.3\%, \\ p_l &= 15\%, \\ p_X &= 0.01\% \\ p_Y &= 0.01\% \\ p_Z &= 0.03\%, \end{aligned} \tag{4.41}$$

where p_d is the photon distinguishability rate, p_m the multiphoton contamination probability, p_l the loss rate, and p_X, p_Y, p_Z the probabilities of a qubit having suffered each of the three Pauli errors before entering the fusion.

Other platforms can be used to implement fusion-based quantum computation. For example, architectures of this kind have been proposed which use photons generated by quantum emitters [182–185]. Each technology has a unique physical noise profile, so it is interesting to observe how the results presented in this thesis would change in different parameter regimes. The error models extracted for the fusion circuits in Chapter 5 make it easy to carry out such an analysis, which only requires evaluating the model functions with different input values.

As an example, consider the following set of physical error parameters:

$$\begin{aligned} p_d &= 1\%, \\ p_m &= 0.02\%, \\ p_l &= 15\%, \\ p_X &= 0.01\% \\ p_Y &= 0.01\% \\ p_Z &= 0.03\%. \end{aligned} \tag{C.1}$$

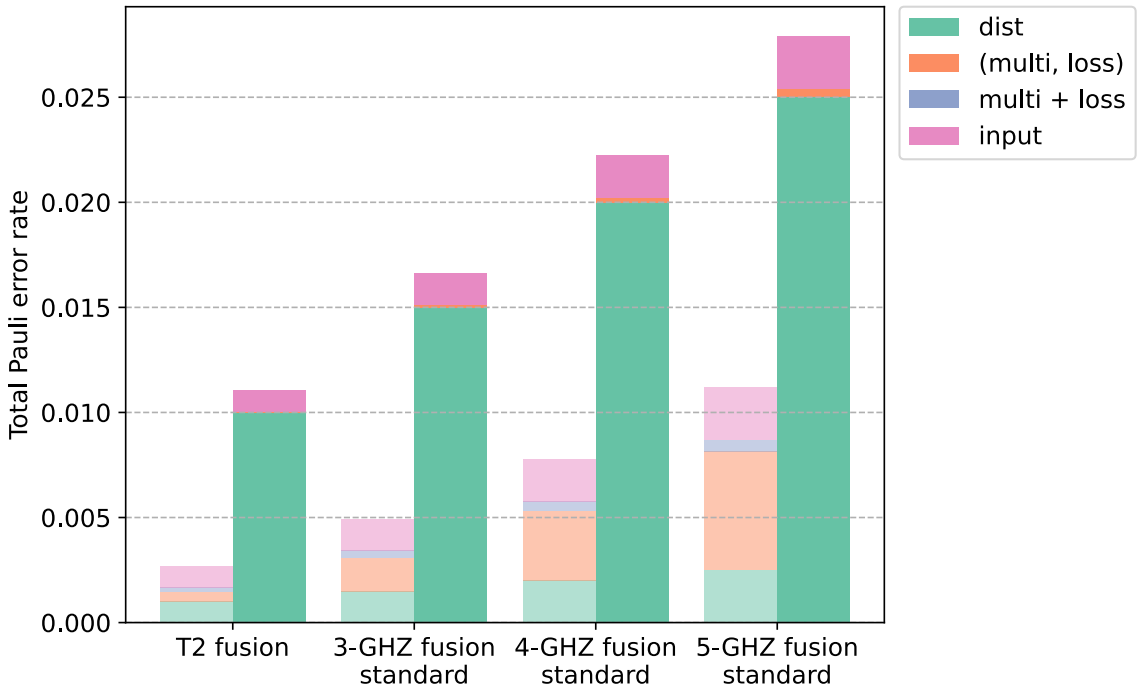

Figure C.1: Total Pauli error rate for standard n-way fusions. Darker bars show the breakdown of the error rate in the parameter regime of interest for photons generated by quantum emitters, assuming the physical parameter values in eq. C.1. Lighter bars are the same as in fig. 5.1, obtained using the physical error parameters in eq. 4.41 for photons from heralded single photon sources. The much higher p_d value in the former set of parameters results in errors from distinguishability dominating the total error rate at all fusion sizes, while the lower p_m makes the contribution from combinations of multiphoton contamination and loss negligible in comparison.

We take these to be indicative of the type of noise which may be expected in an architecture of the kind described in this thesis, but with single photons generated by quantum emitters. The promise of these sources is to significantly reduce the overhead of linear optical quantum computing by providing deterministic single photon emission [186]. The photons they produce are expected to have higher distinguishability rates but lower multiphoton contamination probabilities than those output by heralded single photon sources. The values of p_d and p_m chosen above are current estimates of achievable performance based on experimental data and proposed schemes [182]. Collection efficiency is a challenge for these photons, with predictions of near-term loss targets close to $\sim 25\%$ [187]. However, here we choose to keep p_l fixed at the value used for heralded single photon sources, based on the expectation that a loss rate closer to that value will be necessary for a realistic FBQC architecture [188]. We also keep the input Pauli error rates unchanged between the two sets of parameters, due to the lack of available data on the performance of linear optical seed state generation with photons from quantum emitters. Their impact on the Pauli error rates is simply additive, so the specific distribution of probabilities across the three Pauli errors does not change the overall error rates.

Figures C.1 and C.2 show the total Pauli error rate and marginal per-qubit Pauli error rate for the success outcome of standard n-way fusions with $2 \leq n \leq 5$ obtained using the fusion channels extracted in Chapter 5 and the parameters in eq. C.1. The total Pauli error rates from fig. 5.1, obtained with the parameter values in eq. 4.41, are also shown in lighter colours for reference in fig. C.1. From this, it

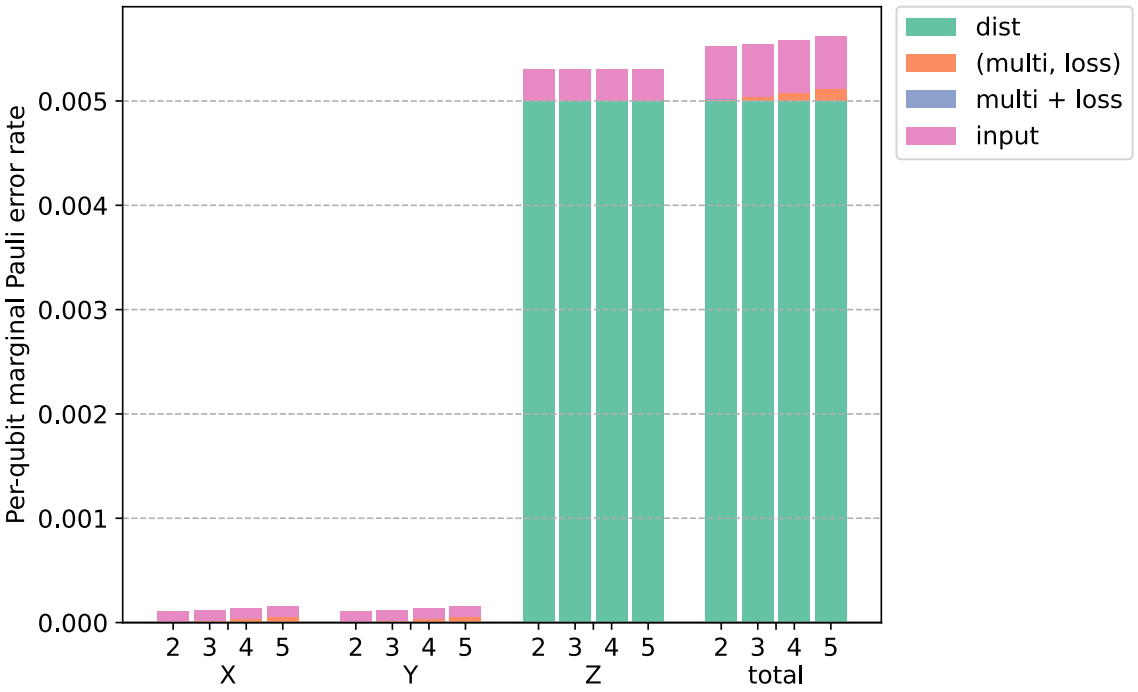

Figure C.2: Marginal Pauli error rate for standard n-way fusions in the parameter regime of interest for photons from quantum emitters, given by the values in eq. C.1. Similarly to the total Pauli error rate in the previous plot, distinguishability dominates the error probability in this case. Since this type of input error translates into single-qubit Z errors regardless of the fusion size, in this scenario correlated errors play a smaller role and the difference in the total per-qubit error rate is smaller than observed in the other parameter regime.

is clear that the contribution from input distinguishability dominates the total error rate in this regime, where p_d is $10\times$ larger than previously considered. Combinations of multiphoton contamination and loss, on the other hand, play a negligible role in comparison due to the much smaller value of p_m. While the total Pauli error rate is significantly higher when considering the example set of physical errors for photons from quantum emitters, this increase would trade off against a large reduction in the resource overhead of the architecture, making the comparison between the two platforms more involved and outside the scope of this thesis. The per-qubit marginal error rates in fig. C.2 confirm the large impact of distinguishability in the parameter regime in eq. C.1. In Chapter 4, it was shown that this kind of error is equivalent to a single-qubit Pauli Z at the input. Therefore, correlated errors play a smaller role in this scenario, resulting in a smaller gap in the per-qubit Pauli error rate between fusions of different sizes, compared to that seen in fig. 5.2.

The adjusted total Pauli error rate was introduced in Chapter 6 to compare the performance of fusions with different numbers of inputs within an RSG, taking into account the total number of fusions performed. Fig. C.3 shows this metric for the standard n-way fusions seen above, alongside the bars for the heralded single photon source parameter regime from fig. 6.3 in lighter colours. The data points to different conclusions in the two cases. As seen previously, when the physical error parameters in eq. 4.41 are used, the different scalings of errors from distinguishability and combinations of multiphoton contamination and loss with fusion size result in the adjusted total Pauli rate decreasing with n at first, reaching a minimum for $n = 4$, and then increasing again for $n = 5$. On the other hand, the set

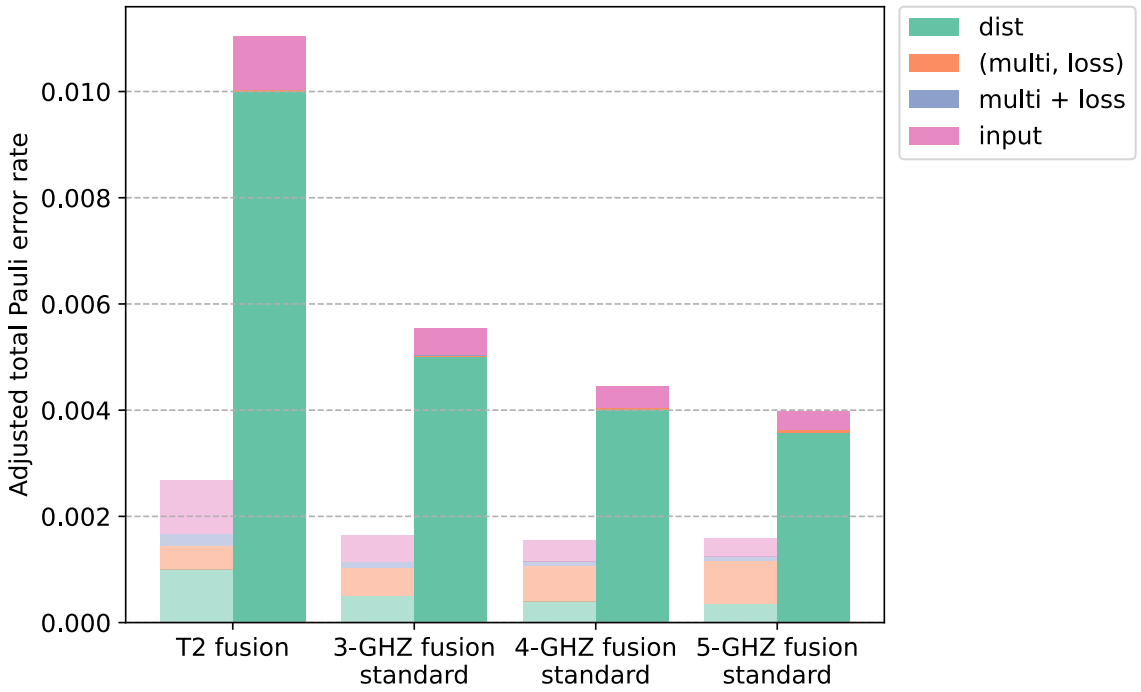

Figure C.3: Adjusted total Pauli error rate for standard n-way fusions in the parameter regime of interest for photons from quantum emitters, given by the values in eq. C.1 (darker bars). The trend with increasing fusion size in this regime differs from the case of photons from heralded single photon sources, given by the parameters in eq. 4.41 (lighter bars). In that scenario, the different scalings of distinguishability and combinations of multiphoton contamination and loss lead to the adjusted total Pauli error rate reaching a minimum at $n = 4$ and then starting to grow again for $n = 5$. On the other hand, the height of the dark bars keeps decreasing as n increases, fully dominated by the reduced impact of distinguishability.

of parameters in eq. C.1 leads to the value of the metric decreasing as the fusion size increases all the way to $n = 5$. This is due to the dominant role played by distinguishability errors in this regime, which favours larger fusions with negligible penalty due to errors caused by multiphoton contamination and loss. This plot does not capture the impact of choosing larger fusions on the resource overhead of RSG, which is crucial to consider for the design of a scheme which can be implemented in practice. The analysis of the impact of different fusion circuits on RSG footprint described in Chapter 6 remains valid in this context.

As evidenced by the data in fig. C.3, it is essential to take into account the physical noise profile when making design decisions for an FBQC architecture. To illustrate this point, here we have chosen to focus on two distinct error parameter regimes representative of two single photon source technologies. Fig. C.4 shows the results of a sweep of the distinguishability rate, p_d, and the probability of a combination of multiphoton contamination and loss, $p_m p_l$. For each set of parameters, the optimal fusion size is shown, defined as the n for which the adjusted total Pauli rate of the standard n-GHZ fusion is lowest, with $2 \leq n \leq 5$ as above. The black star identifies the heralded single photon source parameters in eq. 4.41, and the white star the quantum emitter single photon source parameters in eq. C.1. Three regimes appear in the plot, corresponding to different ratios of p_d to $p_m p_l$. When the former dominates, as in the examples seen above, larger fusions are preferable and $n = 5$, the maximum included in this analysis, is found to be optimal. On the other hand, when combinations of

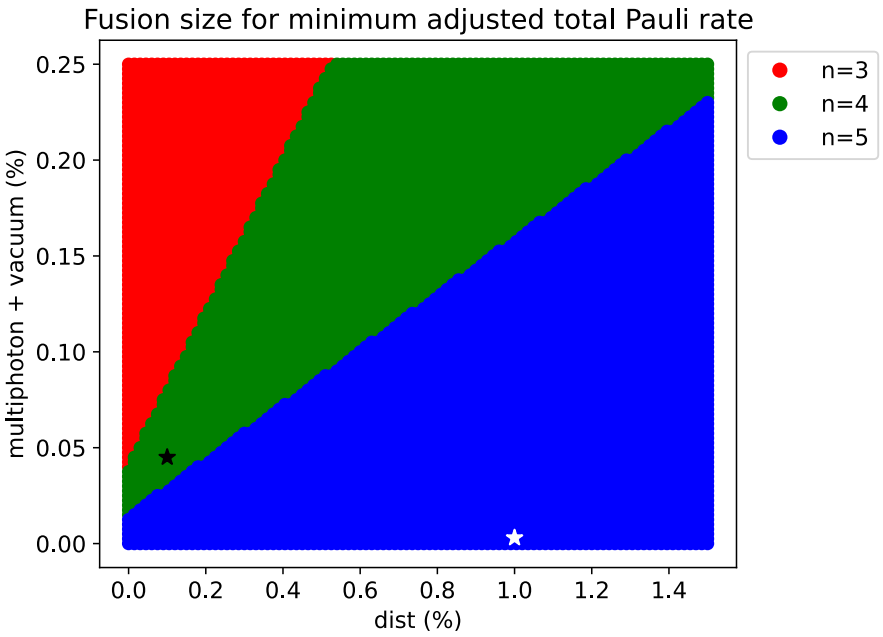

Figure C.4: Fusion size resulting in the minimum adjusted total Pauli error rate for different values of p_d and $p_m p_l$. As p_d increases, larger fusions become preferable due to the linear scaling of distinguishability with fusion size. On the other hand, when $p_m p_l$ is larger, smaller fusions are optimal. The black star shows the heralded single photon source parameters from eq. 4.41, and the white star the quantum emitter single photon source parameters from eq. C.1.

multiphoton and loss are higher, smaller fusions become preferable due to the combinatorial scaling of these terms with fusion size. From this plot, it is clear that when designing an architecture for a given hardware platform, this analysis should be repeated under its specific error model to identify its position on this map. Moreover, the analysis described here is limited to standard multi-way fusions, but other fusion circuits presented in Chapter 3 and analysed in Chapter 5 could also be included.