

SCALABLE ARCHITECTURES FOR PHOTONIC QUANTUM COMPUTING

SARA FRANCO

Centro de Física, Universidade do Minho, Braga 4710-057, Portugal

Supervised by Prof. Ernesto Galvão

Co-supervised by Prof. Mikhail Vasilevskiy

MAP-fis Physics Doctoral Program - Essay Curricular Unit

April 11, 2025

Abstract

Out of various physical systems currently being explored as potential platforms for scalable quantum computing, the distinctive properties of photonic-based approaches make them stand out as notable candidates. Yet, intrinsic limitations curtail large-scale practical implementation, namely, the unavoidable non-determinism of entangling photonic quantum gates and resource intensive multiplexing and active-feed-forwarding requirements. While the past 25 years have brought groundbreaking new models and promising architectures, a large-scale integration between hardware components and computational protocols is yet to be fully realized. In this essay, we propose a research project aimed at tackling the technical challenges involved in implementing scalable photonic quantum computing architectures. We will investigate how to improve the performance of modular building blocks. We will assess resource trade-offs and overhead of different solutions, and identify gaps in knowledge on the performance of recently proposed protocols. We will explore strategies for embedding error-correction into different architectures, an essential ingredient for fault-tolerance.

The present essay is structured in three parts. First, a general introduction contextualizes the field of photonic quantum computation and presents the motivations underlying the proposed research project. Chapter II provides an overview on the state-of-the-art of the field. It will serve to highlight recent results and to identify open questions and problems regarding the most promising proposals for photonic quantum computing architectures. The research project work plan is outlined in chapter III. The main research questions and intended goals are laid out in light of the state-of-the-art review, along with the methodological approach we propose.

Contents

I	Introduction	3
II	State-of-the-art	5
A	Quantum Information with Photons	5
B	Linear-Optical Quantum Computing	6
B.1	The Knill-Lafflame-Milburn scheme	7
B.2	Measurement-Based Quantum Computing	7
C	Alternative architectures	17
C.1	Spin-Optical Quantum Computing	17
C.2	Counterfactual quantum gates	19
D	Quantum Error-Correction	21
D.1	Errors in photonic platforms	22
D.2	Quantum-error correction in photonics	23
III	Research Proposal	24

I Introduction

In recent years, growing research efforts have been directed at exploiting the properties of quantum systems for new technological possibilities. Quantum information science showed that storing, processing and transmitting information via quantum systems can result in gains in speed and functionality. This key idea gave rise to the field of quantum computing, which leverages superposition, entanglement, and other quantum effects in order to gain computational advantage in certain information processing tasks. Its ultimate goal is the construction of a large-scale quantum computer with millions of quantum bits, or qubits, capable of solving certain problems dramatically faster than any classical computer.

A quantum computing architecture is the union of two ingredients: a theoretical quantum computing model, which governs the way in which information processing takes place; and a physical platform, a collection of hardware components that enables the implementation of the model. Small-scale prototypes have been realized in several physical platforms, and current research efforts are being directed towards the development of scalable architectures. Photonic systems in particular present unique advantages, as photons can travel long distances with little decoherence and can be manipulated in miniaturized integrated chips without the need for expensive cryogenic infrastructure. They are not, however, without their challenges. Photons suffer from non-negligible losses, and weak photon-photon interactions encumber the realization of entangling logic gates. These constraints must be taken into account in the development of scalable photonic quantum computing architectures.

A major breakthrough in overcoming these obstacles was the idea of linear optical quantum computing (LOQC), which took off with the 2001 proposal by Knill, Laflamme and Milburn (KLM) [1]. This was the first proposal for a scalable photonic architecture relying only on simple linear optical components, single photon sources and photodetectors. It put forward a solution to the issue of realizing two-qubit entangling gates between photons: a scheme based on probabilistic gates which induce effective nonlinearities through measurement. However, while scalability was shown in theory, estimates on the required number of optical elements to achieve efficient operation indicate this scheme is unfeasible in practice [2].

Since then, more promising models emerged, largely inspired by ideas from measurement-based quantum computation (MBQC) [3]. In MBQC, information processing takes place in rounds of measurements on qubits belonging to large entangled resource states, or cluster states. All entangling gates, used to grow the cluster from smaller seed states [4], can be enacted offline, before the computation takes place. This enables the use of non-deterministic gates of lower success rates, which have a smaller resource footprint. The fusion-based quantum computing (FBQC) approach, recently proposed by California-based company

PsiQuantum [5], takes this simplification one step further. Instead of growing the full cluster first and then executing the computation, alternating steps of entanglement and measurement are performed, with only part of the cluster being "alive" at any given moment. This removes the need of generating large cluster states, with only small entangled states of a constant size being consumed throughout the computation. When paired with quantum error correction and quantum memories (in the form of e.g. optical fibers), FBQC enables, in theory, scalable, fault-tolerant quantum computing.

To date, the main issues faced by MBQC platforms are the lack of efficient, scalable ways to generate photonic entangled states, which often resort to resource intensive multiplexing of imperfect sources, and the need for active feed-forwarding and fast interferometer reconfiguration based on previous measurement outcomes during the computational process. While the FBQC architecture alleviates some of these issues, active research is still being funneled to the development of efficient resource state generators [6], boosted entangling gates [7, 8], and ballistic schemes that aim to altogether avoid the need for active feed-forwarding [9, 10]. Yet another crucial ingredient for scalability is quantum error-correction. Quantum information is fragile, and different proposals for implementing error correction in different architectures are currently being explored in the literature [5, 11, 12, 13].

Outside of the aforementioned models, there have been recent proposals for photonic quantum computing which, while adjacent to the linear-optical framework, take on different approaches. For example, hybrid light-matter schemes have been proposed, which combine matter-based systems with linear-optical setups. Their basic idea is to use quantum single photon emitters with a spin degree of freedom, such as single atoms or quantum dots, which can be manipulated to emit spin-entangled photons. Precise control over these spin systems has been demonstrated to allow the efficient generation of photonic-entangled states in a deterministic manner [14]. Such systems could be used as deterministic resource state generators in exclusively linear-optical processors, or could be exploited as spin qubits themselves, as is the case in the recently proposed spin-optical quantum computing (SPOQC) architecture [11, 12].

In this essay, we will explore what are currently some of the most promising architectures for scalable photonic quantum computing. We will critically assess their strengths and drawbacks, and highlight key open questions and active research directions. Chapter II begins by introducing the main characteristics and challenges associated to the implementation of photonic quantum information platforms. We then introduce linear optical quantum computing architectures, focusing on measurement-based models. We explore techniques for improving resource state generation and alleviating requirements on adaptivity and interferometer reconfiguration, which highlight the importance of considering hardware restrictions when designing experimental schemes. This review will clarify the relevance and innovative aspects of the FBQC

model, whose main architectural features and advantages we discuss. We also explore the aforementioned alternative schemes for photonic quantum computing. We introduce the recently proposed SPOQC model, briefly outlining its unique advantages, and present a counterfactual scheme for photonic gates proposed by Salih et al., which might have applications in client-server quantum computing settings. Finally, we will overview basic concepts regarding quantum error-correction and briefly summarize recent proposals for its implementation in photonics platforms. This will help to identify the most promising routes for fault-tolerance in photonic architectures, fit for exploration in the remainder of my PhD project.

II State-of-the-art

The present chapter aims to be a concise review on recent approaches to photonic quantum computing. We will cover the relevant computational models and promising proposals for physical architectures, presenting a critical outlook on their main features and limitations. We will discuss strategies for overcoming key issues in their practical implementation, and identify open questions surrounding the performance and viability of existing methods. We will briefly introduce the topic of quantum error correction and recent proposals for its implementation in photonic platforms.

A Quantum Information with Photons

Several degrees of freedom of light are at our disposal for quantum information encoding [15]. The preferred strategy in the context of quantum computing tends to be that of *dual-rail* encoding, where each logical state of the qubit corresponds to single photon occupation in one of two orthogonal optical modes. Path- and polarization-encoding are two examples of this strategy. In the former case, two identical spatio-temporal modes (in profile, polarization, center frequency and so on) that are spatially separated in the transverse direction are used to represent each logical bit, while in the latter, two orthogonal polarization states are chosen. Optical encoding approaches can be further categorized as discrete-variable (DV) and continuous-variable (CV). DV approaches leverage the degrees of freedom of single photons, in discrete qubits or qudits. Meanwhile, CV optics exploit the continuous-valued observables in the phase space of quantum harmonic oscillators associated with the modes of the electromagnetic field. A thorough review of quantum computing with CV quantum optics can be found here [16]. For the purposes of this essay, we focus exclusively on photonic quantum computation architectures relying on the DV approach.

Photonic platforms possess unique features that make them excellent candidates for quantum information processing. Unlike other quantum computing technologies, they do not require cryogenic operation,

which enables the construction of modular and easily-networked quantum computers at room temperature. Photons are a naturally low-noise system, suffering from little decoherence due to their limited interactions with the environment or each other. This property, paired with their inherent compatibility with fiber optical communication technology, makes them ideal carriers of quantum information over long distances. Photons will likely be indispensable for transmitting quantum information in communication tasks, or for its distribution between different parts of a quantum network. Despite these compelling features, significant hurdles limit the scalability of photonic quantum computers. The same property that grants photons robustness against noise makes interactions between individual photons difficult to manage, encumbering the realization of two-qubit entangling gates. Photonic systems are highly sensitive to optical losses, due to absorption/scattering in optical components and imperfect detection efficiency, which introduce errors and loss of information. The precise manipulation of light requires fast control over optical switches, which puts a strain on the requirements for classical processing and real-time feedback.

Following the recent review [15], we can list four fundamental requirements for photonic quantum information processors. Firstly, they must be large-scale in order to handle complex tasks. Secondly, they must be universal, i.e., they must enable the implementation of arbitrary transformations, so that the system can be applied to solve a diverse set of problems. Thirdly, optical loss must be maintained to a minimum. And, finally, it must achieve quantum interference, as this ensures the fidelity and reliability of quantum operations. Photonic integrated circuits have recently emerged as a powerful platform for meeting these requirements, offering compactness, scalability, improved stability over bulk optical setups and unprecedented versatility and universality [17].

B Linear-Optical Quantum Computing

Linear optics is the restricted dynamics of light corresponding to unitary evolutions of the creation operators associated to optical modes which preserve the number of photons. Physically, they can be implemented with simple optical components such as beamsplitters and phase shifters, present in any quantum optics laboratory. In this chapter, we will review photonic quantum computing architectures relying solely on linear optical elements. Section B.1 introduces the landmark Knill-Lafflame-Milburn proposal, which laid the groundwork for the measurement-based architectures presented in section B.2. In section B.2.2, we describe several strategies for improving the performance of resource state generators and probabilistic entangling gates, essential in the path towards scalability. Finally, section B.2.3 presents the promising fusion-based quantum computing architecture, and how it incorporates several of those strategies.

B.1 The Knill-Lafflame-Milburn scheme

As we have discussed in section A, a main feature of photonics system are the limited interactions that photons exhibit among themselves and the environment. While this represents an advantage in terms of the low decoherence experienced by photons, limited photon-photon interactions present a challenge in the realization of two-qubit entangling gates between photonic qubits, necessary to achieve universal quantum computation. Early approaches for inducing photon interactions involved harnessing optical nonlinearities in certain materials. However, naturally occurring nonlinear effects are weak in magnitude, such that nonlinear optical processes are extremely inefficient at the single photon level, hindering scalability.

It was not until 2001 that Knill, Laflamme and Milburn (KLM) showed it is possible to achieve efficient, universal quantum computation using only linear optics [1], thereby propelling the field of Linear-Optical Quantum Computing (LOQC). They showed that, in theory, scalable quantum computation can be realized using only linear optical elements, along with single-photon sources and photodetectors. Their scheme circumvents the need for nonlinear interactions between optical modes, rendering it more amenable to practical implementation. In it, one qubit gates result from deterministic linear-optical transformations on dual-rail qubits (see section A). Meanwhile, two-qubit gates are implemented through effective nonlinearities induced by measurement and post-selection, relying on auxiliary entangled photon pairs in additional optical modes and teleportation-based schemes [18]. These gates are non-deterministic, meaning they must be repeatedly attempted until they succeed. Measurement outcomes heralding the success or failure of the gate need to be classically feed-forwarded to other parts of the optical circuit for appropriate reconfiguration. The success probability can be brought to near-unity through the use of larger numbers of auxiliary modes, but not without an unfeasible demand on resources. Thus, while scalability was shown in theory, estimates on the number of required optical elements indicate that the overhead cost of such two-qubit gates is prohibitively large.

In fact, photonic platforms continue to grapple with significant scalability issues, due to this unavoidable probabilistic nature. Nevertheless, intensive work in the field has since produced alternative, more promising models for LOQC, which we will explore in the following section.

B.2 Measurement-Based Quantum Computing

Historically, the earliest quantum computing models were largely inspired by well-known classical computation frameworks. One clear example of this is the circuit model for quantum computing, in which information processing takes place through the unitary evolution of qubit states acted upon by quantum

gates, with conversion of quantum information to classical information through measurement and read out at the end of the circuit. In the quest for the development of scalable quantum computing architectures, alternative classes of quantum computation models have emerged. One such example is the Measurement-Based Quantum Computing (MBQC) model, which presents an alternative framework for quantum information processing, with no direct classical analog. It is grounded in the key result that universal quantum computation is possible using only measurements as the computational steps. Generally speaking, in MBQC models, information processing takes place solely in rounds of measurements on qubits initially prepared in highly entangled states.

We can highlight two main kinds of MBQC models that, while sharing this key feature, take on different approaches. One is the teleportation-based quantum computing model, inspired by the idea of Gottesman and Chuang of teleporting probabilistic quantum gates onto computational states [19]. When probabilistic gates are applied to information-carrying qubits, they destroy the information therein upon failure. If these gates are instead implemented on offline qubits entangled with the computational qubits, they can safely be repeated until success and subsequently teleported onto the computational qubits, safeguarding the data. Teleportation is achieved through projective two-qubit measurements on auxiliary Bell-pairs and post-selection. It was this idea that directly inspired KLM's probabilistic linear-optical entangling gate. The KLM scheme can thus be thought of as an MBQC model itself [20], since in it photons don't interact as in standard circuit models, but instead via the application of these teleportation-based gates. Another kind of MBQC framework is the *cluster-state*, or *one-way quantum computer* (1WQC) model. In this model, first proposed by Raussendorf and Briegel [3], sequences of one-qubit measurements are performed on a particular kind of entangled, multi-qubit state, called a cluster state. It can be shown to be equivalent to the teleportation-based model [21], and it has been the main inspiration for the development of subsequent MBQC approaches. In order to understand this model, we first need to introduce the concept of a *graph state*.

Graph states

In graph theory, a *graph* $\mathcal{G}(V, E)$ is a pair consisting of a finite set of *vertices* $V = \{1, \dots, N\}$ and a set of *edges* $E \subset [V]^2$, tuples of elements in V . Two vertices $a, b \in V$ are said to be *adjacent* if they share an edge, i.e., if $\{a, b\} \in E$. A *Graph state* is the state of a multipartite quantum system which can be represented by a graph. Formally, it is defined as follows [22]:

Definition 1. (*Graph State*) The graph state $|\Psi_{\mathcal{G}}\rangle$ associated to a graph $\mathcal{G}(V, E)$ is the pure state of a multipartite system where: each qubit is labelled by a vertex $i \in V$, and is initially prepared in the state

$|+\rangle_i$; each pair of qubits whose corresponding vertices i, j are adjacent is made to interact via an entangling $CZ = \text{diag}(1, 1, 1, -1)$ gate, The graph state can thus be written as

$$|\Psi_G\rangle = \prod_{\{a,b\} \in E} (CZ_{ab}) |+\rangle^{\otimes a \in V}. \quad (1)$$

In particular, a *cluster state* is a subset of graph states with a regular geometry. For example, a $1D$ cluster state corresponds to a graph which can be represented by a $1D$ chain, a $2D$ cluster state corresponds to a square lattice, and a $3D$ cluster state to a cubic lattice. Graph states are convenient objects, as they have an efficient description in the stabilizer formalism of quantum mechanics [23], a powerful framework used to keep track of eigenstates of multiqubit Pauli operators. While detailed treatment of this formalism lies beyond the scope of this essay, we note its wide-ranging applications in MBQC, quantum error-correction and fault-tolerant quantum computation.

B.2.1 The one-way quantum computing model

The resource state used for 1WQC is typically a two-dimensional cluster state, i.e., a rectangular grid of $|+\rangle$ states generated by entangling CZ gates on each pair of neighbouring qubits. Information can be written onto the lattice, processed and read out by sequences of single-qubit measurements only. In an initial stage, the lattice is "shaped" through computational basis measurements, which disconnect affected qubits from the cluster. Next, information is propagated through the cluster by means of single qubit measurements, in a certain order and basis. The choice of basis is dictated by the results of preceding measurements, i.e., they are adaptive measurements controlled by feed-forwarding of classical outcomes. It can be shown that any quantum gate array can be translated into such a pattern of adaptive one-qubit measurements on certain universal resource states, such as the 2D cluster [3]. In the final stage, the result of the computation is obtained from measurement of the last remaining qubits. The computational power of the model is strongly related to the properties of the initial state, as it is the substrate onto which universal computation can be efficiently carried out. The size of the resource cluster state scales up with the computation to be performed, as it must be large enough that any possible quantum state can be "carved" out of it, and provides all the entanglement involved in the subsequent computation in advance.

The KLM scheme (see section B) overcame a big obstacle faced by photonic quantum information platforms - the apparent need for nonlinear interactions between optical modes. However, its probabilistic, linear-optical entangling gates demand an unfeasible number of optical components, rendering it unachievable in practice. At first sight, it seems that the Raussendorf-Briegel model does not overcome this problem,

and instead simply transfers the resource overhead onto the task of generating the initial entangled state, which in principle would require the same resource intensive entangling gates. In fact, the generation of cluster states is the main scalability hurdle faced by 1WQC models. But there is one key difference in these two approaches. In the KLM scheme, gates are being implemented throughout the computation, so they need to have a high-success probability in order to avoid losses of information and the need for prolonged quantum memories. Meanwhile, in the 1WQC approach, entanglement is only required to grow the cluster state offline. Thus, it is only required that, for every entanglement attempt, the cluster grows *on average*, meaning the success probability of creating a new link in the cluster needs only to be larger than one half. As it turns out, there have been several, increasingly efficient proposals for generating entangled states with only linear optics and photon detection that meet this weaker condition. In the following section, we will present an overview on some of these schemes.

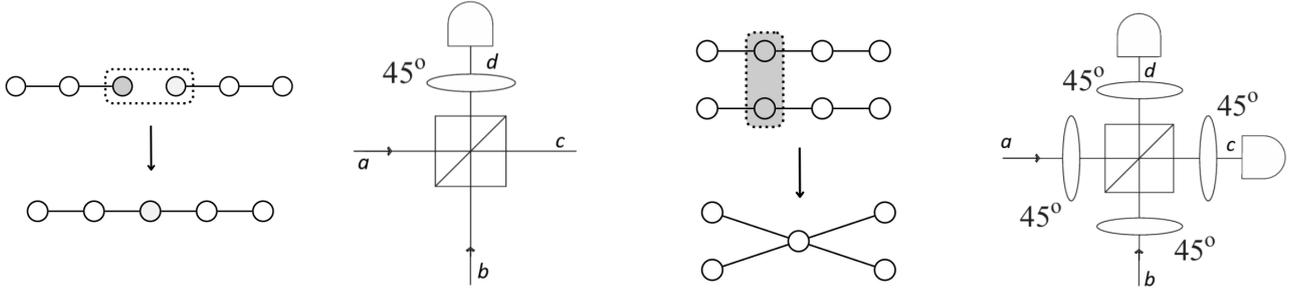
B.2.2 Photonic entangled state generation with linear optics

In this section, we will explore state-of-the-art methods for generating and growing photonic entangled states using only linear-optics, as well as other strategies for dealing with errors introduced by linear-optical entangling gates, which are inherently probabilistic.

Fusion gates

Early approaches for growing cluster states involved the use of non-deterministic entangling gates, with a success probability larger than $1/2$, to link qubits to the cluster one-by-one [24]. However, if instead small chains of qubits are added at a time to the growing cluster, this requirement on the success probability could be relaxed. Indeed, in 2005, Browne and Rudolph introduced a simpler yet more efficient strategy for cluster state growth [4] which involves joining, or *fusing* together, two smaller resource states to create a larger entangled state. They proposed two probabilistic gates, the *Type-I* and *Type-II* fusion gates, which involve only beamsplitters and photodetectors, making them remarkably easy to implement in practice.

Type-I fusion allows us to grow arbitrarily long linear cluster states by linking smaller chains. Starting from two linear clusters of lengths m and n , this gate (when successful) deletes the end-qubit of one of the chains, with the end-qubit from the other inheriting the bonds of the deleted qubit. In this way, a chain of length $m + n - 1$ is formed. Type-I fusion can be realized for polarization-encoded qubits using the simple setup depicted in Fig. 1a, consisting of a polarizing beamsplitter, followed by a 45° polarization rotation on one of the output modes and, finally, a photon number resolving detection in this same mode. The end qubits of each cluster are sent through input modes a and b and made to interfere in the beamsplitter.



(a) **Left** Graph representation of Type-I fusion. Two smaller linear qubit chains are fused into a larger one-dimensional cluster. **Right** Experimental implementation of Type-I fusion. Adapted from [4].

(b) **Left** Graph representation of Type-II fusion. Two linear qubit chains are fused into a two-dimensional graph state. **Right** Experimental implementation of Type-II fusion. Adapted from [4].

Figure 1: Type-I and Type-II fusion gates. Details in main text.

With probability $p = 1/2$, one and only one of the incoming photons is detected at output d and deleted from its original chain, and fusion is accomplished. A failure has the effect of disconnecting them from their respective clusters. Provided the remaining chains are long enough, fusion could be reattempted.

One-dimensional cluster states are not sufficient for universal quantum computation, since their geometry does not permit the implementation of two-qubit gates. Type-II fusion enables us to generate two-dimensional cluster states by fusing together qubits in the middle of two separate linear chains¹. Type-II fusion can be implemented simply by adding 45° rotators to both inputs and outputs of the polarizing beamsplitter used in Type-I fusion, as can be seen from Fig. 1b. One additional ingredient is required for Type-II fusion - redundant encoding. In redundant encoding, a single logical qubit is encoded onto multiple physical qubits, such that a generic cluster state $|\phi_0\rangle|0\rangle + |\phi_1\rangle|1\rangle$ can be encoded as $|\phi_0\rangle|0\rangle^{\otimes n} + |\phi_1\rangle|1\rangle^{\otimes n}$, where n physical qubits encode one logical state. In Type-II fusion, one photon belonging to a redundant encoding from each chain is sent through one input mode, and both are detected at modes c and/or d . Whenever one photon is detected in each detector, which occurs with $p = 1/2$ chance, the two are deleted from their respective chains, which become fused in a 2-dimensional shape. If the gate fails, the two qubits are removed and the two chains remain separate, retaining their structures. Provided enough qubits remain in the redundant encoding of each chain, fusion can be retried.

One advantage of Type-II fusion is that it does not require photon-number resolution to herald a success, and it naturally detects optical losses. It can be paired with Type-I fusion, which can be used to create the

¹Note that Type-I fusion cannot be used to this end, since on a failure, it would split each starting linear chain in two, and thus would not be an efficient strategy.

starting chains to be fused, or it can be used by itself to grow the linear clusters and then join them. In this later case, however, Bell pairs are not a sufficient resource, as we need at least one three-qubit chain to grow a state with Type-II fusion, due to the need for redundant encoding. Browne and Rudolph's protocol improves on previous proposals in several aspects. It replaces complex, teleportation-based entangling gates with a simple apparatus, making it more efficient in the number of optical components required. It also requires fewer photons and measurements to grow a given cluster. Furthermore, since it makes use of Hong-Ou-Mandel interference [25], it only requires stable interferometry over the coherence length of photons, unlike KLM and similar approaches which need stability at the wavelength scale. It is an efficient model for universal quantum computation and has inspired the promising fusion-based quantum computing model.

Boosted fusion

While the Type-I and Type-II fusions gate of the previous section can be used to grow a cluster state, their low success probability of $p = 1/2$ still places strong requirements on multiplexing and quantum memories. It is therefore of great interest to try to boost their success rates.

The problem of enhancing Type-II fusion is intimately related to another important problem in linear optics - that of realizing a Bell State Measurement (BSM). A BSM is a joint measurement of a Bell pair that unequivocally determines which of the four Bell states the pair is in. It is known that a complete BSM cannot be achieved using only linear optical components, and there is a rich history of attempts at enhancing the efficiency of existing protocols [26, 7, 6, 8]. Type-II fusion can be seen as a BSM in a rotated basis, and thus boosted BSM schemes are of great relevance in the context of building scalable LOQC platforms. In [26, 7], BSMs of $3/4$ efficiency² were designed using as auxiliary modes a two-photon entangled state, and four unentangled single photon states, respectively. In both references it was shown that this number can be enhanced using more auxiliary modes and photons, but only with a slight improvement that is not worth the experimental cost. More recently, Bartolucci et al. showed that values of $7/12$ and $2/3$ could be reached using only one or two single photons, respectively [6]. They also proposed the only known scheme to date for boosting Type-I fusion, which reaches $3/4$ efficiency but at the cost of needing an ancilla Bell pair and feed-forwarding. Boosted fusion gates could reduce multiplexing and quantum memory requirements for cluster state growth. As we have seen, however, single photons are not a sufficient starting resource for fusion, and it is imperative that we have some efficient way of creating

²To be precise, these protocols can exactly discriminate two of the possible four Bell states, and determine the other two half of the time.

two- or three-qubit entangled states to serve as seed states. In the next section, we will briefly overview recent protocols for their generation.

Bell state and GHZ state generation

There are several different procedures for entanglement generation in photonics, from exploiting nonlinearities in materials to using linear optical interference and measurement to induce effective nonlinearities. However, they all are inherently probabilistic, and tend to have low success probabilities; for example, [27] reaches a success rate of $p = 2/27$ for Bell-state generation, while [28] achieves an improved but still small $p = 1/9$. In the path towards scalable LOQC architectures, it is urgent that enhanced schemes for entanglement generation are developed.

One interesting approach is the one taken in [6] by Bartolucci et al., where several protocols for linear optical entanglement generation are proposed. One of those is a Bell state generator, which uses four unentangled single photons, or signal modes, each made to interfere at a symmetric beamsplitter with an ancilla vacuum mode. The ancillary modes are then coupled at a particular 4-mode linear interferometer and detected. Probabilistic detection of a certain measurement pattern heralds the creation of a two-photon entangled state in the signal modes. The success probability of generating a Bell pair with this method is of $p = 1/8$, or $p = 3/16$ if active feed-forward is implemented. This success rate can be further pushed to $p = 1/4$ by means of a distillation stage, and introducing a preliminary try-until-success procedure, called *bleeding*, can bring this probability to near unity in the limit of infinite iteration (details in [6]). This bleeding stage requires repeated application of the lossy switch operation, which might undercut the improvement in generation efficiency - a more thorough benefits/cost analysis is required in order to formulate any conclusion. The same is true of the feed-forwarding stage, which can likewise introduce extra errors and complicate technical realization. Another relevant scheme is an n -photon GHZ state generator from $2n$ single photons, which makes use of the Type-I fusion gate introduced in the previous section. Two indistinguishable photons are coupled at a balanced beamsplitter, producing a path-entangled Bell pair through the HOM effect. These Bell-pairs are then fused together iteratively until the resulting graph state is the desired size. Finally, the two outermost qubits are fused in order to build a closed graph, resulting in a GHZ state. Since the success of each fusion is of only $p = 0.5$, the success rate exponentially decays with the size of the GHZ state, but active feed-forwarding and bleeding can also be used in this case to improve this, with the caveat of introducing the additional costs mentioned above.

Probabilistic growth and percolation

Throughout the past sections, we have seen how one main difficulty limiting the growth of large photonic cluster states is the lack of deterministic linear-optical entangling gates. Some clever probabilistic schemes have been devised [4, 7, 6, 8], but ultimately exhibit a trade-off between large success probability and resource overhead. Another challenge is posed by the need for active switching, where upon success or failure of previous gates, quantum systems need to be appropriately rerouted into coherent interactions with other quantum systems. This places a strain on quantum memory resource requirements and classical control systems.

In an attempt to mitigate these issues, Kieling et al. proposed a "novel way to cope with the randomness of non-deterministic quantum gates" [9]. Their approach applies results from classical percolation theory, which describes the behaviour of networks where bonds are added, or opened between nodes with some probability p . Here, nodes correspond to small blocks of entangled states, which are passively linked using fusion gates without the need for feed-forwarding. The final entangled state thus obtained corresponds to a lattice with missing bonds, with open and closed edges representing successful and unsuccessful fusion attempts, respectively. The power of this approach lies in the existence of a critical bond formation probability p_c above which the created state can be renormalized, through a pattern of single qubit measurements, into a universal cluster state for quantum computation. This pattern can be calculated offline through efficient, classical computation. Hence, it suffices that the success probability of the implemented fusion gates is above this threshold, and we can design the initial resource states and lattice geometry to that end. Remarkably, this strategy does away with the need for active switching with a resource overhead that scales as if deterministic gates were available. Yet, it requires at least four-photon entangled states as resource, which as we have discussed, are not trivial to prepare.

Gimeno-Segovia et al. [10] improved on this approach by introducing loss-detecting gates in an architecture that is robust to photon loss and uses only 3-photon GHZ states and boosted fusion. Morley-Short et al. [29] took into account the physical constraints of a realistic architecture with finite and fixed depth. They showed that large universal states can be created by alternating clock cycles of photon generation and entanglement with measurement stages. Essentially, the cluster state is continuously generated in sections interspersed with measurement stages, so that only a finite portion of the lattice is "alive" at a given moment. Universal LOQC can in this way be implemented on a constant-size device of modest depth, and simple algorithms can be employed to shape the percolated lattice, with only small sections of it being analyzed at a time. This idea laid the foundation for the aforementioned fusion-based quantum computing model, which will be described in the next section.

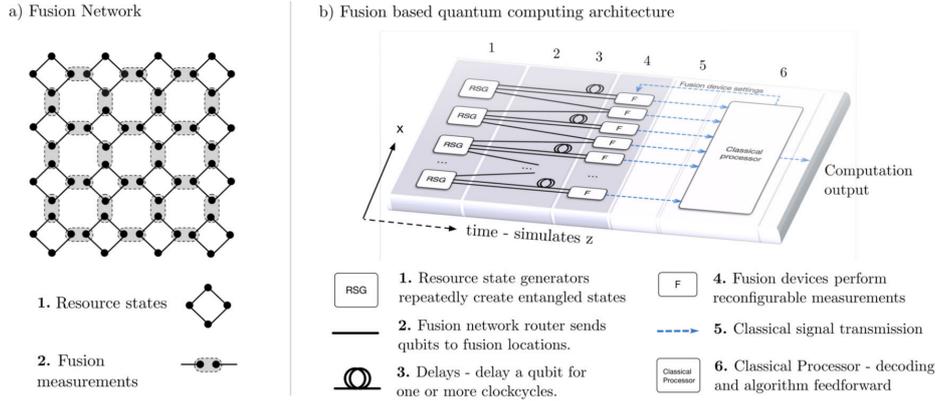


Figure 2: The fusion-based quantum computing (FBQC) architecture. a) A fusion network is the central object of FBQC. It combines the two primitives of the model: small entangled resource states of a constant size and structure (here, the 4-qubit, 4-cycle graph states), and entangling measurements that fuse them together. The network provides the configuration of measurements to be performed. b) A possible physical architecture for implementing FBQC. See details in main text. Adapted from [5].

B.2.3 Fusion-Based Quantum Computing

Fusion-Based Quantum Computing (FBQC) is a MBQC-inspired model for universal fault-tolerant quantum computation, introduced by the California-based company PsiQuantum [5]. It is built on two primitives:

1. **Resource states.** FBQC makes use of small entangled resource states, also called seed states, with a constant size and structure regardless of the scale of the computation to be implemented. These are generated by a Resource State Generator (RSG), a physical system that can produce copies of the seed states to be consumed during the computation.
2. **Fusions.** Fusions are projective entangling measurements on multiple qubits that are performed on qubits belonging to seed states. They can be, for instance, the fusion and boosted fusion gates described in section B.2.2.

These two primitives are used to construct *fusion networks*, the central object of FBQC. A fusion network is a configuration of fusion measurements to be performed on qubits belonging to a collection of resource states. An example is depicted in Fig. 2 a), where resource states correspond to four-qubit graph states in a square geometry which are measured pairwise in fusions, depicted by the shaded ovals. Algorithms can be executed on the network by appropriately choosing the basis of the fusion measurements and combining their measurement outcomes.

Unlike 1WQC, where the size of the resource state scales up with the complexity of the computation being performed, in FBQC the size of the seed states remains constant, and only the number of seed states needed increases. In addition, whereas 1WQC relies only on single-qubit measurements to execute algorithms, here multi-qubit projective measurements, or fusion gates, drive the computation. While not restricted to any specific physical system, FBQC is particularly well suited for photonic platforms, since it is constructed from physical primitives readily accessible in photonic systems.

An example of a physical architecture capable of constructing fusion networks is schematized in Fig. 2 b). This architecture implements the idea introduced by Morley-Short et al.[29] of interleaving photon generation and measurement stages, such that only a part of the network exists at a given point in time. Seed states are generated at a given clock speed in the first layer by RSGs. These states are sent through a fusion network router and distributed across fusion locations. Delay lines just before the fusion devices keep some of the qubits for one or a few more clock cycles. This enables the creation of three-dimensional fusion networks, where time plays the role of the z axis, as indicated in the figure. Overall, individual qubits generated in layer 1 and destructively measured in layer 4 have a very limited lifetime in the device. This is one of the biggest strengths of the FBQC model: not only does it reduce the chance of optical losses, but any given photon only sees a very limited, constant number of optical components before being measured, which significantly reduces the occurrence of errors. Furthermore, it limits how far errors can spread in the network, as errors remain localized. After a round of fusions, classical signals are fed to a classical processor, responsible for implementing the algorithmic feed-forward, reconfiguring fusion measurements based on previous outcomes. There are further architectural simplifications offered by FBQC when compared to previous models. Assuming measurement is fast when compared to the qubit lifetime (which is generally true of photonic systems), a limited number of RSGs can be used and reused to generate the whole fusion network. Moreover, even when components are being reused, the fusion network routing is fixed, which eliminates the need to switch between different device configurations, a process that is error-prone and places a burden on classical control. All feed-forward and reconfiguration is thus restricted to the switching between measurement basis in fusions required to implement logic.

Finally, any actual implementation of a quantum computing architecture will be subject to different sources of error and noise, which can corrupt or erase quantum information. The topic of errors and how we address them in quantum computing is addressed in more detail in section D. We have seen already some features of the FBQC model which limit the impact of errors. Additionally, FBQC is naturally fault-tolerant, as it enables the implementation of quantum error correction protocols, and shows notable performance improvements over previous architectures. Specifically, it exhibits a large tolerance for photon

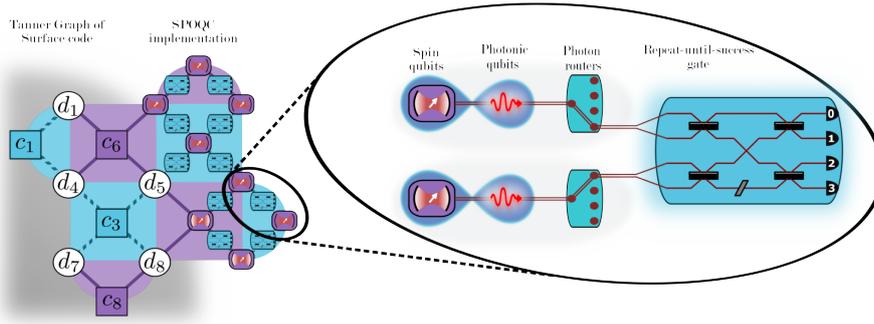


Figure 3: Spin-optical quantum computing architecture. **Left** spin qubits are arranged in the surface code configuration (see section D), divided between data d and check c qubits. **Right** The spin emitters interact indirectly via spin-entangled photons, made to interfere in linear-optical repeat-until-success entangling gates. Adapted from [11].

loss during fusion and for fusion failure. Other articles have since been published [30, 31] which offer further insights into this architecture and its potential hardware implementations. Fault-tolerance, together with low operational depth, simplicity in device architecture and reliance on small constant-sized resource states, among other features, make FBQC a serious contender for a scalable platform for quantum computation.

C Alternative architectures

So far, we have mostly discussed measurement-based and fusion-based models for photonic quantum computation, a promising approach to fault-tolerance. Recently, alternative models have been proposed which could bring new functionalities and applications to the field. In this section, we will briefly introduce two alternative paradigms for photonic quantum computing: the spin-optical quantum computing architecture, a hybrid model combining the strengths of both light- and matter- based systems; and counterfactual quantum gates, a novel way to construct photonic gates with potential applications to client-server quantum computing models.

C.1 Spin-Optical Quantum Computing

Throughout section B.2, we discussed at length the main scalability hurdle faced by current photonic architectures - that of generating large entangled photonic states. In recent years, we have witnessed the emergence of a new kind of highly efficient photon sources based on quantum emitters, such as individual atoms or quantum dots (also known as artificial atoms) [32]. These high quality, on-demand single photon sources have made possible another avenue for photonic entangled state generation. It is possible to couple the spin of the quantum emitter with the emitted photon and generate spin-photon entangled states.

Precise control over the spin then enables the growth of photonic graph states, where the spin mediates the generation of entanglement between photons in a deterministic way [33]. These so-called "photonic machine guns" have already been used to demonstrate the efficient generation of entangled photon states in a controlled and pulsed manner [14]. Quantum emitters could be embedded in fusion-based architectures as resource state generators (see section B.2.3). Their determinism provides a clear scaling advantage over probabilistic schemes, opening up a new road for photonic quantum computation.

Going one step further, Gliniasty et al. proposed a hybrid light-matter architecture [11] that maximizes the utility of these quantum emitter sources. They propose that the spin be leveraged not only as a single photon source but also as a qubit to encode quantum information, in a novel quantum computation model called Spin-Optical Quantum Computing, or SPOQC. In the SPOQC framework, physical qubits are encoded in the spin degree of freedom of the quantum emitters, where full control over the spin qubit's state is assumed. Photonic qubits are used as flying ancilla qubits, mediating two-qubit gates between the spins of non-interacting emitters. Specifically, photons entangled with their respective emitters are routed to linear-optical, repeat-until-success gates, where they are made to interfere. These are heralded, probabilistic controlled gates, meaning it is possible to ascertain whether they have succeeded, failed or aborted. Fig. 3 schematically represents this architecture.

In combining both light- and matter-based schemes, SPOQC benefits from the best of both worlds, exploiting both the spin and photonic qubits to their full potential. It avoids the need for complex linear-optical resource generation schemes such as the ones mentioned in section B.2, which typically rely on resource-intensive multiplexing. The distance between qubits could ensure better isolation and avoid correlated noise, a common issue across several physical platforms for quantum computation. As we will see in more detail in section D, the SPOQC architecture enables the implementation of any quantum error correction code. In particular, the long-range optical links between non-interacting emitters, facilitated by the flying photonic qubits, allow the execution of quantum low-density parity check codes, which are very promising non-local codes with small footprints [34]. The SPOQC architecture was theoretically demonstrated to reach comparable or even improved levels of fault-tolerance when compared to the FBQC approach reviewed in section B.2.3., at least in preliminary studies where the goal was just to preserve the quantum state, rather than apply encoded gates to it. An interesting proposal by Löbl et al. combines these two frameworks [35], in an architecture that uses fusion gates to indirectly fuse the spin qubits through spin-entangled photons in linear-optical gates. It borrows from percolation theory (discussed in section B.2) to avoid the need for adaptivity and extended quantum memories. Hybrid models present themselves as novel, versatile approaches to quantum computing, and certainly, research efforts will soon

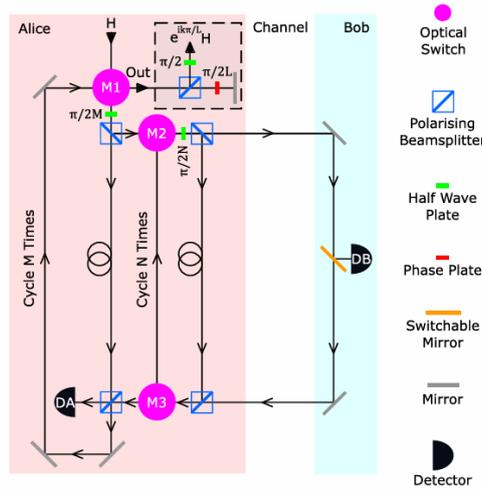


Figure 4: Salih et al.’s scheme for a counterfactual quantum gate. Alice sends the horizontal component of her polarization-encoded qubit through a nested interferometer scheme, a section of which is connected to Bob via a quantum communication channel. Bob can choose to block the channel for a certain number of rounds, thereby controlling a relative phase that is applied to the qubit, without needing to communicate with Alice. Details in main text. Adapted from [36].

start to yield experimental demonstrations of their capabilities.

C.2 Counterfactual quantum gates

The term "counterfactual" refers to something that is contrary to the facts or actual events that have occurred. In the context of quantum mechanics, it is commonly used to describe protocols inspired by the phenomenon of Interaction-free Measurements (IFM), also known as quantum interrogation tasks. The original protocol for an IFM was proposed in 1993 by Elitzur and Vaidman (EV). It consists of a single quantum particle, say, a single photon, traversing a Mach-Zehnder interferometer, where one path may be obstructed by an opaque object. If the object is absent, the interferometer is arranged for destructive interference at one output. If the object is present, even the possibility of interaction disturbs the interference, causing a photon to end up at the "dark" output port with some probability. Detection of a photon at that port thus signals the object’s presence without any photon being absorbed by the object. EV’s IFM scheme is probabilistic, in the sense that, at most, only half of the detections of the object’s presence are successfully interaction-free, with the other half corresponding to absorption of the photon by the object. In 1995, Kwiat et al. [37] proposed a multi-pass interferometric scheme that boosts the success probability arbitrarily close to unity, by leveraging the quantum Zeno effect [38] in multiple interrogations

of the object.

IFM protocols have inspired several practical applications [39, 40, 41, 42]. Of particular interest in the context of the present essay is a scheme for a counterfactual single-qubit quantum gate proposed by Salih et al. [36]. In principle, this gate allows a classical Bob to remotely manipulate an unknown qubit at Alice, without any communication channel between the two and without previous knowledge on the qubit's state. More precisely, it allows Bob to implement an arbitrary relative phase between the vertical and horizontal components of a polarization-encoded photonic qubit's state. The protocol can be summarized as follows: Alice prepares a single photon in an arbitrary polarization state, and sends its horizontal component through a nested-interferometer scheme based on Kwiat et al.'s idea, represented in Fig. 4. She keeps the photon in the setup during a predetermined number of cycles, previously agreed on between her and Bob, after which she switches it out. A section of the interferometer traverses a communication channel between Alice and Bob, which Bob can choose to either block or not block using some classical, opaque object, during a desired number of rounds. The number of rounds in which Bob chooses to block the communication channel controls the relative phase that is applied to the qubit after all rounds. By concatenating several of these gates in a certain arrangement, Bob can apply an arbitrary single-qubit gate on Alice's qubit. This represents a new paradigm for building single-qubit photonic gates, and thorough analysis on the resource overhead of these gates is still missing in the literature. Furthermore, it is possible, in theory, to combine these with a matter-based system in a hybrid model, not unlike the ones discussed in section C.1, in order to construct two-qubit entangling gates. This would make Salih et al.'s gate a universal scheme for quantum computing; however, experimental demonstration of the practical feasibility of such schemes is still needed.

From the practical stand point, one interesting feature of this scheme is that it allows Bob to remotely enact any arbitrary transformation on Alice's qubit through a very elementary mechanism. Simply by choosing whether or not to block the transmission channel with a classical object for a given number of rounds, Bob has full control of the gate he applies, without requiring any prior information about the state of the qubit. It also requires no classical information to be exchanged between Alice and Bob during the application of the gate, or that Alice manipulates the qubit in any way. All that is needed is for both parties to agree on the timing when Alice sends the photon in and the timing of each round. As the authors point out [36], there is a way to arrange the different instances of this gate in such a way that any unitary transformation Bob applies takes a fixed amount of rounds to implement. This could have interesting applications in a server-client quantum computing setting. Alice would play the role of the quantum server, holding the computational power and resources for generating and manipulating the

quantum information. She possesses the single photon source that generates the photonic qubits and controls the optical switches in the gate. Bob acts as the client, who can access this server to run a specific computation he has full control over using only a simple classical mechanism. He does not need to generate or manipulate quantum states at any point. Furthermore, it might be possible for Bob to conduct his remote computation in a secure way, where Alice does not have access to the quantum gate Bob is implementing. If so, this could be used in a blind quantum computation scheme [43], a client-server model for quantum computation where the server is blind to the client's data and algorithms. Counterfactual quantum gates could thus potentially be used in near-term practical applications of photonic quantum computing, and further research is necessary to assess this possibility.

D Quantum Error-Correction

Quantum systems are extremely sensitive, easily losing coherence in the presence of noise. For a quantum computing architecture to be scalable, it must operate at lower error rates than those currently exhibited by individual, physical qubits. Therefore, error-correction needs to be incorporated into any quantum computing model as a crucial ingredient for scalability. Starting in 1995, several proposals for quantum error correction (QEC) began to emerge [44], and currently, significant research efforts are being channeled towards the development of more rigorous theoretical frameworks for QEC and of new and improved codes tailored for different architectures (see [45] for a thorough review). A key theoretical result is the threshold theorem [46], which shows that fault-tolerant quantum computation is possible with only a poly-logarithmic overhead of physical qubits in the length of the computation. In accordance with this theorem, given a specific error model modeling the noise processes in the physical device, any quantum algorithm can be executed fault-tolerantly, so long as the physical errors in the qubits remain below a certain critical value, the so-called error threshold. Generally, this threshold value depends on the chosen QEC code and on the assumed characteristics of the noise.

QEC codes are protocols capable of detecting and correcting errors during the processing and storage of quantum information. While there are many different kinds of codes, they are all based on the same basic principle - redundant encoding, where a single logical qubit is encoded over multiple physical qubits. In this way, any lost or corrupted information can be recovered. Although QEC is largely based on classical error correction theory, there are key challenges we face when trying to apply its principles to quantum systems. Firstly, the no-cloning theorem of quantum mechanics forbids any form of coding based on data-copying, where data is protected by simply making multiple copies of it. Secondly, direct measurement cannot be used to detect errors, as it will destroy any quantum superposition state used for encoding information.

Lastly, unlike classical bits, which can only suffer from bit-flips, qubits are susceptible to a continuum of possible errors. While this might seem daunting, it turns out that, since any single-qubit error E can be expanded as a superposition

$$E = e_0I + e_1X + e_2Z + e_3Y. \quad (2)$$

of bit-flip errors (represented by the Pauli matrix X), phase-flip errors (Z) or a combination of the two ($Y = iXZ$), any QEC code needs only to correct for these two types of errors.

Before implementing error correction in a given platform, it is crucial to understand what types of errors it is most susceptible to, as this knowledge directs the kind of codes that are used. In the next section, we will summarize the main types of errors affecting photonic platforms. Then, we will briefly overview recent proposals for QEC strategies in photonic architectures.

D.1 Errors in photonic platforms

The errors commonly occurring in photonic quantum computing platforms can be grouped into three kinds:

- Generation errors

Most LOQC protocols rely on perfect single-photon sources, sources that lead to a periodic sequence of single detection events occurring in some finite counting time window. This property is known as *anti-bunching*. In practice, all photon sources present some level of photon bunching, leading to errors due to an unexpected number of photons. Another important characteristic of a photon source is the *indistinguishability* of the emitted photons, which should be identical in spatial and temporal mode, frequency, polarization, and so on. Partial distinguishability affects the visibility of interference fringes, leading to errors in the implementation of gates.

- Circuit errors

The non-zero absorption coefficient of optical components can lead to photon loss, one of the most important errors affecting photonic platforms. Furthermore, photons can accumulate errors when traversing an optical circuit. One such example is *mode mismatching*, which occurs when non-identical wave packets are interfered together. These errors can be caused by imperfections in optical components, and are exacerbated by non-ideal photon sources. Yet another source of error is related to classical feed-forwarding processes. These typically involve read-out of a detector, classical processing and conditional switching/reconfiguration of optical components. Optical switches are quite lossy, and if, on top of that, they are slow, they require long-term quantum memories, which

themselves might need additional feed-forwarding. This is an especially important source of error in LOQC platforms, considering the fact that entangling gates are inherently probabilistic and thus require feed-forwarding of heralded outcomes. Finally, quantum memories themselves are fallible. In photonics, quantum memories typically consist of fiber-optical loops that delay the photons. If photons need to be stored for a prolonged amount of time in a fiber, there is a significant chance that they are lost. Besides losing the quantum information, quantum memories might also introduce errors themselves, such as decoherence or mode mismatching.

- Detection errors

Non-ideal detectors can give rise to two error types: *photon loss* errors, where the detector counts fewer photons than the expected number of signal photons.³, and *dark counts*, where the detector counts more photons than were actually present. Photon loss is encapsulated in the *quantum efficiency* $\eta \in [0, 1]$ of the detector, operationally defined as the probability that a single photon input state will result in a count. Conversely, dark counts are defined as the probability that a vacuum input state will result in a count. A standard model for photon loss is to introduce a beam splitter with transmission coefficient η just before a perfect detector, where reflected modes are considered lost.

D.2 Quantum-error correction in photonics

Shortly after the earliest QEC proposals, Kitaev introduced the toric code [48], and the related surface code was developed [49]. Surface codes are currently the leading approach for fault-tolerant quantum computation, due to their high error correction threshold and planar layout [34]. A recent implementation by Google Quantum AI and Collaborators demonstrated below threshold operation and exponential error suppression in a superconducting processor using the surface code [50]. Bartolucci et al. considered its performance when implemented in the FBQC architecture [5] (section B.2.3). They numerically verified its error-correcting capability in a certain region of physical error parameters, for a 4-qubit star and a 6-qubit ring entangled states. They similarly demonstrated a significant photon loss tolerance of 10.4% when using the simpler (2,2)-Shor error-correcting code. Further results on this architecture can be found in [13], which introduces a method to enhance both the fusion success probability and photon loss tolerance.

³The "expected number of signal photons" only has a clear meaning when dealing with photon number states; however, it is still possible to give a more general definition of this term when applied to superpositions of different photon number states [47].

In [11], the performance of a surface code was also numerically investigated in the context of the SPOQC architecture (section C.1), under an error model tailored to the spin-optical system. It reached similar photon loss tolerance to the one exhibited by FBQC model when using the (2,2)-Shor encoding, and an even better performance when using the 4- and 6-qubit states. One advantage of this architecture is that the non-local links between data qubits, mediated by the spin-entangled photons, enable the implementation of low-density, parity check codes, known to have high encoding rate and low footprint, requiring drastically less physical qubits when compared to platforms relying solely on local interactions [51]. More recently, a new family of QEC codes has been developed, known as Floquet codes [52]. Surface codes require the measurement of high weight (typically weight-4) stabilizer Pauli operators, which can be challenging to implement in practice. Floquet codes improve on this drawback by instead performing periodic sequences of weight-2 measurements, dynamically generating effective stabilizers of higher weight over time. Hilaire et al. recently put forward a comparative analysis between the surface code and the so-called honeycomb Floquet code when implemented in the SPOQC architecture [12]. They not only show that the surface code is outperformed by the Floquet code, but demonstrated what is, to date, the highest reported photon loss threshold for photonic platforms without large-scale multiplexing. Floquet codes thus present themselves as a very interesting approach to fault-tolerance in photonic architectures.

Finally, we would like to mention a complementary approach to QEC, that of quantum error *mitigation*. Error mitigation techniques act as a first layer of protection against errors, acting directly on the physical qubits in an attempt to achieve immediate improvements despite hardware noise. Several such techniques have been proposed in the literature, ranging from circuit compilation tricks, randomization, and (counter-intuitively) even increasing error in order to extrapolate to the zero-noise limit [53]. Some of these error mitigation techniques have been demonstrated recently in different platforms [54, 55]. Investigating the use of these techniques in photonic platforms is an interesting open problem with relevance to near-term devices.

III Research Proposal

Throughout chapter II, we have characterized the main features of photonic systems that make them attractive as quantum computing platforms, as well as the main difficulties they pose to the development of scalable photonic quantum computing architectures. Photons will most likely play the essential role of carriers of quantum information over long distances, in communication and distribution tasks and as mediators of interactions between stationary systems. Nevertheless, despite significant theoretical and

technical developments in the field, experimental challenges still stand in the way of universal, large-scale photonic architectures, such as the inherently non-deterministic nature of photonic entangling gates, the sensitivity to photon loss, and resource intensive multiplexing and active-feed-forwarding requirements. We have introduced recent, promising proposals for new models which overcome some of these issues, with preliminary, theoretical demonstrations of their fault-tolerance and potential for scalability. It is expected that small proof-of-principle experimental implementations will soon emerge in the literature. But their modular building blocks can potentially still be improved upon, and open questions remain regarding the optimal strategies for incorporating error-correction protocols. The overarching goal of the hereby proposed research project, to be conducted within the Galvão research group at the International Iberian Nanotechnology Laboratory, is to contribute to the development and improvement of recent proposals for scalable photonic quantum computing architectures. The present essay has helped to identify key research questions and gaps in knowledge on current schemes and approaches, which will be the target of our exploration.

An initial, near-term objective is to continue the thorough literature review we have been conducting in the context of elaborating this essay. We will study in depth two main proposals, which are among the most promising approaches for scalable photonic quantum computing architectures: the fusion-based quantum computing (FBQC) [5] and spin-optical quantum computing (SPOQC) [11] architectures, reviewed in chapter II. We will work out the function of each modular unit and how the different modules interact, and appropriate figures of merit quantifying their performance. This will help us to more precisely identify weak points where there is room for improvement. In parallel, we will study the different kinds of quantum error-correction (QEC) codes that are widely being used in the literature. Namely, we will learn the formal description in the stabilizer formalism of different kinds of surface codes [49] and Floquet codes [52], and investigate proposals for their implementation in different architectures [5, 13, 11, 12], identifying what are currently the most promising strategies. Of particular interest to the SPOQC model are low-density parity check codes [34], enabled by the non-local connectivity of the qubits in this architecture, which have a higher encoding rate and lower resource footprint.

This critical review will be accompanied by background reading and learning about key relevant analytical methods, such as:

- The measurement-based quantum computing model (MBQC), which inspired the FBQC approach.
- The stabilizer formalism and how it is used to describe the MBQC model and quantum error-correcting codes.

- Basics of the theory of contextuality, coherence, and multiphoton indistinguishability, as these have all been identified as resources for photonic quantum computation.
- Recent developments in the theory of photonic quantum computation, in particular: description of linear optics using permanents and other matrix functions, theory of unitary invariants and their use to describe multiphoton indistinguishability.

Another near-term objective will be to assess the possibility of using Salih et al.'s [36] counterfactual gate in a blind computation protocol [43], as described in chapter II. Blind quantum computation is a promising, near-term approach to quantum computation, as it would allow secure quantum computation in a client-server setting. Quantum resources and computational power could in this way be centralized in a powerful server operating in specialized facilities, which clients could then access while keeping their data hidden from the server. We believe it might be possible to enact the Salih et al. gate in a secure way in such a setting, and aim at collaborating with blind quantum computing and quantum cryptography expert Anne Broadbent, a partner of our group in EU-funded project FoQaCia, on assessing the security of such a protocol. If security is verified, we will aim at proposing a protocol and physical architecture for its implementation, with an analysis on required resources. In case the resources are currently feasible, we will propose experimental tests to be implemented with experimental partners of the group, enabled by error mitigation techniques adapted to this photonic setting.

A longer-term objective will be to investigate how we might improve the performance of the modular building blocks of the FBQC and SPOQC architectures we identify in the initial review. Regarding the FBQC model, we will more carefully analyze state-of-the-art schemes for photonic entangled resource state generation and growth, such as Bell-state and GHZ-state generators [6], boosted fusion gates [7, 8] and percolation schemes [29, 10]. We will more precisely quantify associated resource overheads and trade-offs. For instance, the trade-off between higher multiplexing and quantum memory requirements when using less efficient sources and fusion gates vs more resources needed to implement higher efficiency ones (in terms of number of auxiliary optical modes, optical components, feed-forwarding and switching). Another example is the trade-off between using high-quality, but harder to implement quantum memories vs needing faster classical control to quickly transfer information to newly initialized flying qubits. We might also propose new and enhanced schemes, with better efficiency or smaller resource footprint. Preliminary tests of proposed schemes can potentially be carried out using the cloud-accessible integrated photonic chip made available by French company Quandela via the Quandela Cloud service [56], which the student already learned how to use. When it comes to the SPOQC architecture, we will more thoroughly evaluate the effi-

ciency of the deterministic quantum emitter sources of photonic entangled states, as a promising alternative to the inherently probabilistic all-photonic entangling gates. We will also assess resource requirements and performance of the linear-optical repeat-until-success gates mediating the spin interactions, and how they compare to other known schemes.

Another longer-term objective is to identify optimal strategies for implementing QEC in the different architectures. An important tool for assessing the performance of the aforementioned codes will be their numerical simulation, as was done in recent proposals [5, 11, 12, 13]. This will allow us to preliminarily estimate their physical error thresholds and logical error suppression rates under different code distances and decoders, subject to error models tailored specifically for photonic platforms. The student already knows how to use Stim [57], a powerful tool for the fast simulation of such stabilizer codes. The aforementioned proposals only assess error thresholds for codes preserving a logical state over many rounds. We aim to go one step further, by estimating the error threshold under implementation of logical gates as well. Informed by our simulation work, we will be better equipped to evaluate the strengths and weaknesses of different strategies, when applied to the FBQC and SPOQC architectures. It will also allow us to optimize the choice of starting resource state, used for the encoding of logical information, and the choice of code more suitable for each architecture. For example, in [12], a 6-qubit ring state was investigated in a Floquet code implementation, and we will aim to expand on these results by assessing the performance of the 4-qubit states of Fig. 2(a), as well as the smallest theoretically viable option of a 3-qubit GHZ state. We will consider resource trade-offs in the choice of encoding states, where larger entangled states allow for more robust encoding, but require more resources and multiplexing to generate. Work on this topic is currently being pursued by Quandela, a partner of the group in project EPIQUE, and possible collaboration with their team is envisioned. Finally, an alternative path to explore is that of quantum error mitigation [53], a complementary approach to QEC that aims at achieving immediate improvements in existing, noisy hardware with lower technical requirements. This avenue is currently being explored in different platforms [54, 55], and is relevant in the development of near-term photonic processors.

This research project will contribute to the development of both near-term quantum computing models such as the blind quantum computing scheme, as well as longer-term, scalable architectures for photonic quantum computing. The Galvão group has a sizeable component of research in this field, supported by the ERC Advanced grant “Quantum advantage via non-linear boson sampling” (QU-BOSS), and Horizon Europe project “European photonic quantum computer” (EPIQUE). This involvement in several projects provides many opportunities for workshops, seminars and scientific exchanges and collaborations. The group’s expertise is theoretical, with strong collaborations with experimental groups, notably with the

Quantum Lab of Sapienza University of Rome (led by Fabio Sciarrino), and the Milan Polytechnic group (led by Roberto Osellame), as well as quantum technology startups like Quandela, QuiX and Veriqloud. These active collaboration will give us ample opportunities to implement proof-of-principle demonstrations of our schemes.

Besides the PhD thesis itself, the main deliverables will be the publication of academic papers describing original research results and presentations in quantum information and quantum optics conferences. We plan to make available any code developed throughout the project in an open-source repository such as GitHub.

References

- [1] E. Knill, R. Laflamme, and G. J. Milburn. *Nature* 409 (2001), pp. 46–52.
- [2] J. L. O’Brien. *Science* 318.5856 (2007), pp. 1567–1570.
- [3] R. Raussendorf and H. J. Briegel. *Physical Review Letters* 86.22 (2001), pp. 5188–5191.
- [4] D. E. Browne and T. Rudolph. *Phys. Rev. Lett.* 95.1 (2005), p. 010501.
- [5] S. Bartolucci et al. *Nature Communications* 14.1 (2023).
- [6] S. Bartolucci et al. 2021. URL: <http://arxiv.org/abs/2106.13825>. Pre-published.
- [7] F. Ewert and P. Van Loock. *Physical Review Letters* 113.14 (2014), p. 140403.
- [8] T. Yamazaki, R. Ikuta, and T. Yamamoto. 2023. URL: <http://arxiv.org/abs/2301.06551>. Pre-published.
- [9] K. Kieling, T. Rudolph, and J. Eisert. *Physical Review Letters* 99.13 (2007), p. 130501.
- [10] M. Gimeno-Segovia et al. *Physical Review Letters* 115.2 (2015), p. 020502.
- [11] G. de Gliniasty et al. 2023. URL: <http://arxiv.org/abs/2311.05605>. Pre-published.
- [12] P. Hilaire et al. 2024. URL: <http://arxiv.org/abs/2410.07065>. Pre-published.
- [13] W. Song et al. *Physical Review Letters* 133.5 (2024), p. 050605.
- [14] P. Thomas et al. *Nature* 608.7924 (2022), pp. 677–681.
- [15] M. AbuGhanem. *Frontiers of Optoelectronics* 17.1 (2024), p. 33.
- [16] K. Fukui and S. Takeda. *Journal of Physics B: Atomic, Molecular and Optical Physics* 55.1 (2022), p. 012001.

-
- [17] J. Wang et al. *Nat. Photonics* 14 (2020), pp. 273–284.
- [18] S. Barz. *Journal of Physics B: Atomic, Molecular and Optical Physics* 48.8 (2015), p. 083001.
- [19] D. Gottesman and I. L. Chuang. *Nature* 402.6760 (1999), pp. 390–393.
- [20] S. Popescu. *Physical Review Letters* 99.25 (2007), p. 250501.
- [21] F. Verstraete and J. I. Cirac. *Physical Review A: Atomic, Molecular, and Optical Physics* 70.6 (2004), p. 060302.
- [22] M. Hein et al. 2006. URL: <http://arxiv.org/abs/quant-ph/0602096>. Pre-published.
- [23] D. Gottesman. 1998. URL: <http://arxiv.org/abs/quant-ph/9807006>. Pre-published.
- [24] M. A. Nielsen. *Physical Review Letters* 93.4 (2004), p. 040503.
- [25] C. K. Hong, Z. Y. Ou, and L. Mandel. *Phys. Rev. Lett.* 59.18 (1987), pp. 2044–2046.
- [26] W. P. Grice. *Physical Review A* 84.4 (2011), p. 042331.
- [27] F. V. Gubarev et al. *Physical Review A: Atomic, Molecular, and Optical Physics* 102.1 (2020), p. 012604.
- [28] S. A. Fldzhyan, M. Y. Saygin, and S. P. Kulik. *Phys. Rev. Res.* 3.4 (2021), p. 043031.
- [29] S. Morley-Short et al. *Quantum Science and Technology* 3.1 (2018), p. 015005.
- [30] S. Bartolucci et al. 2021. URL: <http://arxiv.org/abs/2109.13760>. Pre-published.
- [31] H. Bombin et al. 2021. URL: <https://arxiv.org/abs/2103.08612>. Pre-published.
- [32] P. Senellart, G. Solomon, and A. White. *Nat. Nanotech.* 12 (2017), p. 1062.
- [33] N. H. Lindner and T. Rudolph. 2008. URL: <http://arxiv.org/abs/0810.2587>. Pre-published.
- [34] N. P. Breuckmann and J. N. Eberhardt. *PRX Quantum* 2.4 (2021), p. 040101.
- [35] M. C. Löbl, S. Paesani, and A. S. Sørensen. *Quantum* 8 (2024), p. 1302.
- [36] H. Salih et al. *New Journal of Physics* 23.1 (2021).
- [37] P. Kwiat et al. *Phys. Rev. Lett.* 74.24 (1995), pp. 4763–4766.
- [38] W. M. Itano et al. *Phys. Rev. A* 41.5 (1990), pp. 2295–2300.
- [39] A. G. White et al. *Phys. Rev. A* 58.1 (1998), pp. 605–613.
- [40] Y. Yang et al. *npj Quantum Information* 9.1 (2023).
- [41] T.-G. Noh. *Phys. Rev. Lett.* 103.23 (2009), p. 230501.

-
- [42] H. Salih et al. *Phys. Rev. Lett.* 110.17 (2013), p. 170502.
- [43] A. Broadbent, J. Fitzsimons, and E. Kashefi. 2008.
- [44] P. W. Shor. *Physical Review A: Atomic, Molecular, and Optical Physics* 52.4 (1995), R2493–R2496.
- [45] S. J. Devitt, W. J. Munro, and K. Nemoto. *Reports on Progress in Physics* 76.7 (2013), p. 076001.
- [46] E. Knill, R. Laflamme, and Z. Wojciech. *Proc. R. Soc. Lond. A.* 454 (), pp. 365–384.
- [47] P. Kok et al. *Reviews of Modern Physics* 79.1 (2007), pp. 135–174.
- [48] A. Yu. Kitaev. *Annals of Physics* 303.1 (2003), pp. 2–30.
- [49] M. H. Freedman and D. A. Meyer. *Foundations of Computational Mathematics* 1.3 (2001), pp. 325–332.
- [50] R. Acharya et al. 2024. URL: <http://arxiv.org/abs/2408.13687>. Pre-published.
- [51] P. Panteleev and G. Kalachev. *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing. Rome Italy: ACM, 2022, pp. 375–388.
- [52] M. B. Hastings and J. Haah. 2021. URL: <http://arxiv.org/abs/2107.02194>. Pre-published.
- [53] Z. Cai et al. *Reviews of Modern Physics* 95.4 (2023), p. 045005.
- [54] A. Salavrakos et al. URL: <https://arxiv.org/abs/2405.02277v1>. Pre-published.
- [55] J. D. Guimarães et al. *PRX Quantum* 4.4 (2023), p. 040329.
- [56] N. Maring et al. *Nature Photonics* 18.6 (2024), pp. 603–609.
- [57] C. Gidney. *Quantum* 5 (2021), p. 497.