

Resource-Constrained Entanglement Distribution and Quantum Compilation

*Distribution d'Intrication et Compilation Quantique
sous Contrainte de Ressources*

Thèse de doctorat de l'université Paris-Saclay

École doctorale n°564: Physique en Île-de-France (PIF)

Spécialité de doctorat: Physique

Graduate School : Physique. Référent : Faculté des sciences d'Orsay

Thèse préparée dans la (ou les) unité(s) de recherche **Institut de Physique Théorique
(Université Paris-Saclay, CNRS, CEA)**, sous la direction de **Nicolas SANGOUARD**

Thèse soutenue à Paris-Saclay, le 01 septembre 2026, par

Bastien GRIVET

Composition du jury

Membres du jury avec voix délibérative

Sébastien Tanzilli

Directeur de recherche, INPHYNI, CNRS, Université Côte d'Azur

Examineur

Mario Krenn

Professeur, Université de Tübingen

Rapporteur & Examineur

Wolfgang Dür

Professeur, Université d'Innsbruck

Rapporteur & Examineur

Hanna Le Jeannic

Chargée de recherche, Laboratoire Kastler Brossel, CNRS, Sorbonne Université

Examinatrice

Title: Resource-Constrained Entanglement Distribution and Quantum Compilation

Keywords: Quantum Information, Quantum communication, Quantum Computation, Quantum Optics

Abstract: Once confined to theoretical proposals, quantum technologies are now approaching practical use. Such a perspective calls for the study of realistic architectures under the limitations of physical implementations. This thesis explores two resource-constrained designs at opposite ends of the quantum information stack: distributing long-distance entanglement in the presence of channel losses, and compiling quantum operations which rely on costly gates. Both share a common goal: optimizing scarce resources to enable near-term implementations of quantum architectures.

In the perspective of a quantum internet implementation, the first part of this manuscript addresses continent-scale entanglement distribution over lossy channels. As a first contribution, we propose an architecture interfacing metropolitan gateways with the continent-scale backbone, addressing the heterogeneity between two physical platforms. The second contribution of this work focuses on the derivation of an exact, non-perturbative model of a DLCZ-type elementary link, valid across the full pump-power range. This model

exposes the fundamental trade-off between distribution rate and entanglement quality. Finally, we introduce an entanglement witness relying on local operations and classical communication, certifying entanglement robustly despite losses.

In the context of fault-tolerant quantum computation, the second part of this manuscript provides an exact quantum circuit synthesis method, where minimizing non-Clifford gates is essential for the realistic implementation of quantum algorithms. Framing quantum circuit synthesis as a single-player game, we develop a reinforcement learning method inspired by AlphaZero, and benchmark the resulting agent on the exact synthesis of costly subroutines. Extending to dynamic circuits, i.e. circuits with ancillary qubits and mid-circuit measurements, the agent finds new decompositions involving fewer non-Clifford gates than standard approaches.

Together, these contributions illustrate how resource optimization at the lowest levels can support the architectures of future quantum networks and computers.

Titre: Distribution d'Intrication et Compilation Quantique sous Contrainte de Ressources

Mots clés: Information Quantique, Communication Quantique, Calcul Quantique, Optique Quantique

Résumé: Longtemps restées à l'état de propositions théoriques, les technologies quantiques se rapprochent aujourd'hui d'applications concrètes. Une telle perspective motive l'étude d'architectures réalistes, tenant compte des contraintes de leur mise en œuvre physique. Cette thèse étudie deux de ces problèmes : distribuer de l'intrication sur de longues distances malgré les pertes de transmission, et compiler des opérations quantiques en limitant les portes logiques coûteuses. Dans les deux cas, l'objectif principal est de tirer le meilleur parti de ressources rares en vue d'une implémentation à court terme.

Dans la perspective du déploiement d'un internet quantique, la première partie de ce manuscrit étudie la distribution d'intrication à l'échelle continentale. La première contribution de cette étude propose une architecture interfaçant réseaux métropolitains et réseau continental, sous contrainte d'hétérogénéité entre ces deux plateformes. La deuxième contribution développe un modèle exact non perturbatif d'un lien élémentaire de type DLCZ, valable sur toute la plage de puissance de pompe. Ce modèle met en évidence le compromis fondamental entre taux de distribution

et qualité de l'intrication. Enfin, ce manuscrit développe un témoin d'intrication basé sur des opérations locales et une communication classique, capable de certifier l'intrication de manière robuste malgré les pertes des mémoires.

Dans un contexte de calcul quantique tolérant aux erreurs, la seconde partie de ce manuscrit porte sur la synthèse de circuits quantiques, où le nombre de portes non-Clifford conditionne la faisabilité des algorithmes. En formulant la synthèse de circuits quantiques comme un jeu à un joueur, ce manuscrit présente une méthode d'apprentissage par renforcement inspirée d'AlphaZero. L'agent ainsi entraîné réalise la synthèse exacte de sous-routines coûteuses. Étendu aux circuits dynamiques, c'est-à-dire des circuits avec qubits auxiliaires et mesures intermédiaires, l'agent s'est révélé capable de trouver de nouvelles synthèses invoquant un nombre de portes non-Clifford plus faible que les architectures standards.

Ces deux travaux montrent qu'une gestion fine des ressources aux niveaux les plus élémentaires peut nourrir les architectures des réseaux et des calculateurs quantiques à venir.

Contents

1	Introduction	11
I	Large Quantum Networks	23
1	Quantum Repeater for long-distance entanglement distribution	29
1.1	Generating entanglement between two distant hubs	29
1.1.1	Idealized low-pump regime	29
1.1.2	Perturbative expansion in the pump parameter	33
1.1.3	Practical example with the two-photon truncation	37
1.2	Scaling quantum repeater	40
1.2.1	The fibre-loss barrier	40
1.2.2	Example with 2 repeater links	41
1.2.3	Multi-link quantum repeater	45
2	Interfacing the long-distance edge with the metropolitan gateway	47
2.1	The ion-swap: bridging backbone and metropolitan network	47
2.1.1	The network protocol	48
2.1.2	Practical example with a 500 km link	51
2.1.3	Numerical study	55
2.1.4	Alternative protocol	57
2.2	Matching photon waveforms	58
2.2.1	The shape-mismatch problem	58
2.2.2	Analytical derivation	59
2.2.3	Results	62
2.2.4	Mixture	62
2.3	Multiplexing	64
2.3.1	Stretching and multiplexing	64
2.3.2	Qualitative study	66
2.3.3	Numerical study	68
3	Exact modeling of an elementary link	71

3.1	Exact computation	71
3.1.1	State shared by Alice and Bob	71
3.1.2	Local detection probabilities	79
3.1.3	Coherence Evaluation	81
3.2	Experimental fit	83
3.2.1	Local probabilities	83
3.2.2	Single-photon interferometry	87
3.3	Entanglement Witness	90
3.3.1	Non-classicality	92
3.3.2	Constructing the Witness	95

II Quantum Circuit Synthesis 107

4 A Game of Unitary Synthesis 111

4.1	Framing synthesis as a single-player game	111
4.2	Reinforcement Learning	112
4.3	Reinforcement Learning for Quantum Circuit Synthesis	113

5 AlphaZero as a long-term strategy agent 119

5.1	Long-term strategies for quantum-circuit synthesis	119
5.2	Monte Carlo Tree Search	121
5.2.1	The exploration-exploitation dilemma	121
5.2.2	The Monte Carlo Tree Search Algorithm	124
5.2.3	Monte Carlo Tree Search as a Reinforcement Learning agent	128
5.3	AlphaZero	129
5.3.1	Guiding MCTS with a Neural Network	129
5.3.2	AlphaZero policy	131
5.3.3	Training AlphaZero	132

6 Results and Perspectives 137

6.1	Learning pipeline	137
6.2	Numerical implementation	138
6.3	Results	141
6.3.1	Quantitative benchmark	141
6.3.2	Qualitative benchmark	143

6.3.3	Human feed-back	143
6.4	Perspectives	144
6.4.1	Limits	144
6.4.2	Improvements	145
7	Conclusion	151
III	Appendix - Quantum Networks	155
A	Helpers	157
A.1	Non-photon-number-resolving detector	157
A.2	Trace of thermal states	158
A.3	Integral calculation	161
B	Perturbative study of an elementary link	163
B.1	Block-diagonal structure of the heralded state	163
B.2	Elementary link state expansion	166
B.3	One- and two-photon contributions to an elementary link	171
B.4	Multi-link swapping	175
C	Study of the full repeater link	179
C.1	Distribution time	179
C.2	Strategy computation for the 4-clicks protocol	181
C.3	Modelling temporal-mode mismatch	184
C.4	Shaping mismatch	186
C.5	Shape-independent construction of the photodetection POVM	193
D	Exact study of an elementary link	195
D.1	Derivation of the state	195
D.2	Fidelity of the elementary link to a Bell pair	199
D.3	Derivation of the displaced local probabilities	201
D.4	Useful coherent states indentities	206
D.5	Perturbative coincidence probability	208
D.6	Exact coincidence probability	210
D.7	Matrix elements of the displacement-based observable	214
D.8	Entanglement witness in the qubit subspace	218

D.9 Entanglement witness beyond the qubit subspace	221
D.10 Bounding multi-photon marginals via twofold coincidences	226

IV Appendix - Quantum Circuit Synthesis 229

D.11 From Markov Decision Problem to Reinforcement Learning	231
---	-----

1 -Introduction

Building the Backbone of Quantum Networks

Motivated by applications in quantum cryptography, long-distance entanglement distribution has rapidly emerged as a promising technique to enable near-term quantum technologies. As a central example, protocols based on Quantum Key Distribution (QKD) leverage the non-local properties of entangled states: when identical measurements are performed on the two halves of a bipartite maximally entangled state, the results are perfectly correlated yet fundamentally unpredictable by any third party [1, 2]. Therefore, keys resulting from an accumulation of measurement outcomes can be used to secure communications with guarantees that do not rely on computational hardness assumptions [3]. Such a physically grounded defence is vital in a world where computers, powered by quantum advantages using Shor’s algorithm [4, 5], could compromise the security of current classical cryptographic systems.

Realizing such long-distance communication requires overcoming a fundamental physical barrier: optical fiber attenuation. Relying on single-photon propagation, entanglement distribution rate is directly limited by the channel transmissivity. Standard telecom fibers exhibit an attenuation of $\alpha = 0.2 \text{ dB/km}$. While this loss is manageable over metropolitan distances, its exponential scaling makes direct long-distance transmission impractical: a $L = 1000 \text{ km}$ link results in a transmission probability of $\eta = 10^{-\alpha \cdot \frac{L}{10}} = 10^{-20}$, effectively reducing the distribution rate to zero. This loss limitation has actually been quantified in [6] with fundamental arguments using relative entropy: the authors showed that, even assuming unlimited local operations and classical communication, the capacity of a lossy bosonic channel is upper-bounded by $C_{\max} = -\log_2(1 - \eta) \approx 1.44 \eta$. This loss-induced limitation motivates the study of quantum repeater based architectures as promising candidates for long-distance quantum communication.

The foundation of quantum repeaters architecture relies on breaking the assumption that entangled photons must originate from the same physical source: [7] shows that Bell inequality violations can be achieved using two independent and spatially separated sources. Inspired by this proposal, [8] highlighted the ability to engineer entanglement between distant nodes that have never interacted. This fundamental concept, known as entanglement swapping, relies on a Bell state measurement and serves as the foundation for quantum repeater architectures, enabling the segmentation of long-distance communication channel into shorter independent elementary sections.

Building upon this concept, the quantum repeater architecture relies on the notion of elementary links. Typically spanning kilometres-scale distances, an elementary link consists of two source nodes, a central station performing Bell state measurement for

entanglement generation, and edges quantum memories storing the heralded state. In this context, an elementary link is considered loaded once the entangled state is successfully heralded and stored. The overall distribution then follows a hierarchical divide-and-conquer strategy: entanglement is generated across independent links in parallel, stored, and subsequently extended over larger distances by performing Bell state measurements on adjacent nodes. As discussed later in this thesis, parallelizing the generation process mitigates the exponential losses inherent to direct fiber transmission, yielding an exponential rate advantage that surpasses the quantum channel capacity limit of direct transmission [6].

The efforts of the research community have led to a standardized classification of quantum repeater protocols into three distinct generations [9].

- The **first generation** (1G) represents the most near-term architecture, closely following the elementary link model previously described. In this framework, elementary links are loaded via heralded entanglement generation where the central station signals a successful Bell state measurement. To mitigate operational errors, 1G repeaters rely on entanglement purification, where multiple noisy pairs are consumed to produce a single higher-fidelity state. Consequently, the process is doubly probabilistic: it requires the simultaneous success of initial heralded entanglement distribution and subsequent purification steps, necessitating numerous trials to achieve a single high quality entangled state. In this context, the Bell pair distribution rate is fundamentally limited by the requirement of two-way classical communications, for both heralding feedback and local operations with classical communication required for purification. This reliance on classical signaling, intrinsic to the probabilistic nature of the protocol, creates a bottleneck where nodes must wait for confirmation signals before proceeding, constraining the overall distribution rate [10–13].
- The **second generation** (2G) overcomes the limitations rate by integrating local quantum information processing. Unlike 1G architectures, 2G repeaters exploit Quantum Error Correction (QEC) to locally correct operation errors. By shifting from purification to local QEC, the need of classical communication for the distillation process is removed. However the 2G still relies on probabilistic heralding of entanglement distribution, maintaining a dependence on two-way classical feedback for link confirmation. While this approach reduces partially the constraints of classical signalling, it requires significantly more sophisticated hardware capable of performing high-fidelity local logical gates [14, 15].
- Finally, the **third generation** (3G) adopts a fundamentally different approach, abandoning heralded entanglement distribution in favour of QEC codes which directly correct losses. In this paradigm, the source node transmits "logical flying

qubits”: a stabilized logical state encoded across a multi-qubit wave packet. Each intermediary station uses local QEC protocols to generate entanglement and correct operational errors on the fly. By entirely eliminating the requirement for two-way classical feedback, 3G repeaters enable high-rate quantum communication conceptually analogous to classical amplifiers. However, such a protocol requires ultra-high communication efficiency: if the transmission losses exceed 50% the non-cloning theorem claims quantum information becomes fundamentally unrecoverable [16–19].

At the time of writing, the experimental quantum repeater community is primarily focused on the realization of first-generation architectures. This generation was introduced by Briegel *et al.* [10] and a practical implementation using atomic ensembles and linear optics was then proposed by Duan, Lukin, Cirac, and Zoller [12]. The latter, now widely known as the DLCZ protocol, brings a built-in purification step, avoiding the necessity of multi-copies for nested purification proposed by Briegel *et al.*. This was the starting point of simple quantum repeater architectures, aiming to beat the direct transmission of photons through fibers [13].

Nowadays, various physical platforms for 1G protocols implementation are being explored [20]. A viable technology requires long coherence times to store states during classical feedback, efficient light-matter interface with indistinguishable emission capabilities for high-fidelity entanglement swapping, telecom-wavelength compatibility to minimize fiber attenuation, and finally multiplexing capabilities to boost entanglement distribution rates. Two main technological families have emerged to meet these requirements: atomic ensembles and individually controlled systems.

- **Atomic ensembles** rely on the collective light-matter interaction of multiple quantum systems, such as cold atomic gas or solid-state ensembles, enabling either the generation of photon pairs with a controllable delay [12] or the storage of single photons from parametric sources with on-demand retrieval [21–25]. The collective behaviour benefits from a strong light-matter coupling and naturally supports multiplexing across multiple degrees of freedom, providing a linear speedup for the entanglement distribution rate with respect to the number of modes [23]. However, as it will be explained in this manuscript, this approach based on probabilistic emissions requires operating at low pump power, which directly limits the distribution rate. Historically, cold atomic ensembles have been a dominant platform for early experimental realizations of remote atomic ensemble entanglement, notably demonstrated in [26, 27]. This technology remains actively used, with recent experiments reporting heralded entanglement of atomic cloud separated by 12.5 km physical distance [28], and over 50 km of optical fiber [29]. On the other hand, rare-earth-doped crystals leverage extreme multimode multiplexing to

compensate low rate from low pump power [30]. Recent works have demonstrated high efficiency on-demand spin-wave memories in such systems [31–33], leading to the realization of entanglement [34] and on-demand spin-wave entanglement demonstration within a 1G repeater architecture [35].

- **Individually controlled systems** rely on single quantum systems for repeater implementations, such as trapped ions, neutral atoms, quantum dots, or color centers. Two main features make these individually controlled systems attractive for implementing first-generation quantum repeaters. First, their high level of control enables the deterministic implementation of entanglement swapping operations [36], which contrast with atomic-ensemble based repeaters inherently probabilistic. Second, they allow the generation of single photons with high purity [37]. However, single quantum system traditionally suffer from a lack of light-matter coupling [38]. Recent progress has partially addressed the former: colored centers leverage nanophotonic integration to improve photon collection efficiency, improving the experimental implementation of remote entanglement between crystal defects [39–41]. Trapped ions benefits from high photon indistinguishability and strong isolation [42–45]. Single neutral atoms have recently achieved major long-distance progress, including entanglement over 33 km of telecom fiber [46], and very recently device-independent QKD with secret key generation at 0.1 bits/s over 11 km [47].

Beyond establishing secure cryptographic keys, entanglement serves as a universal resource for quantum information processing. This motivates a broader architectural vision: the Quantum Internet. Rather than developing point-to-point bipartite links, a quantum network operates as a web of quantum channels, distributing Bell pairs between remote interconnected nodes capable of local quantum processing. This paradigm unlocks a wide range of applications [38, 48], including distributed quantum computing [49], clock synchronization [50], or long-distance quantum sensing [51] such as astronomical interferometry with telescope arrays [52].

In this context of large-scale quantum networks, international collaborations have emerged to coordinate global quantum internet deployment. This thesis is developed within one such initiative: the Quantum Internet Alliance (QIA) [53]. This European project aims to construct continent-scale quantum networks, where metropolitan-scale networks are interconnected with first-generation long-distance quantum repeaters by 2030. The full stack, from physical hardware to software protocols, is being developed collaboratively across multiple research institutions in EU member states.

Specifically, the QIA project structures its network into two complementary layers: a long-distance backbone distributing entanglement over continental scales via a quantum repeater-based architecture, and a metropolitan network handling distribution at the city

scale. Acknowledging the distinct requirements of these layers, the QIA project assigns a dedicated technological approach to each. The sparse long-distance backbone, which requires high-rate entanglement distribution, relies on atomic ensemble technologies with rare-earth-doped crystal memories, enabling large multiplexing capacities across thousands of modes. In contrast, the denser metropolitan network prioritizes individually controlled systems, based on trapped ions, for multi-node synchronization and local processing.

The scope of this thesis focuses on the physical implementation of the long-distance backbone, targeting entanglement distribution over hundreds of kilometers, including the interfaces with metropolitan nodes. A central work of this manuscript is the development of an exact non-perturbative model of an elementary link. This model enables the exploration of high-pump regimes for rate enhancement and provides experimental groups with a reliable framework for parameter calibration and state characterization.

The first contribution of this thesis is to provide theoretical support for the physical implementation of QIA networks. Specifically, we propose a long-distance backbone architecture based on dual repeater chains, each composed of two elementary links using rare-earth-doped crystals, connecting trapped-ion metropolitan edge nodes via dual-rail encoding. A key challenge of such a heterogeneous network lies at the backbone-to-metropolitan interface: rare-earth-doped crystals and trapped-ion systems operate with mismatched photon temporal profiles, differing by roughly one to two orders of magnitude. Based on a standard light storage protocol in rare-earth-doped memories, we enable temporal reshaping of single-photon pulses, providing a practical interface between metropolitan nodes and the quantum repeater backbone. In this thesis, we present a model of this hybrid interface, and quantify the impact of residual mode mismatch on both entanglement fidelity and distribution rate.

The second and central contribution of this thesis is the development of an exact model of a DLCZ-type elementary link. In such links photon-pair sources operate in low-pump regime, avoiding multi-photon contributions which degrades the entanglement fidelity. While providing a simple analytical model, the low-pump approximation breaks down in relevant regimes where high pump powers is required to achieve high entanglement distribution rates. This thesis derives an exact non-perturbative model of the elementary link state, valid across the full pump parameter range, exposing the fundamental trade-off between distribution rate and entanglement quality. In collaboration with the experimental group of *De Riedmatten* in Barcelona, we benchmark the predictions of our model against experimental data, providing a quantitative validation based on local photon detection statistics and visibility measurements.

Finally, this thesis introduces an entanglement witness designed to certify the proper operation of a DLCZ-type elementary link. While conventional witnesses rely on optical interference between remote nodes, such approaches become impractical over long

distances. This manuscript proposes a protocol based on local operations and classical communication for entanglement certification. Specifically, combining displacement operations with photon detection, this method enables robust entanglement certification under arbitrary memory losses within a qubit subspace. Two experimental groups, in Barcelona and Hefei, have expressed interest in testing this entanglement witness.

Quantum compiler for fault-tolerant quantum computation

While information is often treated as an abstract entity, its physical realization, even in classical systems, is inherently subject to noise and errors. For instance, due to thermal fluctuations, hardware degradation, or even cosmic rays [54], modern flash memories must manage bit-error rates as high as 10^{-2} [55]. To mitigate these errors, the computer science community has developed techniques called Error-Correcting Codes: by employing sophisticated redundancy methods, a single bit of information is encoded across multiple physical units [56, 57]. This allows efficient data recovery and ensures information integrity over years-long timescales [58]. The quantum community drew inspiration from these redundancy techniques to protect quantum states in quantum memories. However, due to the no-cloning theorem and measurement-induced disturbance, adapting classical error-correcting codes directly to quantum systems posed fundamental challenges [59].

Building on classical error correcting codes techniques, Shor, Steane, and Calderbank were the first to establish the foundations of Quantum Error Correction (QEC), introducing a class of codes known after their names: the Calderbank-Shor-Steane (CSS) codes [60–62]. Initially focused on protecting quantum states for long storage, QEC was later extended by Shor when he demonstrated that error protection could be extended to quantum operations themselves, shielding the quantum state throughout its noisy evolution [63]. This realization opened the path to Fault-Tolerant Quantum Computation (FTQC): a framework in which operations do not propagate logical errors and correction is applied repeatedly throughout the computation. Building on this, the community exploited algebraic structures to explore a broad landscape of QEC schemes, progressively converging toward practical and robust architectures for FTQC [64–66]. Among these, the surface code, a CSS code defined on a two-dimensional lattice of qubits, has emerged as the dominant candidate for modern QEC architectures, combining a high error threshold with nearest-neighbour connectivity that matches the constraints of physical hardware [67, 68].

Fault-tolerant quantum computation stands in fundamental contrast with the classical paradigm. In classical computing, error correction is almost exclusively dedicated to memory and data transmission: logic gates themselves can be considered noiseless in practice ¹. Therefore, the choice of elementary operations, known as the Instruction

¹In a few critical applications where errors cannot be tolerated, such as space engineering, light

Set Architecture (ISA), is driven by the targeted performance and efficiency²: x86 for high-performance workstations, ARM for energy-constrained mobile devices, or RISC-V for lightweight embedded systems. Error tolerance plays no role in this choice, as the hardware is assumed reliable. In quantum computing, this assumption breaks down: computation is as noisy as storage. Thus, the choice of elementary quantum operations, widely known as quantum gates set, and sometime called quantum ISA (or QISA) to emphasize the classical analogy [69–71], is instead primarily dictated by the constraints of the underlying quantum error correction framework.

In the quest for fault-tolerant quantum computing, Shor noticed in [63] that a specific set of operations, the Hadamard, CNOT, and $\pi/4$ phase gates, could be fault-tolerantly implemented within CSS codes: by acting bitwise on the physical qubits, these gates naturally prevent errors from propagating across the logical block. Building upon this work, Gottesman introduced the stabilizer formalism [72], which remains a bedrock of modern QEC architectures. Within this framework, he introduced the notion of transversal gates: operations that act independently on each physical qubit of a code block, thereby avoiding error propagation between qubits. Formalizing the fault-tolerant structure identified earlier by Shor, Gottesman showed that the entire Pauli group normalizer, known as the Clifford group [73], is transversal for well-chosen CSS codes, making all of these operations fault-tolerant at once.

Yet, while being fault-tolerant, the Clifford group does not generate the set of all unitaries, making the gate set incomplete for universal computation³. In this context, a universal gate set refers to the ability to approximate any unitary operation with arbitrary precision using a finite sequence of gates from this set. Here we understand the choice of elementary operations must simultaneously satisfy two constraints: fault-tolerance, ensuring the operations are compatible with the underlying QEC code, and universality, providing the ability to approximate any unitary operation. To bridge the gap, Boykin introduced in [74] a $\pi/8$ rotation gate, known as the T gate, completing the Clifford group for universality. However, the Eastin-Knill theorem established a fundamental limitation to this approach, stating that no binary QEC code capable of correcting arbitrary single-qubit errors can possess a universal set of transversal gates [75]. This result makes transversality and universality mutually exclusive, necessitating alternative strategies for universal fault-tolerant control. Building on these results, the community established a clear distinction between Clifford and non-Clifford operations, adopting a Clifford generator set extended by a single non-Clifford operation as the standard for universal

error detection methods with single-bit parity checks or redundant computations are sufficient.

²Actually political and historical reasons plays a non negligible role, e.g. CUDA ecosystem keeps NVIDIA dominant in GPU computing, or decades of backward compatibility have locked x86 into consumer software regardless of raw efficiency.

³Aware of the lack of universality, Shor proposed in [63] to reach universality by adding the Toffoli gate.

gate sets.

This distinction shifts the whole complexity to the implementation of non-Clifford operations. Such an implementation is achieved through a complex three-step protocol: magic state injection, magic state distillation, and gate teleportation [76, 77]. First, a multitude of noisy magic states, referring to special non-Clifford resource states, are prepared and injected into logical code blocks. Then, as they are inherently vulnerable to noise, these logical states are distilled, leading to a small number of high fidelity magic state. Finally, the resulting high-fidelity magic states are consumed via gate teleportation to apply the desired non-Clifford operation on the data qubits. This requirement for non-Clifford resources has led to the development of magic state factories: dedicated architectural blocks that consume a significant volume of physical qubits to distill high-fidelity magic states [67, 68, 71]. Despite community efforts to reduce this overhead [78, 79], non-Clifford resources remain the primary bottleneck in modern fault-tolerant designs. As a consequence, the requirement for non-Clifford gates has emerged as the central metric for assessing the practical feasibility of quantum algorithms [71]. In this thesis, we focus on the T gate as the non-Clifford primitive, as it constitutes the standard non-Clifford resource in surface code architectures [67, 68, 80, 81]. This perspective reflects a RISC-like quantum ISA, where Clifford gates are native and a single non-Clifford instruction suffices to promote the set to universality, as it is discussed in [80].

As discussed, the fault-tolerant settings impose a gate availability constraint that quantum software scientist must work within. Thus, the practical implementation of quantum algorithms on physical hardware requires translating abstract unitary operations into concrete sequences of those available gates. Known as quantum circuit synthesis⁴, this process finds a natural parallel with classical computing [82]. In classical computing, compilation bridges the gap between algorithmic descriptions and hardware-level instructions, translating abstract operations into a sequence of instructions supported by the processor’s ISA. A similar abstraction exists in quantum computing: arbitrary quantum evolutions, i.e. unitary transformations, must be decomposed into sequences of elementary gates drawn from a given gate set, with varying gate implementation costs, and under specific hardware constraints such as limited run-time, qubit connectivity or ancilla availability.

More precisely, for a given universal gate set, quantum circuit synthesis aims to find a sequence of gates implementing a target unitary. This task faces an inherently hard combinatorial landscape: calling respectively n and g the number of qubits and the gate-count, at each layer roughly $n \cdot g$ gate placements are available, thus the number of candidate circuits of depth d scales exponentially as $\mathcal{O}((n \cdot g)^d)$ [83]. Meanwhile, the number of circuits implementing the target algorithm is sparse. In particular, finding an optimal decomposition that minimizes circuit depth or total gate count is known to be

⁴Also called unitary synthesis.

NP-hard [84, 85]. Thus, the exponential growth with circuit depth makes, together with the solution sparsity, the synthesis of non-trivial unitaries rapidly intractable, even for moderate system sizes.

In practice quantum algorithms generally require between 10^7 and 10^{11} logical gates [5, 71], making full-algorithm synthesis far beyond reach. Due to the combinatorial hardness of the problem, a more practical approach of quantum circuit synthesis is therefore to focus on smaller subroutines that can be optimized independently and reused throughout the computation. For instance, Shor’s algorithm for n -bits RSA integer factorization relies on arithmetic subroutines involving $\mathcal{O}(n^2)$ modular additions of n -bit numbers [60, 86], where each addition typically requires $\mathcal{O}(n)$ or $\mathcal{O}(\log n)$ gates depending on ancilla availability [87, 88]. This highlights how improving the decomposition of such addition routines can have a significant impact on the overall complexity of the algorithm.

In the context of fault-tolerant quantum computing, modern compilation is not only a computationally hard translation task, it also serves as a resource-constrained optimization problem. As discussed, non-Clifford gates have a significantly higher physical cost due to distillation procedures. Thus, in addition to the intrinsic difficulty of finding a valid decomposition, a central objective of quantum compilation is to minimize the use of non-Clifford gates. More precisely, the optimization of these non-Clifford resources generally targets two distinct metrics: the T -count, referring to the total number of T gates in the circuit, and the T -depth, which corresponds to the number of sequential layers of T gates. While the T -count directly determines the volume of magic states required, the T -depth constitutes a temporal bottleneck of the computation. The choice of optimization metric is thus closely related to hardware constraints: for instance, minimizing T -depth is preferred when the architecture supports parallel magic state production.

Seeking asymptotic efficiency, early work in quantum compilation built upon the Solovay-Kitaev algorithm [89]. The original formulation of this algorithm establishes a method to approximate any unitary in $SU(d)$ using a gate set that generates a dense subset of $SU(d)$, with a sequence length scaling polylogarithmically in the inverse precision. Motivated by fault-tolerant implementations of the Shor’s algorithm, Dawson and Nielsen demonstrated in [90] its application in the approximative decomposition of the $\pi/2^k$ rotations with the Clifford+ T gate set, highlighting the practical relevance of quantum synthesis for quantum compilation. While this algorithm provides approximate decompositions, modern quantum compilation focus on exact synthesis: finding a gate sequence that exactly reproduces the target unitary, or proving that no such decomposition exists. This thesis follows this line of work.

Early modern approaches to exact quantum circuit synthesis rely on deterministic methods such as the meet-in-the-middle search techniques. Introduced in [91], this method enables the synthesis of T -depth-optimal circuits with the support of ancilla

qubits. The core idea is to match forward and backward circuit constructions by decomposing the target unitary as $U = V \cdot W$, effectively splitting the search space in two and yielding a quadratic speedup over naive brute-force methods. In addition, recent methods based on formal problem encodings and solver-based approaches have been proposed. For instance mixed-integer programming formulations solved with Gurobi or CPLEX [84, 92], or more recently boolean encodings using SAT solvers [85], enabling the synthesis of optimal circuits or the formal proof of their non-existence. These solver-based approaches benefit from decades of algorithmic progress in classical computer science. However, the underlying problems are NP-complete, and their runtimes scale exponentially with circuit size, which remains a fundamental bottleneck.

Facing exponentially hard combinatorial search spaces, heuristic methods offer an escape from formal synthesis by trading optimality for scalability. For instance, revisited heuristic variants of meet-in-the-middle have emerged [93, 94], refining the original approach through nested methods with multi-subroutine decompositions and time-space trade-offs to improve runtimes, targeting T -count and T -depth optimization, respectively. Modern prominent hybrid strategy consists of combining discrete search with greedy techniques and continuous numerical optimization. Frameworks such as QSearch [95], QFAST [96], or LEAP [97] successfully adopt this approach by combining search over circuit structures with numerical optimization. In parallel, stochastic methods like *Synthesiq* [98, 99], based on simulated annealing, achieve substantial speedups over exact search.

An alternative approach to these combinatorial methods is machine learning, which has demonstrated strong performance in large constrained search spaces. In this perspective, machine learning offers a natural framework to learn efficient synthesis strategies directly from experience, extracting insights and relevant information for decision-making. More broadly, such approaches are part of a growing trend in physics, where learning-based methods are increasingly used to address complex optimization and control problems.

At the time of writing, machine learning has become a standard tool for scientific discovery. These methods are particularly effective at uncovering structures in high-dimensional spaces where exact approaches become intractable [100]. Over the past decade, and especially with the rise of deep learning [101], they have demonstrated strong performance across a wide range of domains: in mathematics, with AlphaTensor providing a new matrix multiplication algorithms outperforming decades of human-designed constructions [102], in chemistry, with AlphaFold 3 awarded the 2024 Nobel Prize, providing highly accurate predictions for protein structures [103], and in physics, with diverse applications in accelerated materials discovery [104], and even in computer science with AlphaChip, an automated chip design system deployed to design Google’s TPU chips at scale⁵ [105]. More recently, similar ideas have started to impact quantum

⁵Chip designs produced by AlphaChip have been used in production from the TPUv5

computing [106–109], and especially in quantum compilation [83, 110–113].

More formally, machine learning refers to a broad class of methods that tackle a given task by learning from data rather than relying on explicitly programmed rules. Among these approaches, reinforcement learning is known to be suited for sequential decision-making problems under constraints. It differs from conventional machine learning in that it relies on interaction with an environment rather than labeled datasets, which is particularly advantageous for unitary synthesis problem where dataset elements are costly to generate and grow exponentially with the number of qubits. In the reinforcement learning framework, an agent interacts with an environment by taking actions, receiving rewards, and progressively learning a policy that maximizes long-term return [114]. As developed in this thesis, quantum circuit synthesis can be naturally formulated as a sequential decision-making problem over an exponentially large combinatorial space, typically modeled as a Markov Decision Process. Here the agent incrementally builds a circuit gate by gate, getting closer to the targeted unitary via a reward function, and learning through trial-and-error. Crucially, physical and architectural constraints can be directly encoded in the environment rules and reward function, enabling optimization over metrics such as T -count, circuit depth, or qubit connectivity [110].

A particularly successful class of reinforcement learning methods for such discrete search problems is based on AlphaZero style algorithms. These approaches combine Monte Carlo Tree Search with a neural policy-value network to guide the exploration. Originally introduced for board games [115], they have since been applied to complex combinatorial optimization problems [116, 117]. In this setting, learning is driven by self-play, allowing the agent to iteratively improve its decisions without external supervision. This paradigm is especially appealing for quantum circuit synthesis, where the search space is discrete, highly structured, with exponentially large combinatorial possibilities.

As a first contribution, this manuscript introduces a practical reinforcement learning framework for quantum circuit synthesis. This thesis develops an AlphaZero-inspired approach in which circuit construction is formulated as a sequential decision-making problem, guided by Monte Carlo Tree Search and a neural policy-value network. This framework enables exact unitary synthesis with unit fidelity, while remaining versatile with respect to the gate set and qubit connectivity, and flexible through reward design for depth and gate optimization.

The practical relevance of this approach is demonstrated through proof-of-concept results on standard quantum subroutines, in particular via the synthesis of the Toffoli gate with an optimal T -count decomposition. The resulting performance is competitive with state-of-the-art static compilers such as Synthetiq and meet-in-the-middle techniques,

generation, and scaled massively for TPUv6 (Trillium). Since then, AlphaChip has been extended to the design of Google’s Axion CPU and adopted by partners such as MediaTek for the Dimensity series of smartphone SoCs.

with the ability to compile random 3-qubit circuits up to depth 30 in a few seconds.

Motivated by T -count reduction, this thesis also explores the extension to dynamic circuit synthesis, incorporating measurements on ancillary qubits and conditional gates, known to positively impact the T -cost [118–120], from which the agent discovers non-trivial structures. This opens new optimization regimes, enabling further reductions in non-Clifford resources such as the T -count. More broadly, these results highlight the potential of reinforcement learning to uncover unconventional circuit constructions, paving the way toward the automated discovery of quantum protocols under realistic hardware constraints.

Part I

Large Quantum Networks

Introduction

This short introduction collects key definitions, some notations and the outline of the first part of this manuscript. It is not required to follow the thesis as the notations are introduced gradually in the chapters. However, it may help the reader by setting a global picture before diving into the calculations.

Overview of the architecture

The considered quantum repeater architecture can be divided into three stages, as depicted in figure 1.

- The *elementary link* is the building block of the repeater. It generates entanglement between two distant *hubs*, noted H, separated by intermediate distances of the order of $L_{\text{Lnk}} = 250 \text{ km}$. Localized at the edges of the elementary link, a hub combines a photon-pair source with a quantum memory.
- The *repeater chain* is a linear arrangement of elementary links. It propagates entanglement along the chain through *entanglement swapping*, a local operation on neighbouring hubs, reaching continent-scale distances of the order of $L_{\text{chn}} = 500 \text{ km}$.
- The *ion-swap* interfaces the long-distance repeater chain with the metropolitan network. Based on a dual-rail architecture, it aims to swap entanglement from the atomic-ensemble memories of the repeater chain edges onto the trapped-ion memories of the metropolitan gateway.

Modeling

In the considered architecture, the elementary link uses spontaneous parametric down-conversion (SPDC) sources for entanglement generation. Such SPDC source consists of a non-linear crystal pumped by a laser, emitting correlated photon pairs via annihilation of pump photons. The state is described by

$$|\psi_{\text{SPDC}}\rangle = \sqrt{1-p} \exp\left(\sqrt{p} a^\dagger b^\dagger\right) |00\rangle = \sqrt{1-p} \sum_n \sqrt{p}^n |nn\rangle, \quad (1)$$

where $a^\dagger$ and $b^\dagger$ are the photon-pair modes. The *pump parameter* p is directly related to the experimental pump power and sets the probability of emitting a photon pair, typically $p \approx 10^{-2}$. Operating at low pump $p \ll 1$ suppresses multi-photon emissions but reduces the distribution rate, a central trade-off in this manuscript.

The architecture relies on non-photon-number-resolving detectors (non-PNRD), which only discriminate a *no-click* from a *click* event. Following appendix A.1, the associated POVM on mode a reads

$$\begin{aligned} \text{no-click: } \quad \Pi_0^a(\eta, \text{dc}) &= (1 - \text{dc}) \cdot (1 - \eta)^{a^\dagger a}, \\ \text{click: } \quad \Pi_1^a(\eta, \text{dc}) &= \mathbb{1} - (1 - \text{dc}) \cdot (1 - \eta)^{a^\dagger a}, \end{aligned} \tag{2}$$

with η the detection efficiency and dc the dark-count probability. We sometime use the notation Π_0^a and Π_1^1 , keeping in mind it corresponds to an efficiency η and dark-count dc . For clarity, we define the loss $R = 1 - \eta$ and $\kappa = 1 - \text{dc}$. In this manuscript, all imperfections are pushed onto the detectors. For instance, a memory read out by a detector is modelled by a single efficiency $\eta = \eta_{\text{Mem}} \cdot \eta_{\text{Det}}$ combining the memory and detection losses.

The modelling lies between two formalisms: the SPDC sources with linear optics naturally fits with Gaussian formalism, while the heralding POVM is diagonal in the Fock basis. The computation is therefore delicate, as it requires translating between the two pictures: a Gaussian description of the source against a photon-number description of the measurement. We carry the derivation in the Fock formalism, since the generated entanglement lives in the qubit subspace $|ij\rangle$ with $i, j = 0, 1$. This makes the state representation easier to interpret physically, and clarifies the influence of the pump regime on the repeater performance.

Notations and numerical values

Depending on the context, we refer to the two remote parties as Alice and Bob with the indices A and B, adding Tom and Jerry with the indices T and J, when further parties are needed. Although the derivations hold in the general asymmetric case, we most often present the symmetric case where Alice and Bob share the same parameters, preventing heavy reading with unnecessary complexity. We then denote the shared value with the index AB, for instance $\eta_A = \eta_B = \eta_{AB}$. Similarly, we sometimes set the dark counts to zero in the final expressions, even though the derivation is carried out in the general case.

The architecture involves several devices: detectors, memories, displacement operators, etc. each carrying their own numerical values. In this work we use values from the literature in the perspective of a functional repeater, see table 1. Following the convention above, the efficiency of each stage of the repeater decomposes into elementary contributions. The central station combines transmission and detection, $\eta_{\text{Lnk}} = \eta_{\text{Trs}} \cdot \eta_{\text{Det}}$, the swap combines memory efficiency and detection, $\eta_{\text{Swp}} = \eta_{\text{Mem}} \cdot \eta_{\text{Det}}$, and likewise the ion interface, $\eta_{\text{Ion}} = \eta_{\text{Mem}} \cdot \eta_{\text{Det}}$, up to refinements in specific cases.

Outline

In Chapter 1, we introduce a perturbative model of an elementary link for the low-pump regime. This perturbative treatment explicitly shows the contribution of each Fock basis vector to the elementary link state and captures the main trade-off between entanglement quality and distribution rate. We then expose the efficiency wall from fibre attenuation, which leads to link segmentation for practical long-distance entanglement distribution, following the DLCZ architecture presented earlier in the main introduction.

Building on this architecture, Chapter 2 focuses on the interface between the edge hubs of the long-distance backbone with the metropolitan gateway nodes, still in the perturbative regime. We introduced a dual-rail protocol where the entanglement stored in the edge atomic-ensemble memories of the backbone is swapped onto the trapped-ion memories of the metropolitan gateways. We study the impact of photon shape mismatch between the two platforms, focusing on the improvement of a protocol stretching the atomic-ensemble photon to match the ion waveform [121].

Finally, we develop in Chapter 3 an exact model of the elementary link, accounting for all multi-photon contributions and valid across the full pump range. We derive and study the figures of merit, fidelity and heralding probability, beyond the reach of the perturbative expansion. In collaboration with the experimental group of *De Riedmatten* in Barcelona, we benchmark the predictions of our model against experimental data, providing a quantitative validation based on local photon detection statistics and visibility measurements. From this exact description, we then build an entanglement witness based on local operations and classical communication, able to certify entanglement for any memory loss.

Parameter	Symbol	Value	Reference
Pump pair probability	p_A, p_B	10^{-2}	[13, 122]
Fiber transmission	α	0.2 db/Km	[123]
Detection efficiency	η_{Det}	0.9	[124, 125]
Memory efficiency	η_{Mem}	0.5	[33, 35]
Displacement efficiency	η_{Disp}	≈ 1	[126, 127]
Dark-count probability	dc	10^{-5}	[122, 124]

Table 1: Experimental parameters used in the DLCZ repeater model.




Figure 1: Illustration of a full repeater link, made of two parallel chains, each composed of two elementary links. The three stages are highlighted: elementary link (blue), entanglement swapping (green), and ion-swap (red). Each hub combines an atomic-ensemble memory with a spontaneous parametric down-conversion source. The ion memories act as the gateways to the metropolitan networks.

1 -Quantum Repeater for long-distance entanglement distribution

1.1 .Generating entanglement between two distant hubs

1.1.1 .Idealized low-pump regime

Generating entanglement between two hubs, H_A and H_B , separated by kilometre-scale distances must assume no prior interaction. Consequently, the initial state of the joint system is factorised, $|\psi_{\text{ini}}\rangle = |\psi_A\rangle \otimes |\psi_B\rangle$, living in the Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$, where $\mathcal{H}_A$ and $\mathcal{H}_B$ are associated to Alice's and Bob's hubs H_A and H_B , respectively. The strategy to bring entanglement between $\mathcal{H}_A$ and $\mathcal{H}_B$ is then to prepare a local entangled pair inside each hub, where short-range interactions are easy, and extend this entanglement over the whole system with an operation acting on one party of each pair. Concretely, the hub H_A (resp. H_B) carries two modes a and a' (resp. b and b'). First an initial entangled state is locally generated between a and a' (resp. b and b'). Then, an operation consumes one mode from each hub through a joint measurement, for instance on the modes a' and b' , which transfers the two independent entanglements a - a' and b - b' to the remaining pair a - b . Figure 1.1 illustrates such a protocol. In a general scope, this procedure is called *entanglement swapping* since the initial entanglement of the two independent hubs is swapped across the system [8].

Set-up and Heralding

In practice, each hub is equipped with a quantum memory and a spontaneous parametric down-conversion (SPDC) photon-pair source producing a correlated *signal-idler* pair [12, 13]. The signal photon is stored in the hub memory while the idler photon is sent through a long-distance optical fibre. The two idler photons meet at a central station composed of a beam-splitter followed by a single-click heralding measurement. Such a single-click measurement, called *Bell measurement*, transfers the photonic correlations onto the memory excitations:

$$|\psi_A\rangle \otimes |\psi_B\rangle \rightarrow |\psi_{\text{entangled}}\rangle. \quad (1.1)$$

As a probabilistic measurement, the central station emits a classical feedback signal to both hubs, informing Alice and Bob whether the entanglement generation has succeeded. After a successful herald, the two memories store an entangled state ready to be consumed.

Concretely, the central station is composed of a 50-50 beam-splitter U_{BS} followed by a single-click heralding measurement Π_{Lnk} consumes the idler photon and swaps the entanglement from the two signal-idler photon pairs to the memory-memory system. See




Figure 1.1: Illustration of the entanglement swapping protocol. Local entanglement is generated in the two hubs. The joint measurement consumes the photons a' and b' , swapping the entanglement from a - a' and b - b' to the remaining modes a - b .

figure 1.2. In this preliminary study, we neglect dark counts, i.e. having a count without any input photon, and denote by η_{Lnk} the overall link efficiency, which combines the long-distance transmission efficiency η_{Trs} with the detection efficiency η_{Det} , so that $\eta_{\text{Lnk}} = \eta_{\text{Trs}} \cdot \eta_{\text{Det}}$. We denote by m_A and c_A (resp. m_B and c_B) the memory mode and the long-distance channel mode, which respectively refers to the signal and idler photon¹. More generally, the indices A and B refer to Alice's and Bob's hub respectively.

As an upper bound on the reachable entanglement quality, the initial entanglement quality is critical. To ensure clean photon-pair generation, the SPDC sources operate in a low-pump regime $p_A, p_B \ll 1$, typically $p_A, p_B = 10^{-2}$, where p_A and p_B denote the *pump parameter* of H_A and H_B . In this regime, each source generates either zero or one photon-pair, neglecting multiple pairs emissions. For clarity, we consider the symmetric situation $p_A = p_B = p$. The initial unnormalized state then reads

$$\left. \begin{aligned} |\psi_A\rangle &= |00\rangle + \sqrt{p_A} |11\rangle + o(\sqrt{p_A}) \\ |\psi_B\rangle &= |00\rangle + \sqrt{p_B} |11\rangle + o(\sqrt{p_B}) \end{aligned} \right\} \xrightarrow{|\psi_A\rangle \otimes |\psi_B\rangle} |0000\rangle + \sqrt{p} (|0011\rangle + |1100\rangle) + o(\sqrt{p}), \quad (1.2)$$

where the modes are ordered as $|m_A, c_A, c_B, m_B\rangle$, grouping the idler photons in the central pair. Using the standard 50-50 beam-splitter transformation $c_A^\dagger, c_B^\dagger \rightarrow \frac{c_A^\dagger \pm c_B^\dagger}{\sqrt{2}}$, the state becomes

$$|0000\rangle + \sqrt{p} \left(|0\rangle \frac{|01\rangle + |10\rangle}{\sqrt{2}} |1\rangle + |1\rangle \frac{|01\rangle - |10\rangle}{\sqrt{2}} |0\rangle \right) + o(\sqrt{p}). \quad (1.3)$$

Re-ordering the modes $|m_A, c_A, c_B, m_B\rangle \rightarrow |m_A, m_B\rangle |c_A, c_B\rangle$ to isolate the memory subsys-

¹In this notation, we use m referring to the memories, and c the long-distance channels.

tem, we get

$$|0000\rangle + \sqrt{p} \left(\frac{|01\rangle + |10\rangle}{\sqrt{2}} |01\rangle + \frac{|01\rangle - |10\rangle}{\sqrt{2}} |10\rangle \right) + o(\sqrt{p}). \quad (1.4)$$

This makes the purpose of the single-click measurement explicit: a single click on either c_A or c_B certifies the successful transmission of one idler photon to the central station, filtering the vacuum term, and selects one of the Bell states present in the superposition. A single click thus heralds either the Bell state $|\Psi^+\rangle$ or $|\Psi^-\rangle$, the sign $\pm$ depending on which of the two detectors clicked. The classical feedback must therefore carry both the click event and the information on which detector has clicked. Without a which-detector feedback, Alice and Bob only know they share a $|\Psi^\pm\rangle$ Bell state without any information on the sign, and the resulting state is a balanced mixture of the two Bell pair which carries no entanglement

$$|\Psi^+\rangle \langle \Psi^+| + |\Psi^-\rangle \langle \Psi^-| = |01\rangle \langle 01| + |10\rangle \langle 10|. \quad (1.5)$$

With which-detector feedback, a local correction can be applied on the wrong-parity events, and the shared state between Alice's and Bob's memories is a pure Bell pair

$$\rho_{\text{Lnk}} = |\Psi^\pm\rangle \langle \Psi^\pm|. \quad (1.6)$$

Focusing on the positive Bell pair, the fidelity, a standard metric quantifying the entanglement quality, then reads

$$\mathcal{F}_{\text{Lnk}} = \text{Tr} [|\Psi^+\rangle \langle \Psi^+| \rho_{\text{Lnk}}] = 1, \quad (1.7)$$

The heralding success probability, that is, the probability to get a single click detection on the detectors is given by

$$\mathbb{P}_{\text{Lnk}} = 2p \cdot \eta_{\text{Lnk}}, \quad (1.8)$$

where the factor 2 stands by considering both right or left single-click event, and with $\eta_{\text{Lnk}} = \eta_{\text{Trs}} \eta_{\text{Det}}$.

With a typical pump parameter $p \approx 10^{-2}$, a detection efficiency $\eta_{\text{Det}} \approx 0.9$, and a transmission probability η_{Trs} that depends on the fibre attenuation coefficient α and on the distance $L/2$ between each hub and the central station

$$\eta_{\text{Trs}} = 10^{-\alpha L/20}. \quad (1.9)$$

Standard telecom fibres exhibit an attenuation $\alpha = 0.2 \text{ dB/km}$, so that long-distance entanglement distribution at $L = 250 \text{ km}$ leads to a transmission $\eta_{\text{Trs}} \approx 10^{-5}$ and a heralding probability

$$\mathbb{P}_{\text{Lnk}} \approx 5 \cdot 10^{-5}. \quad (1.10)$$




Figure 1.2: Optical setup of an elementary link. Each hub hosts an SPDC source and an atomic-ensemble memory. The source, parameterized by a pump parameter p_A (resp. p_B), generates signal-idler photon pairs in the memory mode m_A (resp. m_B) and the long-distance channel mode c_A (resp. c_B). The central station mixes the modes c_A and c_B on a 50 – 50 beam-splitter, and a single click projects the memory modes onto an entangled state.

Entanglement quality and distribution rate as figures of merit

This preliminary derivation illustrates how the optical setup heralds pure entanglement in the idealized regime $p \ll 1$. Keeping p small aims to suppress the SPDC multi-photon contributions from higher orders in p , which may deteriorate the purity of the generated Bell pair. For instance, a second order in p brings a multi-photon terms of the form $|22\rangle$ to the system, passing the single-click heralding when one of the two idler photons is lost. This heralding contaminates the state with a p^2 term in the mixture which lower the wanted $|11\rangle$ contribution. Thus, the fidelity of the heralded state reads

$$\mathcal{F}_{\text{Lnk}} = 1 - \alpha p + o(p), \quad (1.11)$$

hence corresponding to a lower the entanglement quality.

The requirement for high fidelities imposes low photon-pair emission probabilities, resulting in a reduced success probability: beyond the fundamental limit set by the transmission losses η_{Trs} , the heralding probability scales linearly with p , thus a small p directly reduces the distribution rate. This trade-off between the quality of generated entanglement and distribution rate lies at the heart of the repeater architecture, and is a guiding thread throughout this manuscript. We therefore focus on the two following metrics, which will be tracked consistently throughout the study

- **Fidelity:** the fidelity of the heralded state to a Bell pair, denoted $\mathcal{F}_{\text{Lnk}}$, is the

entanglement quality metric. It quantifies how close the heralded state is to the target Bell state $|\Psi^+\rangle = \frac{|01\rangle + |10\rangle}{\sqrt{2}}$, and is defined as

$$\mathcal{F}_{\text{Lnk}} = \text{Tr} [\rho_{\text{Lnk}} |\Psi^+\rangle \langle \Psi^+|] . \quad (1.12)$$

- **Heralding probability:** the probability of a successful heralding event, denoted $\mathbb{P}_{\text{Lnk}}$, is the **distribution rate metric**. It quantifies the success rate in distributing an entangled pair across the link, which directly determines the entanglement distribution rate. It is defined as

$$\mathbb{P}_{\text{Lnk}} = \text{Tr} [\mathcal{E}_{\text{Lnk}} (\rho_{\text{ini}})] , \quad (1.13)$$

with ρ_{ini} the initial state from the SPDC sources, and $\mathcal{E}_{\text{Lnk}}$ the non-trace preserving quantum channel from the swapping protocol, it will be properly described later.

1.1.2 . Perturbative expansion in the pump parameter

As discussed, the trade-off between entanglement quality and distribution rate is driven by the multi-photon contributions of the SPDC sources. A quantitative understanding of this trade-off therefore requires a multi-photon study, explicitly tracking the higher-order terms in p . We thus extend the previous calculation to higher orders, quantifying how the multi-photon contributions degrade the entanglement fidelity heralded by the protocol.

In the following paragraphs, we derive a perturbative expression of the state, fidelity and the heralding probability up to the N^{th} order in p . We first establish an algebraic result: despite the complexity of the SPDC initial state, the computation reduces to independent study of the block-diagonal structure of the initial density matrix, where each diagonal block, we call *i-photon sector*, carries states with exactly i photons. From this result, we derive a perturbative expression of the heralding probability $\mathbb{P}_{\text{Lnk}}$, the heralded state ρ_{Lnk} and the fidelity $\mathcal{F}_{\text{Lnk}}$. While this expansion provides a recipe to study the multi-photon contributions sector by sector, the expressions rapidly become intractable beyond the 2nd order. Nevertheless, the expansion remains useful for numerical implementations, when experimental setup requires a higher-order model for a fine description. This is illustrated in a collaboration with the De Riedmatten group at ICFO [35], where this approach up to 2nd order were needed to reproduce the observed photon statistics.

To begin with, let us give a more formal definition of the swapping protocol. Let U_{BS} denote the beam-splitter evolution acting on the idler modes $|c_A, c_B\rangle$, and Π_{Lnk} the positive operator-value measurement (POVM) element associated with the single-click heralding event at the central station. The heralded state then reads

$$\rho_{\text{Lnk}} = \frac{1}{\mathcal{N}} \text{Tr}_{c_A, c_B} \left[\Pi_{\text{Lnk}} U_{\text{BS}} \rho_{\text{ini}} U_{\text{BS}}^\dagger \right] , \quad (1.14)$$

with $\mathcal{N}$ the normalisation constant, $\rho_{\text{ini}} = \rho_{\text{SPDC}} \otimes \rho_{\text{SPDC}}$ the initial factorized state from

the SPDC hub sources, and Tr_{c_A, c_B} the partial trace over the channel modes. Following appendix A.1, the beam-splitter transformation and the single-click POVM element acts on the channel modes c_A and c_B as

$$c_A^\dagger, c_B^\dagger \longrightarrow \frac{c_A^\dagger \pm c_B^\dagger}{\sqrt{2}}, \quad \Pi_{\text{Lnk}} = \kappa_{\text{Lnk}} \cdot R_{\text{Lnk}}^{c_A^\dagger c_A} \left(\mathbb{1} - R_{\text{Lnk}}^{c_B^\dagger c_B} \right), \quad (1.15)$$

with $R_{\text{Lnk}} = 1 - \eta_{\text{Lnk}}$ the overall loss, and $\kappa_{\text{Lnk}} = 1 - \text{dc}_{\text{Lnk}}$ with dc_{Lnk} the dark-count probability. To keep light derivation, we assume symmetric losses on c_A and c_B , concretely, the central station is placed at the middle of the link. For clarity, we define $\mathcal{E}_{\text{Lnk}}$ as the completely positive trace-non-preserving channel of the swapping protocol as

$$\mathcal{E}_{\text{Lnk}} : \rho_{\text{ini}} \longmapsto \text{Tr}_{c_A, c_B} \left[\sqrt{\Pi_{\text{Lnk}}} U_{\text{BS}} \rho_{\text{ini}} U_{\text{BS}}^\dagger \sqrt{\Pi_{\text{Lnk}}} \right]. \quad (1.16)$$

The heralding probability and fidelity reads

$$\rho_{\text{Lnk}} = \frac{\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}})}{\mathbb{P}_{\text{Lnk}}}, \quad \mathbb{P}_{\text{Lnk}} = \text{Tr}[\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}})], \quad \mathcal{F}_{\text{Lnk}} = \langle \Psi^+ | \rho_{\text{Lnk}} | \Psi^+ \rangle. \quad (1.17)$$

The derivation below includes non-zero dark counts dc_{Lnk} , while keeping the symmetric assumption $p_A = p_B = p$.

Algebraic study From a general point of view, the initial state from the two SPDC sources can be decomposed as a block matrix where each block lives in a $2i$ -photon Hilbert subspace $\tilde{\mathcal{H}}_{\text{idler}}(i) \otimes \tilde{\mathcal{H}}_{\text{signal}}(i)$, with the definition

$$\tilde{\mathcal{H}}(i) = \text{Span} \{ |kl\rangle \quad s.t. \quad k + l = i \}, \quad (1.18)$$

see appendix B.1 for a detailed study. In this context, each block i is associated with the order p^i , referring to the subspace spanned by the SPDC contribution $\tilde{\mathcal{H}}_{\text{idler}}(i) \otimes \tilde{\mathcal{H}}_{\text{signal}}(i)$, carrying $2i$ photons across the system, i photons in the idler modes c_A and c_B , and i photons in the memory modes m_A and m_B . We call these blocks the *i-photon sectors*, or the *i-photon contributions*. The factor two reflects the fact that photons are emitted pairwise by the sources, so that the order p^i produces a state of exactly $2i$ photons across the system. More precisely, $2i$ photons localised in each hub before the Bell measurement, and i photons delocalised across the two hubs system after the measurement, as the Bell measurement consume the idler photons. Concretely, the expansion of the initial state in powers of p reads

$$\rho_{\text{ini}} = \frac{1}{\mathcal{N}} \sum_{i=0}^N p^i \rho_{\text{ini}}^{(i)} + \rho_{\text{ini}}^*(p) + o(p^N), \quad \mathcal{N} = \sum_{i=0}^N p^i, \quad (1.19)$$

where N is the truncation order, $\mathcal{N}$ the normalisation constant, each $\rho_{\text{ini}}^{(i)}$ is the density matrix of the i -photon sector with $\text{Tr}[\rho_{\text{ini}}^{(i)}] = 1$, and ρ_{ini}^* collects all off-block-diagonal

cross-terms. The appendix B.1 provides formal arguments, it also shows that the cross-terms ρ_{ini}^* do not contribute to the heralded state, namely $\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^*) = 0$. Furthermore, each i -photon sector is stable under $\mathcal{E}_{\text{Lnk}}$, so that any i -photon contribution $\rho_{\text{ini}}^{(i)}$ is mapped to a heralded state $\rho_{\text{Lnk}}^{(i)}$ that also carries exactly i photons

$$\rho_{\text{ini}} = \frac{1}{\mathcal{N}} \begin{pmatrix} p^0 \rho_{\text{ini}}^{(0)} & \star & \star & \dots \\ \star & p \rho_{\text{ini}}^{(1)} & \star & \dots \\ \star & \star & p^2 \rho_{\text{ini}}^{(2)} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad (1.20)$$

$$\xrightarrow{\mathcal{E}_{\text{Lnk}}} \frac{1}{\mathcal{N}} \begin{pmatrix} p^0 \mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(0)}) & 0 & 0 & \dots \\ 0 & p \mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(1)}) & 0 & \dots \\ 0 & 0 & p^2 \mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(2)}) & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

This highlights that the computation reduces to the independent study of the diagonal blocks $\rho_{\text{ini}}^{(i)}$, which directly corresponds to the i^{th} order of the expansion. Thus, the couple $(\mathbb{P}_{\text{Lnk}}^{(i)}, \rho_{\text{Lnk}}^{(i)})$ is an intrinsic object of the i -photon sector, independent of the other sectors.

Dilution factor From the algebraic decomposition, the fidelity $\mathcal{F}_{\text{Lnk}}$ and the heralding probability $\mathbb{P}_{\text{Lnk}}$ therefore expand in powers of p , where each order isolates the contribution of one i -photon sector. This brings a natural way of studying the multi-photon contributions on the desired metrics. The normalisation requires however some care: the constant $\mathcal{N} = \sum_i p^i$ explicitly depends on p , and must itself be expanded order by order. This produces an additional expansion in p which brings a correction on top of the bare contribution of each sector. Physically, each i -photon sector sees its relative weight diluted by the presence of the higher-order sectors. We therefore refer these corrections as *dilution corrections*, since they capture how the bare contribution of an i -photon sector is diluted by the presence of higher-order sectors. This correction also guarantees that $\mathbb{P}_{\text{Lnk}}$ and $\mathcal{F}_{\text{Lnk}}$ remain bounded between 0 and 1 at every truncation order. To distinguish the two contributions, we adopt the following convention: a quantity without a tilde, $\mathbb{P}_{\text{Lnk}}^{(i)}$, $\rho_{\text{Lnk}}^{(i)}$ or $\mathcal{F}_{\text{Lnk}}^{(i)}$, denotes the bare contribution of the i -photon sector, while the same quantity with a tilde, $\tilde{\mathbb{P}}_{\text{Lnk}}^{(i)}$, $\tilde{\rho}_{\text{Lnk}}^{(i)}$ or $\tilde{\mathcal{F}}_{\text{Lnk}}^{(i)}$, includes the dilution correction.

Probability expansion We define the heralding probability of the i -photon sector as

$$\mathbb{P}_{\text{Lnk}}^{(i)} = \text{Tr} \left[\mathcal{E}_{\text{Lnk}} \left(\rho_{\text{ini}}^{(i)} \right) \right], \quad (1.21)$$

which captures the bare contribution of the i -photon sector to the heralding event. From appendix B.2, the total heralding probability reads

$$\mathbb{P}_{\text{Lnk}} = \sum_{i=0}^N p^i \tilde{\mathbb{P}}_{\text{Lnk}}^{(i)} + \mathcal{O}(p^N), \quad \tilde{\mathbb{P}}_{\text{Lnk}}^{(i)} = \mathbb{P}_{\text{Lnk}}^{(i)} - \mathbb{P}_{\text{Lnk}}^{(i-1)}, \quad (1.22)$$

where $\tilde{\mathbb{P}}_{\text{Lnk}}^{(i)}$ denotes the probability of the i -photon sector after the dilution correction. We additionally define the relative quantity

$$\hat{\mathbb{P}}_{\text{Lnk}}^{(i)} = \frac{\mathbb{P}_{\text{Lnk}}^{(i)}}{\mathbb{P}_{\text{Lnk}}^{(0)}}, \quad (1.23)$$

which measures the contribution of the i -photon sector relatively to the leading order $\mathbb{P}_{\text{Lnk}}^{(0)}$. This notation will appear naturally in the derivation of the heralded state.

State expansion We define the heralded state of the i -photon sector as

$$\rho_{\text{Lnk}}^{(i)} = \frac{\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(i)})}{\mathbb{P}_{\text{Lnk}}^{(i)}}, \quad (1.24)$$

with $\text{Tr}[\rho_{\text{Lnk}}^{(i)}] = 1$. After calculation (cf. appendix B.2), the heralded state reads

$$\rho_{\text{Lnk}} = \sum_{i=0}^N w_i \rho_{\text{Lnk}}^{(i)} + \mathcal{O}(p^N), \quad w_i = \frac{p^i \mathbb{P}_{\text{Lnk}}^{(i)}}{\sum_{j=0}^N p^j \mathbb{P}_{\text{Lnk}}^{(j)}}. \quad (1.25)$$

The weights w_i admit a Bayesian interpretation: w_i is the probability that the heralded contribution originates from the i -photon sector, given the heralding event. Finally, expanding the normalisation factor in p yields the perturbative form

$$\rho_{\text{Lnk}} = \sum_{i=0}^N p^i \tilde{\rho}_{\text{Lnk}}^{(i)} + \mathcal{O}(p^N), \quad \tilde{\rho}_{\text{Lnk}}^{(i)} = \hat{\mathbb{P}}_{\text{Lnk}}^{(i)} \rho_{\text{Lnk}}^{(i)} - \sum_{k=0}^{i-1} c_{i-k} \hat{\mathbb{P}}_{\text{Lnk}}^{(k)} \rho_{\text{Lnk}}^{(k)}, \quad (1.26)$$

where $\tilde{\rho}_{\text{Lnk}}^{(i)}$ denotes the contribution of the i -photon sector after the dilution correction. The coefficients c_k encode the Bayesian back-reaction from the expansion of w_i :

$$c_0 = 1, \quad c_k = \sum_{j=1}^k \hat{\mathbb{P}}_{\text{Lnk}}^{(j)} c_{k-j}. \quad (1.27)$$

Fidelity expansion Finally, the heralded state derived above directly yields the fidelity to the target Bell state $|\Psi^+\rangle$. We define the fidelity of the i -photon sector as

$$\mathcal{F}_{\text{Lnk}}^{(i)} = \langle \Psi^+ | \rho_{\text{Lnk}}^{(i)} | \Psi^+ \rangle, \quad (1.28)$$

which is the fidelity conditioned on the heralded contribution originating from the i -photon sector. By linearity, the fidelity propagates through the mixture, leading to

$$\mathcal{F}_{\text{Lnk}} = \sum_{i=0}^N p^i \tilde{\mathcal{F}}_{\text{Lnk}}^{(i)} + o(p^N), \quad \tilde{\mathcal{F}}_{\text{Lnk}}^{(i)} = \hat{\mathbb{P}}_{\text{Lnk}}^{(i)} \mathcal{F}_{\text{Lnk}}^{(i)} - \sum_{k=0}^{i-1} c_{i-k} \hat{\mathbb{P}}_{\text{Lnk}}^{(k)} \mathcal{F}_{\text{Lnk}}^{(k)}, \quad (1.29)$$

with the same coefficients c_k as before. The structure mirrors the heralded state: the sectorial fidelities $\mathcal{F}_{\text{Lnk}}^{(i)}$ play the role of the $\rho_{\text{Lnk}}^{(i)}$, and $\tilde{\mathcal{F}}_{\text{Lnk}}^{(i)}$ is the contribution after the dilution correction.

1.1.3 . Practical example with the two-photon truncation

The previous expansion provides a recipe to compute a perturbative expression of the heralded state, probability, and fidelity by focusing on the action of $\mathcal{E}_{\text{Lnk}}$ on each term $\rho_{\text{ini}}^{(i)}$. Here we give a practical example of this expansion at order $N = 2$. This decomposition up to the 2-photon sector was actually required by the *de Riedmatten group* (ICFO, Barcelona), where the review of [35] needed a finer modelization of the link were needed.

First, let us derive the i -photon sector $\rho_{\text{ini}}^{(i)}$ of the initial state up to 2 photons. The output state of each source is given by

$$|\psi_{\text{SPDC}}\rangle = \frac{1}{\sqrt{1-p}} \left[|00\rangle + \sqrt{p} |11\rangle + p |22\rangle + o(p) \right], \quad (1.30)$$

leading to the global initial state

$$|\psi_{\text{ini}}\rangle = \frac{1}{1-p} \left[|0000\rangle + \sqrt{p} \cdot (|1100\rangle + |0011\rangle) + p \cdot (|2200\rangle + |0022\rangle + |1111\rangle) + o(p) \right], \quad (1.31)$$

with the usual mode ordered $|m_A, c_A, c_B, m_B\rangle$. This decomposition makes explicit the contributions of the p expansion, with $\rho_{\text{ini}}^{(i)} = |\psi_{\text{ini}}^{(i)}\rangle \langle \psi_{\text{ini}}^{(i)}|$ and

$$\begin{aligned} \text{0-photon sector:} \quad & |\psi_{\text{ini}}^{(0)}\rangle = |0000\rangle, \\ \text{1-photon sector:} \quad & |\psi_{\text{ini}}^{(1)}\rangle = \frac{1}{\sqrt{2}} (|1100\rangle + |0011\rangle), \\ \text{2-photon sector:} \quad & |\psi_{\text{ini}}^{(2)}\rangle = \frac{1}{\sqrt{3}} (|2200\rangle + |0022\rangle + |1111\rangle). \end{aligned} \quad (1.32)$$

0-photon sector Obviously, the vacuum cannot generate an entangled state, but it can trigger a fake heralding through a dark-count at the central station. Formally, the heralding probability from the vacuum reads

$$\mathbb{P}_{\text{Lnk}}^{(0)} = \mathcal{E}_{\text{Lnk}}(|\underline{0}\rangle \langle \underline{0}|) = \text{dc}_{\text{Lnk}} (1 - \text{dc}_{\text{Lnk}}), \quad (1.33)$$

which heralds the state

$$\rho_{\text{Lnk}}^{(0)} = |\underline{0}\rangle \langle \underline{0}|, \quad (1.34)$$

and contributes to the fidelity as

$$\mathcal{F}_{\text{Lnk}}^{(0)} = 0. \quad (1.35)$$

1-photon contribution: To avoid heavy calculation, we only display the results here, the detailed calculation is available in appendix B.3. The 1-photon sector heralds the state

$$\rho_{\text{Lnk}}^{(1)} = \frac{1}{1 + \frac{R_{\text{Lnk}}(1 - \kappa_{\text{Lnk}})}{1 - \kappa_{\text{Lnk}}}} |\Psi_1^+\rangle \langle \Psi_1^+| + \frac{1}{1 + \frac{1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}}{R_{\text{Lnk}}(1 - \kappa_{\text{Lnk}})}} |\Psi_1^-\rangle \langle \Psi_1^-|, \quad (1.36)$$

with a heralding probability of

$$\mathbb{P}_{\text{Lnk}}^{(1)} = \kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}) + \kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}}), \quad (1.37)$$

and a fidelity

$$\mathcal{F}_{\text{Lnk}}^{(1)} = \frac{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}})}{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}) + \kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}})}. \quad (1.38)$$

2-photon sector As for the 1-photon sector, the derivation of $\rho_{\text{Lnk}}^{(2)}$ is rather heavy and is detailed in appendix B.3. The 2-photon contribution heralds the state

$$\begin{aligned} \rho_{\text{Lnk}} = & \frac{1}{1 + \frac{\kappa_{\text{Lnk}}(1 - \kappa_{\text{Lnk}})R_{\text{Lnk}}^2 + \kappa_{\text{Lnk}}(1 - \kappa_{\text{Lnk}}R_{\text{Lnk}})R_{\text{Lnk}}}{\kappa_{\text{Lnk}}(1 - \kappa_{\text{Lnk}})R_{\text{Lnk}}^2}} |\Psi_{12}^-\rangle \langle \Psi_{12}^-| \\ & + \frac{1}{1 + \frac{\kappa_{\text{Lnk}}(1 - \kappa_{\text{Lnk}}R_{\text{Lnk}}^2) + \kappa_{\text{Lnk}}(1 - \kappa_{\text{Lnk}}R_{\text{Lnk}})R_{\text{Lnk}}}{\kappa_{\text{Lnk}}(1 - \kappa_{\text{Lnk}})R_{\text{Lnk}}^2}} |\Psi_{12}^+\rangle \langle \Psi_{12}^+|, \\ & + \frac{1}{1 + \frac{\kappa_{\text{Lnk}}(1 - \kappa_{\text{Lnk}}R_{\text{Lnk}}^2) + \kappa_{\text{Lnk}}(1 - \kappa_{\text{Lnk}})R_{\text{Lnk}}^2}{\kappa_{\text{Lnk}}(1 - \kappa_{\text{Lnk}}R_{\text{Lnk}})R_{\text{Lnk}}}} |\Psi_2^-\rangle \langle \Psi_2^-| \end{aligned}$$

in the basis

$$|\Psi_{12^\pm}\rangle = \frac{|02\rangle \pm \sqrt{2}|11\rangle + |20\rangle}{2}, \quad |\Psi_2^-\rangle = \frac{|02\rangle - |20\rangle}{\sqrt{2}}, \quad (1.39)$$

with a heralding probability

$$\begin{aligned} \mathbb{P}_{\text{Lnk}}^{(2)} = & \frac{2}{3} \left[\frac{1}{4} \kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}^2) + \frac{1}{2} \kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}) + \frac{1}{4} \kappa_{\text{Lnk}} R_{\text{Lnk}}^2 (1 - \kappa_{\text{Lnk}}) \right] \\ & + \frac{1}{3} \left[\frac{1}{2} \kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}^2) + \frac{1}{2} \kappa_{\text{Lnk}} R_{\text{Lnk}}^2 (1 - \kappa_{\text{Lnk}}) \right], \end{aligned} \quad (1.40)$$

and the fidelity reads

$$\mathcal{F}_{\text{Lnk}}^{(2)} = 0. \quad (1.41)$$

Final expression Combining the contributions of the three sectors in the absence of dark counts $\text{dc}_{\text{Lnk}} = 0$ (for lighter expression), the heralded state of the elementary link

reads

$$\rho_{\text{Lnk}} = |\Psi^+\rangle\langle\Psi^+| + \left[(2\eta_{\text{Lnk}} - 3) \cdot |\Psi^+\rangle\langle\Psi^+| + (1 - \eta_{\text{Lnk}}) |\Psi_2^-\rangle\langle\Psi_2^-| + (2 - \eta_{\text{Lnk}}) |\Psi_{12}^+\rangle\langle\Psi_{12}^+| \right] \cdot p + o(p). \quad (1.42)$$

with the heralding probability

$$\mathbb{P}_{\text{Lnk}} = 2\eta_{\text{Lnk}} \cdot p + 2\eta_{\text{Lnk}}(1 - 2\eta_{\text{Lnk}}) \cdot p^2 + o(p^2), \quad (1.43)$$

and the fidelity to the target Bell pair

$$\mathcal{F}_{\text{Lnk}} = \langle\Psi^+|\rho_{\text{Lnk}}|\Psi^+\rangle = 1 - (3 - 2\eta_{\text{Lnk}}) \cdot p + o(p). \quad (1.44)$$

A short calculation for a 250 km link gives a fidelity of $\mathcal{F}_{\text{Lnk}} = (1 - 0.03) \tilde{\mathcal{F}}_{\text{Lnk}}^{(1)} = 0.97$ and a probability of $\mathbb{P}_{\text{Lnk}} = (1 - 0.01) \tilde{\mathbb{P}}_{\text{Lnk}}^{(1)} = 5.8 \cdot 10^{-5}$. Thus, considering the contribution of the 2-photon sectors decrease the fidelity by 3% and increase the heralding probability by 1%.

Discussion

- **Entanglement quality** As expected, the leading contributing order comes from the 1-photon sector. In the absence of dark counts $\text{dc}_{\text{Lnk}} = 0$, this sector heralds a pure Bell pair $\mathcal{F}_{\text{Lnk}}^{(1)} = 1$, which is consistent with the preliminary study. In the presence of dark counts, a heralding event can be triggered from a noise photon, which mixes the $|\Psi^+\rangle$ and $|\Psi^-\rangle$ Bell components (see equation (1.36)) and deteriorate the off-diagonal coherence of the heralded state. This coherence is actually a decreasing function of the dark-count: $\text{dc}_{\text{Lnk}} = 0$ heralds a pure Bell state, while $\text{dc}_{\text{Lnk}} = 1$ heralds the perfectly balanced mixture $|\Psi^+\rangle\langle\Psi^+| + |\Psi^-\rangle\langle\Psi^-|$ which carries no entanglement. In practice, the dark-count is low $\text{dc}_{\text{Lnk}} \ll 1$, so that $\mathcal{F}_{\text{Lnk}}^{(1)}$ remains close to unity. The computation also shows that 0, 2-photons sectors do not contribute to the entanglement generation. Indeed, since the fidelity is measured with the single-excitation Bell state $|\Psi^+\rangle$, any sector carrying more than one excitation has zero overlap with it

$$\mathcal{F}_{\text{Lnk}}^{(i)} = \langle\Psi^+|\rho_{\text{Lnk}}^{(i)}|\Psi^+\rangle = 0 \quad \text{for all } i \neq 1. \quad (1.45)$$

The full fidelity expansion therefore reduces to a single positive term with negative contributions from the dilution corrections of higher-order sectors

$$\mathcal{F}_{\text{Lnk}} = \mathcal{F}_{\text{Lnk}}^{(1)} - \sum_{i \neq 1}^N p^i \mathcal{F}_{\text{corr}}^{(i)}, \quad \mathcal{F}_{\text{corr}}^{(i)} = \mathcal{F}_{\text{Lnk}}^{(i)} - \tilde{\mathcal{F}}_{\text{Lnk}}^{(i)}. \quad (1.46)$$

In other words, the contribution of sectors $i \neq 1$ is purely destructive: they dilutes the single positive contribution of the 1-photon sector, lowering the heralded fidelity.

- **Heralding probability** Here, two competing effects coexist. On the one hand, a higher photon number means more photons reaching the central station, increasing the probability of triggering a detection event. On the other hand, the photons from an i -photon sector are spread across the two input modes of Π_{Lnk} , lowering the probability of finding no photon in one of the modes, as required for the single-click heralding. In the long-distance regime, the transmission efficiency η_{Trs} is small and only a few photons reach the central station and the single-click constraint is never saturated, thus multi-photon contributions boost the heralding probability in practice.

1.2 .Scaling quantum repeater

1.2.1 .The fibre-loss barrier

In the previous discussion, we exposed the *entanglement generation* protocol as a practical way for entanglement generation between distant hubs that have never interacted. Based on the *swapping protocol*, this entanglement generation extends the entanglement from two local pairs to the full link through a swapping method, resulting in a pure Bell state shared between the two distant hubs in the idealised low-pump regime [12, 13]. In a long-distance scenario, however, this architecture has to overcome a fundamental physical barrier: optical fibre attenuation. As just discussed, entanglement swapping relies on single-photon propagation through the leading 1-photon sector contribution, thus entanglement generation rate is directly limited by the channel transmissivity.

Concretely, in the first-order perturbative regime (see equation (1.8)), the heralding probability reads

$$\mathbb{P}_{\text{Lnk}} = 2p \eta_{\text{Lnk}} = 2p \eta_{\text{Det}} \eta_{\text{Trs}}, \quad (1.47)$$

with a typical pump parameter $p \approx 10^{-2}$, a detection efficiency $\eta_{\text{Det}} \approx 0.9$, and a transmission probability η_{Trs} that depends on the fibre attenuation coefficient α and on the distance $L/2$ between each hub and the central station

$$\eta_{\text{Trs}} = 10^{-\alpha L/20}. \quad (1.48)$$

Standard telecom fibres exhibit an attenuation $\alpha = 0.2 \text{ dB/km}$, so that long-distance entanglement distribution at $L = 500 \text{ km}$ leads to a transmission $\eta_{\text{Trs}} \approx 10^{-5}$ and a heralding probability

$$\mathbb{P}_{\text{Lnk}} \approx 2 \cdot 10^{-7}, \quad (1.49)$$

effectively reducing the entanglement distribution rate to impractical values.

The standard solution is to segment the long-distance channel into shorter sub-channels, called *elementary links*, each based on the optical setup described above [12]. Each elementary link generates entanglement independently between their two end memory hubs, and two neighbouring elementary links share a common intermediate hub. After a repeated trial-and-error process, once two neighbouring links have both heralded entanglement generation, a short-distance swapping operation is performed locally at the shared hub with near-unit transmission efficiency, transferring the entanglement across the two links. Recursively swapping the loaded entanglement, a series of elementary links is able to spread entanglement over the long distances, defining a *multi-links repeater*, also called *repeater chain*. Due to its exponential dependence on L , splitting the link into N shorter links brings an exponential advantage on the transmission, $\exp(-\alpha \frac{L}{20}) \rightarrow \exp(-\alpha \frac{L}{20N})$, while the independent heraldings only cost a linear slowdown.

For clarity, we adopt the following terminology. We refer to as *entanglement generation* the protocol that spreads entanglement over an elementary link, since from the multi-link repeater perspective it produces entanglement between two distant hubs that have never interacted. We say that an elementary link is *loaded* when a delocalized excitation has been heralded and stored, ready to be consumed. We refer to as *shared hubs*, or directly *swap*, the hubs shared between two neighbouring elementary links, carrying their respective end-memories. The local swap performed at shared hubs is referred to as *entanglement swapping*, in line with the standard repeater literature [12, 13].

1.2.2 . Example with 2 repeater links

To start with, let us consider a segmentation of a 500 km link into two 250 km elementary links, as depicted in figure 1.3. Concretely, each elementary link runs the entanglement generation protocol in parallel. Once a link has successfully been heralded, the corresponding Bell pair is stored in the memories of its end hubs and waits for the second link to successfully herald entanglement. When both links are loaded, a short-distance entanglement swap is performed at the central shared hub, transferring the entanglement across the full 500 km link. Halving the link length turns the per-photon transmission η_{Trs} into $\sqrt{\eta_{\text{Trs}}}$ on each elementary link, providing a quadratic gain at the cost of an additional near-unit-efficiency swap. Concretely, a single $L = 500$ km link gives $\eta_{\text{Trs}} \approx 10^{-5}$, leading to a heralding probability of the order of $\mathbb{P}_{\text{Lnk}} \approx 2 \cdot 10^{-7}$. Splitting the link into two $L/2 = 250$ km sub-links raises the per-link transmission to $\eta_{\text{Trs}} \approx 3 \cdot 10^{-3}$, and the per-link heralding probability to $\mathbb{P}_{\text{Lnk}} \approx 6 \cdot 10^{-5}$, two orders of magnitude higher.

Description

Formally, we denote with the indices L and R the quantities associated with the *left* and *right* elementary links, respectively. For instance, the memory modes of the left elementary link are denoted $|m_A^L, m_B^L\rangle$, and the right elementary link $|m_A^R, m_B^R\rangle$. The short-distance entanglement swap is structurally identical to the long-distance entanglement generation protocol: once both elementary links are loaded, the two neighbouring memory modes m_B^L and m_A^R are mixed on a beam-splitter U_{BS} , followed by a single-click heralding measurement Π_{Swp} . From appendix A.1, the entanglement swap POVM element reads

$$\Pi_{\text{Swp}} = \kappa_{\text{Swp}} \cdot R_{\text{Swp}}^{m_B^R \dagger m_B^R} \left(\mathbb{1} - R_{\text{Swp}}^{m_A^L \dagger m_A^L} \right), \quad (1.50)$$

with $\kappa_{\text{Swp}} = 1 - \text{dc}_{\text{Swp}}$ the dark-count and $R_{\text{Swp}} = 1 - \eta_{\text{Swp}}$ the loss of the swap. As the losses have been pushed onto the detector, the swap loss actually captures the memory efficiency and the detection efficiency, so that $\eta_{\text{Swp}} = \eta_{\text{Mem}} \cdot \eta_{\text{Det}}$. As $\text{dc}_{\text{Swp}} \ll 1$, we assume $\kappa_{\text{Swp}} = 1$ for a matter of clarity. This operation swaps the two independent entangled states $m_A^L - m_B^L$ and $m_A^R - m_B^R$ of the two 250 km elementary links onto the single 500 km entangled state $m_A^L - m_B^R$, propagating the entanglement across the full repeater. As for entanglement generation, we define the non-trace-preserving operator $\mathcal{E}_{\text{Swp}}$ such that

$$\mathcal{E}_{\text{Swp}} : \rho_{\text{Lnk}}^L \otimes \rho_{\text{Lnk}}^R \mapsto \text{Tr}_{m_B^L, m_A^R} \left[\sqrt{\Pi_{\text{Swp}}} U_{BS} \rho_{\text{Lnk}}^L \otimes \rho_{\text{Lnk}}^R U_{BS}^\dagger \sqrt{\Pi_{\text{Swp}}} \right], \quad (1.51)$$

thus

$$\mathbb{P}_{\text{Swp}} = \text{Tr} [\mathcal{E}_{\text{Swp}} (\rho_{\text{Lnk}}^L \otimes \rho_{\text{Lnk}}^R)], \quad \rho_{\text{Swp}} = \frac{\mathcal{E}_{\text{Swp}} (\rho_{\text{Lnk}}^L \otimes \rho_{\text{Lnk}}^R)}{\mathbb{P}_{\text{Swp}}}. \quad (1.52)$$

Using the result from the 1st order expansion, once loaded the state of the elementary links read $\rho_{\text{Lnk}}^L = \rho_{\text{Lnk}}^R = |\Psi^+\rangle \langle \Psi^+|$. The full state of the two-link repeater before the swap therefore reads

$$\rho_{\text{Lnk}}^L \otimes \rho_{\text{Lnk}}^R \longrightarrow |\psi\rangle = \frac{|1010\rangle + |1001\rangle + |0110\rangle + |0101\rangle}{2}, \quad (1.53)$$

with the mode ordering $|m_A^L, m_B^L, m_A^R, m_B^R\rangle$, for clarity the state is written as a ket. The beam-splitter U_{BS} acting on the central memory modes m_B^L and m_A^R gives

$$\begin{aligned} |\psi_{BS}\rangle &= \frac{1}{2} \left[|1\rangle \frac{|01\rangle + |10\rangle}{\sqrt{2}} |0\rangle + |1\rangle |00\rangle |1\rangle \right. \\ &\quad \left. + |0\rangle \frac{|02\rangle - |20\rangle}{\sqrt{2}} |0\rangle + |0\rangle \frac{|01\rangle - |10\rangle}{\sqrt{2}} |1\rangle \right] \\ &\rightarrow \frac{1}{2} \left[\frac{|01\rangle + |10\rangle}{\sqrt{2}} |01\rangle + \frac{|01\rangle - |10\rangle}{\sqrt{2}} |10\rangle \right. \\ &\quad \left. + \frac{|00\rangle |02\rangle}{\sqrt{2}} - \frac{|00\rangle |20\rangle}{\sqrt{2}} + |11\rangle |00\rangle \right]. \end{aligned} \quad (1.54)$$




Figure 1.3: Schematic of the entanglement swapping for a repeater chain made of two elementary links. Once the links are loaded with ρ_{Lnk}^L and ρ_{Lnk}^R , the swapping protocol, a 50 – 50 beam-splitter U_{BS} followed by a single-click heralding Π_{Swp} , propagates the entanglement onto the two edge memories of the chain.

with the mode re-ordering $|m_A^L, m_B^L, m_A^R, m_B^R\rangle \rightarrow |m_A^L, m_B^R\rangle |m_B^L, m_A^R\rangle$ to isolate the modes involved in the swap measurement. Applying the single-click measurement Π_{Swp} and tracing over the central modes yields the heralded state of the full repeater

$$\rho_{\text{Swp}} = \frac{2}{4 - \eta_{\text{Swp}}} |\Psi^+\rangle \langle \Psi^+| + \frac{2 - \eta_{\text{Swp}}}{4 - \eta_{\text{Swp}}} |00\rangle \langle 00|. \quad (1.55)$$

with the swap probability

$$\mathbb{P}_{\text{Swp}} = \frac{\eta_{\text{Swp}}(4 - \eta_{\text{Swp}})}{4}, \quad (1.56)$$

and a fidelity

$$\mathcal{F}_{\text{Swp}} = \frac{2}{4 - \eta_{\text{Swp}}}. \quad (1.57)$$

Physical Interpretation

To quantify the distribution rate advantage of the segmented architecture, we compare the expected distribution time of an entangled pair over the full repeater in the two scenarios. Let L denote the total length of the repeater. We recall the long-distance transmission efficiency from a hub to the central station of an elementary link of length L :

$$\eta_{\text{trs}}^{(L)} = 10^{-\alpha \frac{L}{2 \cdot 10}}. \quad (1.58)$$

We first focus on the two-link scenario. The two elementary links run their heralding protocol in parallel. As explained in appendix C.1, this parallelization leads to a probability

of having both links simultaneously loaded

$$\mathbb{P}_{\text{Load}} = \frac{2}{3} \mathbb{P}_{\text{Lnk}}^{(L/2)}, \quad (1.59)$$

where $\frac{2}{3}$ is the *synchronisation factor* from the parallelization, and $\mathbb{P}_{\text{Lnk}}^{(L/2)}$ is the heralding probability of an elementary link of length $L/2$. We assume the trial-and-error rate is bounded by the idler photon travel and classical feedback time. Concretely, each trial takes a time $L/2c$, where c is the speed of light in the fibre and $L/2$ is the round-trip length of the idler photon and classical feedback over an elementary link. Considering that the memory can store M modes, an elementary link can attempt $4Mc/L$ entanglement generations per second. Combining the trial rate, the heralding probability $\mathbb{P}_{\text{Lnk}}^{(L/2)}$, and the swap success probability $\mathbb{P}_{\text{Swp}}$, the expected time to distribute an entangled pair over a distance L reads

$$T_{2\text{link}} = \frac{1}{\nu \mathbb{P}_{\text{Load}} \mathbb{P}_{\text{Swp}}} = \frac{L}{4cM} \cdot \frac{3}{2p\eta_{\text{Lnk}}^{(L/2)}} \cdot \frac{4}{\eta_{\text{Swp}}(4 - \eta_{\text{Swp}})}. \quad (1.60)$$

For comparison, the expected distribution time of a single L -long elementary link reads

$$T_{1\text{link}} = \frac{1}{r \mathbb{P}_{\text{Lnk}}^{(L)}} = \frac{L}{2cM} \cdot \frac{1}{p\eta_{\text{Lnk}}^{(L)}}. \quad (1.61)$$

The relative speed-up of the segmented architecture is then

$$\mathcal{S} = \frac{T_{1\text{link}}}{T_{2\text{link}}} = \frac{\eta_{\text{Swp}}(4 - \eta_{\text{Swp}})}{3} \cdot 10^{\alpha L/20}. \quad (1.62)$$

This last result exhibits an exponential gain in L that justifies the segmented architecture.

However, this distribution time improvement comes with a reduction on the entanglement quality. On the first hand, the swap operation degrades the quality of the heralded entanglement as the fidelity $\mathcal{F}_{\text{Swp}}$ is a decreasing function of η_{Swp} . Specifically, in the limit $\eta_{\text{Swp}} = 0$ the state does not carry any coherence and the fidelity reduces to $\mathcal{F}_{\text{Swp}} = 1/2$, leaving no entanglement in the heralded state. On the other hand, the degradation goes beyond the swap inefficiency itself. In the ideal limit $\eta_{\text{Swp}} = 1$, the fidelity remains bounded by $\mathcal{F}_{\text{Swp}} = 2/3$. This bound is intrinsic to the swap protocol: a single click swap heralding can be triggered by two photons, leaving the edge memories of the repeater empty. This two-link example therefore recovers the same trade-off between distribution rate and entanglement quality already encountered at the elementary link level. At the elementary link, this trade-off was governed by the pump parameter p . At the repeater scale, it is now also governed by the repeater architecture and swap efficiency η_{Swp} , which adds an additional trade-off on top of the elementary link one.

1.2.3 . Multi-link quantum repeater

The previous calculation shows that the segmentation of a quantum repeater into sub-links leads to a strong improvement of the distribution rate at the cost of a reduction of entanglement quality. We propose here to study a N -links repeater architecture where a chain of length L_{chn} is divided into N sub-links of length L_{lnk} . We thus denote L_{lnk} the length of an elementary link, and L_{chn} the length of the repeater chain, such that $L_{\text{chn}} = N L_{\text{lnk}}$. The main idea is to parallelize the entanglement generation in elementary links and gradually merge the entanglement across the whole repeater via entanglement swapping. The entanglement propagation is done block-by-block following a divide-and-conquer architecture [13].

Formally, we consider a repeater segmented into $N = 2^n$ elementary sub-links, organized as a binary tree of depth $n = \log_2(N)$, thus $L = 2^n \cdot L_{\text{lnk}}$. At each depth $d \in \{1, \dots, n\}$, 2^{n-d} entanglement swaps are performed in parallel, each swapping two adjacent Bell pairs of length $2^{d-1} L_{\text{lnk}}$ into a single Bell pair of length $2^d L_{\text{lnk}}$. Concretely, at depth $d = 0$, the N sub-links run the entanglement generation protocol in parallel, each producing a Bell pair over a length $L_{\text{lnk}} = L/N$. At depth $d = 1$, a swap is triggered as soon as two neighbor elementary links are loaded, propagating the entanglement over a distance $2^1 L_{\text{lnk}}$. At depth $d + 1$ a swap is triggered once both depth- d adjacent Bell pairs are loaded, propagating entanglement over a distance $2^{d+1} L_{\text{lnk}}$. The recursion terminates at depth $d = n$, where a single Bell pair spans the full link of length L_{chn} .

Building upon this protocol, we can compute the state at every intermediate step. Based on the induction detailed in appendix B.4, at every depth d the intermediate state propagated through 2^d sub-links reads

$$\rho_{\text{swp}}^{(d)} = (1 - \alpha_d) |00\rangle \langle 00| + \alpha_d |\Psi^+\rangle \langle \Psi^+|, \quad (1.63)$$

where the Bell weight of the mixture α_d reads

$$\alpha_d = \frac{1}{\frac{\alpha_L + \alpha_R}{2\alpha_L\alpha_R} + 1 - \frac{\eta_{\text{swp}}}{2}}, \quad (1.64)$$

and with a heralding probability

$$\mathbb{P}_{\text{swp}}^{(d)} = \frac{\eta_{\text{swp}}}{4} [(2 - \eta_{\text{swp}}) \alpha_{d-1}^L \alpha_{d-1}^R + \alpha_{d-1}^L + \alpha_{d-1}^R]. \quad (1.65)$$

The coefficients α_{d-1}^L and α_{d-1}^R correspond to the weight of the Bell state $|\Psi^+\rangle \langle \Psi^+|$ in the previous mixture $\rho_{\text{swp}}^{(d-1)}$, respectively from the left and right sides of the swap. At the end of the protocol, the fidelity of the full repeater state reads

$$\boxed{\mathcal{F}_{\text{swp}}^{(n)} = \alpha_n}, \quad (1.66)$$

and the total loading probability

$$\mathbb{P}_{\text{Swp}}^{(n)} = \mathbb{P}_{\text{Lnk}}^{(L/N)} \left(\frac{2}{3}\right)^n \prod_{d=1}^n \mathbb{P}_{\text{Swp}}^{(d)}. \quad (1.67)$$

Let us point out that $n = \log_2(N)$ is the depth of the binary tree, hence a power of n actually means a linear dependence on the number of sub-links N . Concretely, the total probability gets an exponential speed-up on the transmission $\eta_{\text{Trs}} \rightarrow \eta_{\text{Trs}}^{1/N}$, at the cost of a linear slow-down through the synchronisation factor $\left(\frac{2}{3}\right)^n = N^{\log_2(2/3)} \approx 1/\sqrt{N}$ and swapping probability $\prod_{d=1}^n \mathbb{P}_{\text{Swp}}^{(d)}$. This is a typical speed-up a divide-and-conquer architecture, where complexity reduces to the depth of the binary tree, $\log_2 N$.

2-Interfacing the long-distance edge with the metropolitan gateway

The previous chapter introduced a method to generate entanglement between two distant hubs that have never interacted. Each hub independently prepares an entangled pair, and a swapping protocol heralds the generation of entanglement between the two hubs. This method, called *entanglement generation*, enables the creation of entanglement between two distant independent parties. However, at continent-scale distances, the swapping becomes impractical: relying on single-photon exchange, the heralding probability collapses to zero. We therefore introduced link segmentation and parallel entanglement generation, offering an exponential advantage at the cost of a vacuum component growing linearly with the number of segments. In this chapter, we assume such a segmented architecture and study the interface between a long-distance backbone repeater, based on atomic ensemble memories, and a metropolitan network, based on individual controlled system such as trapped ion, focusing on swapping the entanglement from the end-hub long-distance atomic ensemble memories to the metropolitan entry-hub ion memories. This interface swap relies on a 4-clicks protocol that filters out the vacuum and transfers only the Bell-state component of the long-distance system to the metropolitan network. This filtering ensures the distribution of high-quality entangled pair to remote metropolitan networks. The first section exposes the protocol and provides preliminary calculations to develop intuition behind it. We then bring numbers by studying a practical case: a 500 km full repeater composed of two elementary links. Finally, we take a closer look at the interface by studying the impact of a possible photon-shape mismatch between ion and atomic ensemble memories.

2.1 .The ion-swap: bridging backbone and metropolitan network

First, let's recall some context. The overall network can be decomposed into two stages. A first stage where the long-distance backbone distributes entanglement at the continental scale through a chain of elementary links made with atomic-ensemble memories. It can be viewed as a continent-scale graph with city-nodes. A second stage, the metropolitan network, which lies at the nodes of the backbone graph and aim to distribute entanglement to end users at the city scale. The long-distance backbone generates entanglement between its edge hubs, and a dedicated swapping protocol, based on a 4-clicks heralding, transfers it to the metropolitan gateways. This swap, that we will call *ion-swap* or *4-clicks protocol*, is the central object of this chapter.

As established previously, segmenting the long-distance link mitigates exponential photon loss at the cost a vacuum component growing linearly with the number of




Figure 2.1: Setup of the dual-rail full repeater with ion swapping. Both chains are loaded with a state $\rho_{\text{Chn}}^{\text{U/D}}$ following the derivation of section 1.2. Once loaded, the entanglement is swapped from the two chains onto the ion memories via a 4-click heralding protocol. Note we dropped the tilde on the mode notation for clarity.

elementary links. We refer to such a multi-segment structure as a *repeater chain*, or simply *chain*. In this context, we denote with the index Chn the quantities related to the repeater chain: $\mathcal{F}_{\text{Chn}}$, $\mathbb{P}_{\text{Chn}}$, etc. Bridging the gap with the notations from section 1.2, for a chain made of n elementary links we have: $\rho_{\text{Chn}} = \rho_{\text{Swp}}^{(n)}$, $\rho_{\text{Chn}} = \rho_{\text{Swp}}^{(n)}$ and $\mathcal{F}_{\text{Chn}} = \mathcal{F}_{\text{Swp}}^{(n)}$. The 4-clicks protocol serves two purposes: transferring the entanglement from the backbone to the metropolitan gateways, and filtering out the vacuum component to select the Bell-state contribution of the mixture. For simplicity, in the whole section we neglect dark-count.

The architecture relies on two parallel repeater chains, each providing a Bell-state/vacuum mixture at the chain edges. At each metropolitan-network gateway, a trapped-ion memory emits a photon whose polarization is entangled with its internal energy state. This photon is routed with the backbone photon to a local Bell-state measurement, where a 4-clicks heralding swaps the entanglement from the backbone to the metropolitan ions. This event requires one click per rail and one click per ion, respectively provided by the backbone photon Bell pairs and the ion-photon pairs, which filter out the vacuum component. Thus, the 4-clicks heralding consumes two delocalized long-distance Bell pairs and produces a pure Bell state in two remote metropolitan gateway hubs.

2.1.1 .The network protocol

We now get a closer look at the network architecture. As depicted in figure 2.1, the architecture consists of two long-distance repeater chains [13], each generating entanglement between two edge solid-state quantum memories, four balanced beam splitters (BS), eight non-photon-number-resolving photodetectors Π_{Ion} , two polarising beam splitters (PBS), and two ion nodes (denoted with the index A and B) [45], assumed to be located in two distinct metropolitan networks far apart from one another. In the following analysis, we assume long-distance repeater chains operate in the ideal regime

restricted to the single-photon sector, as discussed in the previous chapter. We thus start from loaded long-distance repeater chains, each producing a mixed state of the form

$$\begin{aligned}\rho_{\text{Chn}}^U &= (1 - \alpha^U) |00\rangle \langle 00| + \alpha^U |\Psi^+\rangle \langle \Psi^+| \\ \rho_{\text{Chn}}^D &= (1 - \alpha^D) |00\rangle \langle 00| + \alpha^D |\Psi^+\rangle \langle \Psi^+|,\end{aligned}\quad (2.1)$$

where the U and D indices refer to the upper and lower repeater chain, respectively. We introduce the optical modes $(\tilde{b}^\dagger, \tilde{c}^\dagger, \tilde{d}^\dagger, \tilde{e}^\dagger, \tilde{f}^\dagger, \tilde{g}^\dagger, \tilde{h}^\dagger, \tilde{i}^\dagger)$, which refer to the photon modes at the BS input ports, the un-tilde counterparts refer to the optical modes at the BS output. Driving from the state $|g\rangle$ of any of the ions leads to the emission of a photon of polarization V or H, with the ion decaying either to $|g_V\rangle$ or to $|g_H\rangle$. For convenience, we introduce the mode-like operators (a_H, a_V) and (j_H, j_V) such that $a_{H/V}^\dagger |g\rangle_A = |g_{H/V}\rangle_A$ and $j_{H/V}^\dagger |g\rangle_B = |g_{H/V}\rangle_B$, where the indices A and B refer to Alice's and Bob's ions, respectively.

Each ion is prepared in a superposition of $|g_V\rangle$ and $|g_H\rangle$, emitting two photons of orthogonal polarizations $_H$ and $_V$, distributed over the modes $\tilde{b}$ and $\tilde{c}$ through the PBS,

$$|g\rangle_A \rightarrow \frac{|g_H\rangle_A \otimes |1\rangle_{\tilde{b}} + |g_V\rangle_A \otimes |1\rangle_{\tilde{c}}}{\sqrt{2}}, \quad (2.2)$$

and similarly for Bob's ion. Additionally, the two repeater chains are loaded in parallel with a Bell state and vacuum mixture, with a triggered emission from the memories. For simplicity we only track the Bell state in the protocol since the vacuum cannot trigger a heralding event¹, but one keeps in mind an extra dilution factor $\alpha^U \cdot \alpha^D$ on the calculation, taking into account the weight of the Bell state in the initial mixture. Precisely, the protocol starts from the states

$$\begin{aligned}|\psi_{\text{Ion}}^A\rangle &= \frac{|g, 10\rangle + |e, 01\rangle}{\sqrt{2}} \rightarrow \frac{b^\dagger + a^\dagger \cdot c^\dagger}{\sqrt{2}} \\ |\psi_{\text{Ion}}^B\rangle &= \frac{|g, 10\rangle + |e, 01\rangle}{\sqrt{2}} \rightarrow \frac{h^\dagger + j^\dagger \cdot i^\dagger}{\sqrt{2}} \\ |\psi_{\text{Chn}}^U\rangle &= \frac{|10\rangle + |01\rangle}{\sqrt{2}} \rightarrow \frac{d^\dagger + e^\dagger}{\sqrt{2}} \\ |\psi_{\text{Chn}}^D\rangle &= \frac{|10\rangle + |01\rangle}{\sqrt{2}} \rightarrow \frac{f^\dagger + g^\dagger}{\sqrt{2}},\end{aligned}\quad (2.3)$$

leading to

$$|\psi_{\text{ini}}\rangle = \left[\frac{a_H^\dagger \tilde{b}^\dagger + a_V^\dagger \tilde{c}^\dagger}{\sqrt{2}} \right] \left[\frac{\tilde{d}^\dagger + \tilde{e}^\dagger}{\sqrt{2}} \right] \left[\frac{\tilde{f}^\dagger + \tilde{g}^\dagger}{\sqrt{2}} \right] \left[\frac{\tilde{h}^\dagger j_H^\dagger + \tilde{i}^\dagger j_V^\dagger}{\sqrt{2}} \right] |\bar{0}\rangle. \quad (2.4)$$

The network is loaded with four excitations in total: one delocalized in $(\tilde{b}, \tilde{c})$ from Alice's ion, one delocalized in $(\tilde{d}, \tilde{e})$ from the upper repeater chain, one delocalized in $(\tilde{f}, \tilde{g})$ from the lower repeater chain, and one delocalized in $(\tilde{h}, \tilde{i})$ from Bob's ion.

The four beam splitters generate interference between the pairs of modes $(\tilde{b}, \tilde{d})$, $(\tilde{c}, \tilde{f})$, $(\tilde{h}, \tilde{e})$, and $(\tilde{i}, \tilde{g})$. For now, we assume perfect interference and introduce the output modes (b, d) , (c, f) , (h, e) , and (i, g) , such that the unitary action of the beam splitters

¹As we neglect dark-count in this section.

corresponds to the mapping

$$\tilde{b}, \tilde{d} \rightarrow b, d = \frac{\tilde{b} \pm \tilde{d}}{\sqrt{2}}, \quad (2.5)$$

and similarly for the other modes. The BS unitary action transforms $|\psi_{\text{ini}}\rangle$ into

$$|\psi\rangle = \frac{1}{4} \left[a_{\text{H}} \cdot \frac{b+d}{\sqrt{2}} + a_{\text{V}} \cdot \frac{c+f}{\sqrt{2}} \right]^{\dagger} \left[\frac{b-d}{\sqrt{2}} + \frac{h-e}{\sqrt{2}} \right]^{\dagger} \left[\frac{c-f}{\sqrt{2}} + \frac{i-g}{\sqrt{2}} \right]^{\dagger} \left[j_{\text{H}} \cdot \frac{h+e}{\sqrt{2}} + j_{\text{V}} \cdot \frac{i+g}{\sqrt{2}} \right]^{\dagger} |\bar{0}\rangle. \quad (2.6)$$

The BS transformation is followed by a heralding event consisting of four single-click measurements, one after each beam splitter. Since the network contains exactly four photons delocalised across the optical modes, two scenarios are possible. Clicks on b , c , h , and i leave the ions in $|g_{\text{V}}\rangle_{\text{A}} \otimes |g_{\text{H}}\rangle_{\text{B}}$ or in $|g_{\text{H}}\rangle_{\text{A}} \otimes |g_{\text{V}}\rangle_{\text{B}}$, depending on whether the clicks originate from the memories' or the ions' photons, leaving both possible states in superposition. One ends up with the ion-encoded Bell state

$$|\psi_{\text{ion}}\rangle = \frac{|g_{\text{V}}\rangle_{\text{A}} \otimes |g_{\text{H}}\rangle_{\text{B}} + |g_{\text{H}}\rangle_{\text{A}} \otimes |g_{\text{V}}\rangle_{\text{B}}}{\sqrt{2}}. \quad (2.7)$$

The case where two photons (here $\tilde{b}$ and $\tilde{d}$) bunch and produce the same single-click event is discarded by restricting the heralding to four clicks rather than fewer. Based on appendix A.1, the POVM element of such a heralding reads

$$\Pi_{\text{ion}} = \prod_{(\alpha, \beta) \in \mathcal{S}} \Pi_1^{\alpha}(\eta_{\text{ion}}) \Pi_0^{\beta}(\eta_{\text{ion}}), \quad (2.8)$$

with η_{ion} the efficiency heralding, and $\mathcal{S} = \{(b, d), (c, f), (h, e), (g, i)\}$ the set of click/no-click mode pair. Note that, as for the entanglement generation, multiple heralding events are possible: one click after each BS yields $2^4 = 16$ possibilities, all leading to the same state up to a local phase from the interference, which can be handled by local corrections. Thus, we just focus on one specific event, keeping in mind the factor 2^4 from this degeneracy.

From this preliminary study, one can compute the resulting fidelity, heralding probability, and distribution time. The distributed state is a Bell pair since the 4-clicks protocol can only herald a delocalized excitation between the two ions:

$$\mathcal{F}_{\text{ion}} = 1. \quad (2.9)$$

Considering a symmetric chain $\alpha^{\text{U}} = \alpha^{\text{D}} = \alpha$, the probability of measuring four single clicks given loaded chains reads

$$\mathbb{P}_{\text{ion}} = \alpha^2 \frac{\eta_{\text{ion}}^4}{8}, \quad (2.10)$$

and the distribution time reads

$$T = \frac{2}{3} \frac{1}{\mathbb{P}_{\text{ion}} \mathbb{P}_{\text{Chn}}} \frac{1}{\nu_{\text{Lnk}}}. \quad (2.11)$$

with $\mathbb{P}_{\text{chn}}$ the probability of having a chain loaded, and ν_{Lnk} the elementary link distribution rate. This unit fidelity is a direct consequence of the idealised low-pump scenario: we assumed a vacuum/Bell-state mixture in the repeater chains, that itself comes from the low-pump regime where only the 1-photon sector is considered. Higher-order contributions are expected to decrease the fidelity. This is precisely what we observe in the next section, where we study a practical implementation: a full dual-rail 500 km repeater built from two-link chains, including the contribution of the 2-photon sector.

2.1.2 .Practical example with a 500 km link

In this section, we study a practical implementation of a full repeater link over 500 km, based on the dual-rail architecture described above. For an accurate evaluation, we include contributions up to the 2-photon sector, making the evaluated metrics p -dependent. This p -dependence couples the fidelity $\mathcal{F}_{\text{ion}}$ and the entanglement distribution time T : the pump parameter p_{Target} is fixed by a target fidelity $\mathcal{F}_{\text{Target}}$, yielding an ion-ion distribution time $T(p_{\text{Target}})$.

As depicted in figure 2.2, the architecture is divided into three stages: entanglement creation in the elementary links (1), entanglement propagation along the repeater chain (2), and edges-to-ion swap (3). Each of these steps has been already studied, and we now combine the partial results from the previous sections. For simplicity, we neglect dark counts and assume symmetry between the upper and lower chains, as well as between Alice's and Bob's hubs.

Entanglement creation (1) Starting with entanglement creation, based on the previous perturbative study and equation (1.42), the contribution of the 1,2-photons sectors yield the state

$$\rho_{\text{Lnk}} = |\Psi^+\rangle\langle\Psi^+| + \left[(2\eta_{\text{Lnk}} - 3) \cdot |\Psi^+\rangle\langle\Psi^+| + (1 - \eta_{\text{Lnk}}) |\Psi_2^-\rangle\langle\Psi_2^-| + (2 - \eta_{\text{Lnk}}) |\Psi_{12}^+\rangle\langle\Psi_{12}^+| \right] \cdot p + o(p). \quad (2.12)$$

with a fidelity

$$\mathcal{F}_{\text{Lnk}} = \langle\Psi^+| \rho_{\text{Lnk}} |\Psi^+\rangle = 1 - (3 - 2\eta_{\text{Lnk}}) \cdot p + o(p). \quad (2.13)$$

and a heralding probability

$$\mathbb{P}_{\text{Lnk}} = 2\eta_{\text{Lnk}} \cdot p + 2\eta_{\text{Lnk}}(1 - 2\eta_{\text{Lnk}}) \cdot p^2 + o(p^2). \quad (2.14)$$

We quickly recall η_{Lnk} is the efficiency of the heralding POVM element of the central station elementary link, accounting for the long-distance transmission η_{Trs} and detection efficiency η_{Det} , thus $\eta_{\text{Lnk}} = \eta_{\text{Det}} \cdot \eta_{\text{Trs}}$.

Entanglement propagation (2) Once two elementary links are loaded, the entanglement can be propagated through the whole chain. Focusing on a single chain and considering a second-order expansion, we expect an expanded density matrix of the form

$$\rho_{\text{Chn}} = \tilde{\rho}_{\text{Chn}}^{(1)} + \tilde{\rho}_{\text{Chn}}^{(2)} \cdot p + o(p), \quad (2.15)$$

where $\rho_{\text{Chn}}^{(1)}$ comes from the one-photon contribution and $\rho_{\text{Chn}}^{(2)}$ from the two-photon contribution. Using a computer algebra program, the computation yields

$$\begin{aligned} \tilde{\rho}_{\text{Chn}}^{(1)} &= \frac{2 - \eta_{\text{Chn}}}{4 - \eta_{\text{Chn}}} \cdot |00\rangle \langle 00| + \frac{2}{4 - \eta_{\text{Chn}}} \cdot |\Psi^+\rangle \langle \Psi^+| \\ \tilde{\rho}_{\text{Chn}}^{(2)} &= \frac{(-2 + \eta_{\text{Chn}})(48 - 30\eta_{\text{Lnk}} + \eta_{\text{Chn}}(-20 + 13\eta_{\text{Lnk}}))}{2(4 - \eta_{\text{Chn}})^2} \cdot |00\rangle \langle 00| \\ &+ \frac{8\eta_{\text{Lnk}} + \eta_{\text{Chn}}(-80 + \eta_{\text{Chn}}(28 - 17\eta_{\text{Lnk}}) + 44\eta_{\text{Lnk}})}{4(-4 + \eta_{\text{Chn}})^2} \cdot |\Psi^+\rangle \langle \Psi^+| \\ &+ \frac{(4 - 3\eta_{\text{Chn}})(4 - 3\eta_{\text{Lnk}})}{4(4 - \eta_{\text{Chn}})} \cdot |\Psi^-\rangle \langle \Psi^-| \\ &+ \frac{4 - 3\eta_{\text{Lnk}}}{8 - 2\eta_{\text{Chn}}} \cdot |\Psi_2^-\rangle \langle \Psi_2^-| \\ &+ \frac{12 - 7\eta_{\text{Lnk}} - \sqrt{80 - 72\eta_{\text{Lnk}} + 17\eta_{\text{Lnk}}^2}}{4(4 - \eta_{\text{Chn}})} \cdot |\Psi_{\gamma+}^-\rangle \langle \Psi_{\gamma+}^-| \\ &+ \frac{12 - 7\eta_{\text{Lnk}} + \sqrt{80 - 72\eta_{\text{Lnk}} + 17\eta_{\text{Lnk}}^2}}{4(4 - \eta_{\text{Chn}})} \cdot |\Psi_{\gamma-}^-\rangle \langle \Psi_{\gamma-}^-|, \end{aligned} \quad (2.16)$$

with η_{Chn} the entanglement swapping efficiency, as defined in section 1.2 and where we defined the orthonormal basis

$$|\Psi_2^-\rangle = \frac{|02\rangle - |20\rangle}{\sqrt{2}} \quad |\Psi_{\gamma\pm}^\pm\rangle = \gamma_\pm \cdot |02\rangle \pm \sqrt{1 - 2\gamma_\pm^2} \cdot |11\rangle + \gamma_\pm \cdot |20\rangle,$$

under the condition

$$\left(\frac{\gamma_+ + \gamma_-}{2}\right)^2 + \left(\frac{\gamma_+ - \gamma_-}{2}\right)^2 = \frac{1}{4}. \quad (2.17)$$

In our situation, a proper decomposition of the state gives

$$\gamma_\pm = \frac{4 - \eta_{\text{Lnk}} \pm \sqrt{80 + \eta_{\text{Lnk}}(-72 + 17\eta_{\text{Lnk}})}}{\sqrt{2}\sqrt{16(2 - \eta_{\text{Lnk}})^2 + (4 - \eta_{\text{Lnk}} + \sqrt{80 + \eta_{\text{Lnk}}(17\eta_{\text{Lnk}} - 72)})^2}}. \quad (2.18)$$

When both elementary links are loaded, the probability of achieving the swap is

$$\mathbb{P}_{\text{Chn}} = \eta_{\text{Chn}} \frac{4 - \eta_{\text{Chn}}}{4} + \eta_{\text{Chn}} \frac{2\eta_{\text{Lnk}} - \eta_{\text{Chn}}(8 - 5\eta_{\text{Lnk}})}{8} \cdot p + o(p). \quad (2.19)$$

and the fidelity of ρ_{chn} with a Bell pair reads

$$\mathcal{F}_{\text{chn}} = \frac{2}{4 - \eta_{\text{chn}}} + \frac{8\eta_{\text{Lnk}} + \eta_{\text{chn}}(80 - \eta_{\text{chn}}(28 - 17\eta_{\text{Lnk}}) - 44\eta_{\text{Lnk}})}{4(4 - \eta_{\text{chn}})^2} \cdot p + o(p). \quad (2.20)$$

Ion swap (3) Once the two chains are loaded, the 4-clicks heralding swaps the excitation from the chains to the ions, with probability $\mathbb{P}_{\text{lon}}$, resulting in a state with the fidelity $\mathcal{F}_{\text{lon}}$. Due to the complexity of the computation, we rely on the strategy developed in section C.2. Briefly, this strategy decomposes the state ρ_{chn} over the orthogonal basis $\rho_i = |\Psi_i\rangle\langle\Psi_i|$, rather than in terms of n -photon sectors. Concretely, for each chain we write

$$\rho_{\text{chn}} = \sum_i a_i \cdot \rho_i, \quad (2.21)$$

with $\{a_i\}$ and $\{\rho_i\}$ defined from this decomposition. The coefficients a_i carry the p -expansion up to second order, such that $a_i = \alpha_i + \beta_i \cdot p + o(p)$. The index i is defined in table 2.1, for instance $i = 2$ yields

$$\begin{aligned} \rho_2 &= |\Psi^+\rangle\langle\Psi^+| \\ a_2 &= \frac{2}{4 - \eta_{\text{chn}}} + p \cdot \frac{8\eta_{\text{Lnk}} + \eta_{\text{chn}}(-80 + \eta_{\text{chn}}(28 - 17\eta_{\text{Lnk}}) + 44\eta_{\text{Lnk}})}{4(-4 + \eta_{\text{chn}})^2}. \end{aligned} \quad (2.22)$$

We introduce $\rho_i \otimes \rho_j$ for the components of $\rho_{\text{chn}} \otimes \rho_{\text{chn}}$. Based on this decomposition, we build table 2.1 and identify the coefficients $\{\alpha_i\}$ and $\{\beta_i\}$. Table 2.2 then shows the heralding probability $\mathbb{P}_{i,j}$ and the fidelity $\mathcal{F}_{i,j}$ for every term $\rho_i \otimes \rho_j$. Since the expansion cut-off is at order p^2 , $\mathcal{F}_{\text{lon}}$ and $\mathbb{P}_{\text{lon}}$ do not involve $\beta_i \cdot \beta_j$ terms, and we therefore restrict the table to $i = 1, 2$. Furthermore, to avoid duplication, the symmetric pairs $i \longleftrightarrow j$ are omitted. Finally, combining the two tables with the following expressions yields the desired ion-swap heralding probability and associated fidelity:

$$\boxed{\mathbb{P}_{\text{lon}} = \mathbb{P}_{\text{lon}}^{(1)} + \mathbb{P}_{\text{lon}}^{(2)} \cdot p + o(p)} \quad (2.23)$$

$$\boxed{\mathcal{F}_{\text{lon}} = \frac{\mathcal{F}_{\text{lon}}^{(1)}}{\mathbb{P}_{\text{lon}}^{(1)}} + \left(\frac{\mathcal{F}_{\text{lon}}^{(2)}}{\mathbb{P}_{\text{lon}}^{(1)}} - \frac{\mathcal{F}_{\text{lon}}^{(1)}}{\mathbb{P}_{\text{lon}}^{(1)}} \cdot \frac{\mathbb{P}_{\text{lon}}^{(2)}}{\mathbb{P}_{\text{lon}}^{(1)}} \right) \cdot p + o(p)} \quad (2.24)$$

with

$$\begin{aligned} \mathcal{F}_{\text{lon}}^{(1)} &= \sum_{i,j} \alpha_i \alpha_j \cdot \mathbb{P}_{i,j} \mathcal{F}_{i,j} & \mathcal{F}_{\text{lon}}^{(2)} &= \sum_{i,j} (\alpha_i \beta_j + \alpha_j \beta_i) \cdot \mathbb{P}_{i,j} \mathcal{F}_{i,j} \\ \mathbb{P}_{\text{lon}}^{(1)} &= \sum_{i,j} \alpha_i \alpha_j \cdot \mathbb{P}_{i,j} & \mathbb{P}_{\text{lon}}^{(2)} &= \sum_{i,j} (\alpha_i \beta_j + \alpha_j \beta_i) \cdot \mathbb{P}_{i,j}. \end{aligned} \quad (2.25)$$

Due to the length, we do not provide the full analytical result and restrict the study to a numerical results. Note that we recover the general results of section 1.1.2 for the fidelity, with an explicit expression of the 2-photon contribution and its associated dilution

correction:

$$\tilde{\mathcal{F}}_{\text{lon}}^{(1)} = \frac{\mathcal{F}_{\text{lon}}^{(1)}}{\mathbb{P}_{\text{lon}}^{(1)}}, \quad \tilde{\mathcal{F}}_{\text{lon}}^{(2)} = \frac{\mathcal{F}_{\text{lon}}^{(2)}}{\mathbb{P}_{\text{lon}}^{(1)}} - \frac{\mathcal{F}_{\text{lon}}^{(1)}}{\mathbb{P}_{\text{lon}}^{(1)}} \cdot \frac{\mathbb{P}_{\text{lon}}^{(2)}}{\mathbb{P}_{\text{lon}}^{(1)}}. \quad (2.26)$$

The probability, on the contrary, does not carry any dilution factor, this is a direct consequence of the assumption that the initial state from the repeater chains is already normalized

$$\tilde{\mathbb{P}}_{\text{lon}}^{(1)} = \mathbb{P}_{\text{lon}}^{(1)}, \quad \tilde{\mathbb{P}}_{\text{lon}}^{(2)} = \mathbb{P}_{\text{lon}}^{(2)}. \quad (2.27)$$

Also note that the zeroth-order coefficient of the fidelity is unity by construction: $\tilde{\mathcal{F}}_{\text{lon}}^{(1)} = \mathcal{F}_{\text{lon}}^{(1)} / \mathbb{P}_{\text{lon}}^{(1)} = 1$.

Distribution time From section C.1, the entanglement distribution time T reads

$$T = \left(\frac{3}{2}\right)^2 \frac{1}{M} \cdot \frac{2L_{\text{chn}}}{c} \cdot \frac{1}{\mathbb{P}_{\text{lon}} \cdot \mathbb{P}_{\text{chn}} \cdot \mathbb{P}_{\text{lnk}}},$$

where M is the number of modes, $\nu_{\text{lnk}} = c/L_{\text{lnk}} = 2c/L_{\text{chn}}$ the elementary link distribution rate, $\mathbb{P}_{\text{lnk}}$ is the heralding probability of a single elementary link, $\mathbb{P}_{\text{chn}} = \mathbb{P}_{\text{swp}}$ is the probability of loading one chain (i.e achieving the entanglement swapping while the two elementary links are loaded), and $\mathbb{P}_{\text{lon}}$ is the heralding probability of the 4-clicks protocol while having the two repeater-chain loaded. The $(3/2)^2$ factor is the synchronization factor from both parallelization of the elementary link loading and link-chain loading, formally

$$\mathbb{P}_{\text{TOT}} = \left(\frac{2}{3}\mathbb{P}_{\text{chn}}\right) \mathbb{P}_{\text{lon}} = \left(\frac{2}{3}\right) \left(\frac{2}{3}\mathbb{P}_{\text{lnk}}\right) \mathbb{P}_{\text{chn}} \mathbb{P}_{\text{lon}}. \quad (2.28)$$

From Equation 2.14, Equation 2.19 and Equation 2.23, the probabilities decompose as

$$\begin{aligned} \mathbb{P}_{\text{lnk}} &= \tilde{\mathbb{P}}_{\text{lnk}}^{(1)} \cdot p + \tilde{\mathbb{P}}_{\text{lnk}}^{(2)} \cdot p^2 + \mathbf{o}(p^2), \\ \mathbb{P}_{\text{chn}} &= \tilde{\mathbb{P}}_{\text{chn}}^{(1)} + \tilde{\mathbb{P}}_{\text{chn}}^{(2)} \cdot p + \mathbf{o}(p), \\ \mathbb{P}_{\text{lon}} &= \tilde{\mathbb{P}}_{\text{lon}}^{(1)} + \tilde{\mathbb{P}}_{\text{lon}}^{(2)} \cdot p + \mathbf{o}(p). \end{aligned} \quad (2.29)$$

We thus obtain the first-order expansion

$$\begin{aligned} \frac{1}{\mathbb{P}_{\text{lnk}} \cdot \mathbb{P}_{\text{chn}} \cdot \mathbb{P}_{\text{lon}}} &= \frac{1}{\tilde{\mathbb{P}}_{\text{lnk}}^{(1)} \cdot \tilde{\mathbb{P}}_{\text{chn}}^{(1)} \cdot \tilde{\mathbb{P}}_{\text{lon}}^{(1)}} \cdot p^{-1} \\ &\quad - \frac{\tilde{\mathbb{P}}_{\text{lnk}}^{(2)} \cdot \tilde{\mathbb{P}}_{\text{chn}}^{(1)} \cdot \tilde{\mathbb{P}}_{\text{lon}}^{(1)} + \tilde{\mathbb{P}}_{\text{lnk}}^{(1)} \cdot \tilde{\mathbb{P}}_{\text{chn}}^{(2)} \cdot \tilde{\mathbb{P}}_{\text{lon}}^{(1)} + \tilde{\mathbb{P}}_{\text{lnk}}^{(1)} \cdot \tilde{\mathbb{P}}_{\text{chn}}^{(1)} \cdot \tilde{\mathbb{P}}_{\text{lon}}^{(2)}}{\left(\tilde{\mathbb{P}}_{\text{lnk}}^{(1)} \cdot \tilde{\mathbb{P}}_{\text{chn}}^{(1)} \cdot \tilde{\mathbb{P}}_{\text{lon}}^{(1)}\right)^2} + \mathbf{o}(1), \end{aligned} \quad (2.30)$$

leading to the distribution time

$$\boxed{T = T^{(1)} \cdot \frac{1}{p} + T^{(2)} + \mathbf{o}(1)}, \quad (2.31)$$

with

$$\begin{aligned} T^{(1)} &= \left(\frac{3}{2}\right)^2 \frac{1}{M} \cdot \frac{L}{2c} \cdot \frac{1}{\mathbb{P}_{\text{Lnk}}^{(1)} \cdot \mathbb{P}_{\text{Chn}}^{(1)} \cdot \mathbb{P}_{\text{lon}}^{(1)}}, \\ T^{(2)} &= - \left(\frac{3}{2}\right)^2 \frac{1}{M} \cdot \frac{L}{2c} \cdot \frac{\mathbb{P}_{\text{Lnk}}^{(2)} \cdot \mathbb{P}_{\text{Chn}}^{(1)} \cdot \mathbb{P}_{\text{lon}}^{(1)} + \mathbb{P}_{\text{Lnk}}^{(1)} \cdot \mathbb{P}_{\text{Chn}}^{(2)} \cdot \mathbb{P}_{\text{lon}}^{(1)} + \mathbb{P}_{\text{Lnk}}^{(1)} \cdot \mathbb{P}_{\text{Chn}}^{(1)} \cdot \mathbb{P}_{\text{lon}}^{(2)}}{\left(\mathbb{P}_{\text{Lnk}}^{(1)} \cdot \mathbb{P}_{\text{Chn}}^{(1)} \cdot \mathbb{P}_{\text{lon}}^{(1)}\right)^2}. \end{aligned} \quad (2.32)$$




Figure 2.2: Practical example of a dual-rail, two-link, full repeater chain of length $L_{\text{Chn}} = 500$ km. The three stages are highlighted in blue (entanglement creation), green (entanglement swapping), and red (ion swapping).

2.1.3 . Numerical study

In this numerical study we expect the traditional trade-off between entanglement quality and distribution time. To make this trade-off explicit, we plot the distribution time T as a function of the ion-ion fidelity $\mathcal{F}_{\text{lon}}$, giving a numerical estimate of the distribution time for a target fidelity $\mathcal{F}_{\text{Target}}$. The linear dependence in p makes this straightforward. Given a target fidelity $\mathcal{F}_{\text{Target}}$, the p -expansion from equation (2.24) fixes the pump parameter

$$\mathcal{F}_{\text{Target}} = \tilde{\mathcal{F}}_{\text{Target}}^{(1)} + \tilde{\mathcal{F}}_{\text{Target}}^{(2)} \cdot p_{\text{Target}} \implies p_{\text{Target}} = \frac{\mathcal{F}_{\text{Target}} - \tilde{\mathcal{F}}_{\text{Target}}^{(1)}}{\tilde{\mathcal{F}}_{\text{Target}}^{(2)}}. \quad (2.33)$$

The distribution time thus becomes a function of the target fidelity $\mathcal{F}_{\text{Target}}$ such that

$$T = T^{(1)} \cdot \frac{\tilde{\mathcal{F}}_{\text{Target}}^{(2)}}{\mathcal{F}_{\text{Target}} - \tilde{\mathcal{F}}_{\text{Target}}^{(1)}} + T^{(2)} + o(1). \quad (2.34)$$

This highlights the advantages of including the 2-photon contribution in the model: the fidelity is now pump parameter-dependent, coupled with the distribution time through a shared degree of freedom. In the ideal 1-photon scenario, the ion-ion entanglement is a pure Bell state by construction, leaving fidelity and distribution time decorrelated.

Figure 2.3 shows the distribution time T as a function of $\mathcal{F}_{\text{lon}}$ for different memory efficiencies η_{Mem} , for a repeater length $L_{\text{Chn}} = 500$ km and $L_{\text{Chn}} = 250$ km respectively.


i	ρ_i	α_i	β_i
1	$ 00\rangle$	$\frac{2-\eta_{\text{Swp}}}{4-\eta_{\text{Swp}}}$	$\frac{(-2+\eta_{\text{Swp}})(48-30\eta_{\text{Lnk}}+\eta_{\text{Swp}}(-20+13\eta_{\text{Lnk}}))}{2(4-\eta_{\text{Swp}})^2}$
2	$ \Psi^+\rangle$	$\frac{2}{4-\eta_{\text{Swp}}}$	$\frac{8\eta_{\text{Lnk}}+\eta_{\text{Swp}}(-80+\eta_{\text{Swp}}(28-17\eta_{\text{Lnk}})+44\eta_{\text{Lnk}})}{4(4-\eta_{\text{Swp}})^2}$
3	$ \Psi^-\rangle$	0	$\frac{(4-3\eta_{\text{Swp}})(4-3\eta_{\text{Lnk}})}{4(4-\eta_{\text{Swp}})}$
4	$ \Psi_2^-\rangle$	0	$\frac{4-3\eta_{\text{Lnk}}}{8-2\eta_{\text{Swp}}}$
5	$ \Psi_{a_+}^-\rangle$	0	$\frac{12-7\eta_{\text{Lnk}}-\sqrt{80-72\eta_{\text{Lnk}}+17\eta_{\text{Lnk}}^2}}{4(4-\eta_{\text{Swp}})}$
6	$ \Psi_{a_-}^-\rangle$	0	$\frac{12-7\eta_{\text{Lnk}}+\sqrt{80-72\eta_{\text{Lnk}}+17\eta_{\text{Lnk}}^2}}{4(4-\eta_{\text{Swp}})}$

Table 2.1: Table built from the expression of ρ_{chn} in Equation 2.20. Each term ρ_i of the mixture ρ_{chn} and their respective coefficients α_i and β_i . We recall $a_i = \alpha_i + \beta_i \cdot p + o(p)$.

i	j	$\mathcal{F}_{i,j} \cdot N_{i,j}$	$\mathbb{P}_{i,j}$
1	1	0	0
1	2	0	0
1	3	0	0
1	4	0	0
1	5	0	$(1-2a^2) \cdot \frac{\eta_{\text{lon}}^4}{8}$
1	6	0	$(1-2a^2) \cdot \frac{\eta_{\text{lon}}^4}{8}$
2	1	0	0
2	2	$\frac{\eta_{\text{lon}}^4}{8}$	$\frac{\eta_{\text{lon}}^4}{8}$
2	3	0	$\frac{\eta_{\text{lon}}^4}{8}$
2	4	$(4-3\eta_{\text{lon}}) \frac{\eta_{\text{lon}}^4}{8}$	$(4-3\eta_{\text{lon}}) \frac{\eta_{\text{lon}}^4}{8}$
2	5	$\frac{\eta_{\text{lon}}^4}{8} \left(2 - a_+ \sqrt{4-8a_+^2} (2-\eta_{\text{lon}}) - \eta_{\text{lon}} - a_+^2 \eta_{\text{lon}} \right)$	$\frac{\eta_{\text{lon}}^4}{4} \left(2(1-2a_+)(2-\eta_{\text{lon}}) + a_+^2(4-3\eta_{\text{lon}}) \right)$
2	6	$\frac{\eta_{\text{lon}}^4}{8} \left(2 - a_- \sqrt{4-8a_-^2} (2-\eta_{\text{lon}}) - \eta_{\text{lon}} - a_-^2 \eta_{\text{lon}} \right)$	$\frac{\eta_{\text{lon}}^4}{4} \left(2(1-2a_-)(2-\eta_{\text{lon}}) + a_-^2(4-3\eta_{\text{lon}}) \right)$

Table 2.2: Values of the fidelity $\mathcal{F}_{i,j}$ and 4-clicks heralding probability $\mathbb{P}_{i,j}$ for each term $\rho_i \otimes \rho_j$ of $\rho_{\text{chn}} \otimes \rho_{\text{chn}}$. The indices i and j refers to the index of Table 2.1.

Targeting a fidelity $\mathcal{F}_{\text{ion}} \approx 0.9$ with a memory efficiency $\eta_{\text{Mem}} \approx 0.5$, the distribution time reaches the order of 10 s at 250 km and 1 min at 500 km, demonstrating the feasibility of the architecture. As it will be discussed in the next section, we noticed a $1/\eta_{\text{Mem}}^5$ dependence on the memory efficiency, reflecting the six simultaneous memory uses along the protocol: one for the chain, and four for the 4-clicks heralding. This sensitivity to η_{Mem} highlights the central role of memory lifetime in the system: a long-lived memory keeps loaded subsystems available across multiple trials of the remaining ones, reducing the probability of starting from scratch. A different parallelization scheme could lower this exponent from six to four, as discussed in the next section.



(a) Distribution time T as a function of the targeted ion-pair fidelity $\mathcal{F}_{\text{ion}}$, for different memory efficiencies η_{Mem} , at a total distance of 250 km (elementary links of 125 km).

(b) Distribution time T as a function of the targeted ion-pair fidelity $\mathcal{F}_{\text{ion}}$, for different memory efficiencies η_{Mem} , at a total distance of 500 km (elementary links of 250 km).

Figure 2.3: Distribution time T of a full repeater as a function of the targeted ion-pair fidelity $\mathcal{F}_{\text{ion}}$, for several memory efficiencies η_{Mem} . We use practical parameters are taken from table 1.

Panel (a) corresponds to a total repeater distance of 250 km, and panel (b) to 500 km. For a target fidelity $\mathcal{F}_{\text{ion}} \approx 0.9$ and a memory efficiency $\eta_{\text{Mem}} \approx 0.5$, the distribution time reaches the order of 10 s at 250 km and 50 s at 500 km, demonstrating the feasibility of the architecture.

2.1.4 .Alternative protocol

The numerical results exhibit a $T \propto 1/\eta_{\text{Mem}}^5$ scaling for the distribution time. This exponent reflects the parallelization scheme of the architecture: the two repeater chains are loaded independently, contributing a single simultaneous memory use, while the ion swap requires the four edge memories of the chains to be released simultaneously, contributing to four joint memory uses. Memory efficiency η_{Mem} is part of the dominant experimental constraints of the protocol. Despite substantial efforts to mitigate memory loss, it remains a challenging aspects of the repeater architecture, and its impact on the distribution time motivates the search for parallelization schemes that lower the exponent [13]. As seen, the 1-photon contribution of such a parallelization gives the

probabilities

$$\begin{aligned}
\mathbb{P}_{\text{Lnk}} &= 2 \cdot \eta_{\text{Lnk}} \cdot p \\
\mathbb{P}_{\text{Swp}} &= \eta_{\text{Swp}} \cdot \frac{4 - \eta_{\text{Swp}}}{4} \\
\mathbb{P}_{\text{Ion}} &= \left(\frac{2}{4 - \eta_{\text{Swp}}} \right)^2 \cdot \frac{\eta_{\text{Ion}}^4}{8}
\end{aligned}
\quad \xrightarrow{\eta_{\text{Mem}} \ll 1} \quad
\begin{aligned}
&1 \\
\eta_{\text{Swp}} &\propto \eta_{\text{Mem}} \\
\eta_{\text{Ion}}^4 &\propto \eta_{\text{Mem}}^4
\end{aligned}
\tag{2.35}$$

We therefore consider an alternative scheme in which the ion swap is performed before the chain swap. This choice trades parallelization at the chain level for parallelization at the ion level: loading the chain now requires two simultaneous memory uses $\eta_{\text{Mem}}^1 \rightarrow \eta_{\text{Mem}}^2$, but each gateway swaps its ion independently, reducing the four simultaneous memory uses of the ion swap to two independent uses of two $\eta_{\text{Mem}}^4 \rightarrow \eta_{\text{Mem}}^2$. The overall scaling thus becomes $T \propto 1/\eta_{\text{Mem}}^4$, mitigating the memory-efficiency bottleneck at the cost of a more demanding chain-loading stage. Concretely, the 1-photon contribution of such a parallelization gives the probabilities

$$\begin{aligned}
\mathbb{P}_{\text{Lnk}} &= 2 \cdot \eta_{\text{Lnk}} \cdot p \\
\mathbb{P}_{\text{Swp}} &= \frac{2\eta_{\text{Swp}}^2}{(6 - \eta_{\text{Ion}})^2} \\
\mathbb{P}_{\text{Ion}} &= \frac{\eta_{\text{Det}}^2}{4} (3 - \eta_{\text{Ion}})
\end{aligned}
\quad \xrightarrow{\eta_{\text{Mem}} \ll 1} \quad
\begin{aligned}
&1 \\
\eta_{\text{Ion}}^2 &\propto \eta_{\text{Mem}}^2 \\
\eta_{\text{Ion}}^2 &\propto \eta_{\text{Mem}}^2
\end{aligned}
\tag{2.36}$$

In figure 2.4, we plot for both scenarios the distribution time T as a function of η_{Mem} for several values of $\mathcal{F}_{\text{Ion}}$. The respective $1/\eta_{\text{Mem}}^5$ and $1/\eta_{\text{Mem}}^4$ dependences are visible confirming the analysis above. We actually observe a power 5.3 and 4.3 dependance, and not a strict 5 and 4, since the derivation is only valid for the asymptotic limit $\eta_{\text{Mem}} \rightarrow 0$: the black curves match the asymptotic behavior at small η_{Mem} and deviate from the model curves as η_{Mem} increases. Still, scenario 2 outperforms scenario 1 over the whole range of η_{Mem} . This conclusion holds for the particular architecture studied here, with only two elementary links per chain. In a more general situation with $N = 2^n$ elementary links, loading a chain requires η_{Mem}^n simultaneous memory uses, so parallelizing the ion swap is advantageous only when $n \leq 2$.

2.2 .Matching photon waveforms

2.2.1 .The shape-mismatch problem

The heterogeneous network described above involves ion-based metropolitan networks with atomic-ensemble-based long-distance link. Atomic ensembles benefit from strong light-matter coupling and large multiplexing capacity, making them suitable for entanglement generation over long distances [12, 38], while denser metropolitan networks rely on individually controlled systems for multi-node synchronization and local processing [37, 38]. Such a heterogeneity raises the issue of interfacing disparate quantum




Figure 2.4: Distribution time as a function of the memory efficiency for different target fidelities, over a 500 km repeater link. The two scenarios are compared: protocol 1 (ion swap last) in solid lines, and protocol 2 (atom swap last) in dashed lines. The asymptotic power-law dependences in η_{Mem} are overlaid.

systems, where photons of different origin must interfere on a beam-splitter. Precisely, mismatched photon shapes degrades indistinguishability, reducing the heralded ion-ion entanglement quality [121].

In practice, considering $^{40}\text{Ca}^+$ ion metropolitan nodes [45] and $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ rare-earth-doped for long-distance backbone hubs memory [35], photons emitted by the ion are roughly one order of magnitude longer than the atomic-ensemble ones, with a relative overlap of 10%. Seeking high photon indistinguishability, [121] proposes to exploit the collective behaviour of the atomic frequency comb protocol (AFC) of the atomic memory to reshape the photon, stretching its temporal profile to match the ion one. Concretely, the protocol replaces the single readout π -pulse of the standard AFC echo by a sequence of weaker partial readout pulses, each releasing a controlled fraction of the stored excitation, so that the output photon is reconstructed bin by bin into a target temporal envelope matching the ion-emitted shape. See figure 2.5 for illustration. Such a procedure reaches an overlap of 90% with the ion waveform, while preserving most of the memory efficiency. The following paragraphs study the impact of such a stretching on the overall repeater performance.

2.2.2 .Analytical derivation

Let us provide analytical expressions of the ion-ion fidelity $\mathcal{F}_{\text{ion}}$ and the ion-swap probability $\mathbb{P}_{\text{ion}}$. First, let us focus on a single interface $(b^\dagger, d^\dagger)$, where the mode b carries the ion photon shape and d the AFC photon shape. Taking the atomic-ensemble photon




Figure 2.5: Illustration of the stretching procedure. A short input photon is stored in the atomic-ensemble memory. Based on the AFC protocol, a sequence of partial read-out pulses releases fractions of the stored excitation, shaping the output to match the photon emitted by the ion.

shape d as reference, we decompose b into two temporal modes: b_0 fully overlapping with d , and $b_\perp$ the residual mode orthogonal to d . Denoting $x = |\langle 0|b \cdot a^\dagger|0\rangle|^2$ the overlap parameter, the decomposition reads

$$b^\dagger = \sqrt{x} b_0^\dagger + \sqrt{1-x} b_\perp^\dagger. \quad (2.37)$$

We similarly denote η_{ion} and $\eta_{\text{ion}}^\perp$ the detection efficiencies of the reference and orthogonal modes, with $\eta_{\text{ion}} \approx \eta_{\text{ion}}^\perp$ in practice. From appendix C.3, the beam-splitter relation then reads

$$\begin{aligned} b^\dagger &\xrightarrow{\text{BS}} \frac{b_0^\dagger + d_0^\dagger}{\sqrt{2}}, \\ d^\dagger &\xrightarrow{\text{BS}} \sqrt{x} \cdot \frac{b_0^\dagger - d_0^\dagger}{\sqrt{2}} + \sqrt{1-x} \cdot \frac{b_\perp^\dagger - d_\perp^\dagger}{\sqrt{2}}. \end{aligned} \quad (2.38)$$

and the no-click POVM element follows as

$$\begin{aligned} \Pi_0^b &= (1 - \eta_{\text{ion}})^{b_0^\dagger b_0} \cdot (1 - \eta_{\text{ion}}^\perp)^{b_\perp^\dagger b_\perp} \\ \Pi_1^b &= \mathbf{1} - \Pi_0^b \end{aligned} \quad (2.39)$$

Restricting the model to the 1-photon contribution, we recall that the link-chain state is actually a mixture between a Bell state and the vacuum due to the swaps, and we additionally assume an imperfect ion emission

$$\rho_{\text{chn}}^{\text{ini}} = (1 - \alpha_{\text{chn}}) |00\rangle \langle 00| + \alpha_{\text{chn}} |\Psi^+\rangle \langle \Psi^+|, \quad (2.40)$$

$$\rho_{\text{ion}}^{\text{ini}} = (1 - \alpha_{\text{ion}}) |00\rangle \langle 00| + \alpha_{\text{ion}} |\Psi^+\rangle \langle \Psi^+|. \quad (2.41)$$

The same modelization applies to the other interfaces $(c^\dagger, f^\dagger)$, $(e^\dagger, h^\dagger)$, and $(g^\dagger, i^\dagger)$, providing a suitable framework to extend the calculation.

Based on the modelization above and the full derivation of appendix C.4 the ion-swap probability reads

$$\mathbb{P}_{\text{ion}} = \alpha_{\text{chn}}^2 \alpha_{\text{ion}}^2 \frac{\eta_{\text{ion}}^2 \cdot (\eta_{\text{ion}}^\perp + x \cdot \delta)^2}{8}, \quad (2.42)$$

heralding the state

$$\begin{aligned} \rho_{\text{ion}} = & \left(\frac{1}{2} + \alpha_{\text{chn}}^2 \alpha_{\text{ion}}^2 \frac{\eta_{\text{ion}}^2 \cdot x^2}{2 \cdot (\eta_{\text{ion}}^\perp + x \cdot \delta)^2} \right) \cdot |\Psi^+\rangle \langle \Psi^+| \\ & + \left(\frac{1}{2} - \alpha_{\text{chn}}^2 \alpha_{\text{ion}}^2 \frac{\eta_{\text{ion}}^2 \cdot x^2}{2 \cdot (\eta_{\text{ion}}^\perp + x \cdot \delta)^2} \right) \cdot |\Psi^-\rangle \langle \Psi^-|, \end{aligned} \quad (2.43)$$

yielding the fidelity

$$\mathcal{F}_{\text{ion}} = \frac{1}{2} + \alpha_{\text{chn}}^2 \alpha_{\text{ion}}^2 \frac{\eta_{\text{ion}}^2 \cdot x^2}{2 \cdot (\eta_{\text{ion}}^\perp + x \cdot \delta)^2}. \quad (2.44)$$

A short qualitative study of the extreme scenarios $x = 0, 1$ gives:

- **Perfect overlap:** as expected, we recover the result from section 2.1 when considering a perfect overlap $x = 1$. The ion-ion state is a pure Bell state, that is, a completely entangled state

$$\rho_{\text{ion}} = |\Psi^+\rangle \langle \Psi^+|, \quad (2.45)$$

and the ion swap probability reads

$$\mathbb{P}_{\text{ion}} = \frac{\eta_{\text{ion}}^4}{8}. \quad (2.46)$$

- **No overlap:** on the contrary, considering orthogonal modes $x = 0$, the ion and links-chain photons do not interfere, resulting in an incoherent state which does not carry entanglement

$$\rho_{\text{ion}} = \frac{1}{2} |01\rangle \langle 01| + \frac{1}{2} |10\rangle \langle 10|,$$

with a probability

$$\mathbb{P}_{\text{ion}} = \frac{\eta_{\text{ion}}^2 \cdot \eta_{\text{ion}}^{\perp 2}}{8}.$$

In practice, the efficiency of detecting a photon in the reference mode or the orthogonal mode are very close. The approximation $\frac{\delta}{\eta_{\text{ion}}} \ll 1$ with a first order expansion shows

$$\mathcal{F}_{\text{ion}}(x) = \frac{1}{2} + \frac{x^2}{2} \left(1 + (1-x) \frac{\delta}{\eta_{\text{ion}}} \right) \quad \text{and} \quad \mathbb{P}_{\text{ion}}(x) = \frac{\eta_{\text{ion}}^4}{8} \cdot \left(1 + (x-1) \frac{\delta}{\eta_{\text{ion}}} \right) \quad (2.47)$$

since $0 \leq x \leq 1$ the fidelity and the probability exhibit opposite behavior with respect to δ . This exhibits the usual trade-off between entanglement quality and distribution rate.

2.2.3 .Results

The calculation above gives a practical estimation of the influence of the overlap on the repeater efficiency which can be easily implemented numerically. In this section, we provide the numerical implementation of a full repeater link of two elementary links, as studied in section 2.1.2, with a shape mismatch x . Figure 2.6 shows the distribution time according to the shaping mismatch x for different values of the ion-pair fidelity. In this numerical implementation, we use a derivation up to 2-photon contribution, providing a finer modelization of the impact of the mismatch. In practice, without stretching, one has $x = 0.1$, the stretching method from [121] (shaping with cropped-echo) the overlap reaches $x = 0.9$.

Figure 2.6 shows that some fidelities become unreachable for low overlap values. This bound can be derived analytically from equation (2.47) providing the 1-photon contribution to the ion-pair fidelity

$$\tilde{\mathcal{F}}_{\text{lon}}^{(1)}(x) = \frac{1}{2} + \frac{x^2}{2} \left(1 + (1-x) \frac{\delta}{\eta_{\text{lon}}} \right), \quad (2.48)$$

under the assumption $\delta \ll 1$. Since the 1-photon sector is the only one contributing to the Bell-state component (higher-order sectors only dilute the pair), the maximum reachable fidelity is set by $\tilde{\mathcal{F}}_{\text{lon}}^{(1)}(x)$. This gives a minimum overlap to reach a target fidelity $\mathcal{F}_{\text{lon}}$:

$$\mathcal{F}_{\text{lon}}(x) \leq \tilde{\mathcal{F}}_{\text{lon}}^{(1)}(x) \approx \frac{1+x^2}{2} \implies x_{\min} = \sqrt{2\mathcal{F}_{\text{lon}} - 1}. \quad (2.49)$$

In the bare scenario without stretching, $x = 0.1$, the maximum reachable fidelity is bounded to $\mathcal{F}_{\text{lon}}^{\max} = 0.505$. With the stretching method, $x = 0.9$ raises this bound to $\mathcal{F}_{\text{lon}}^{\max} = 0.9$. This quantitative gap motivates the necessity of a shaping procedure in any practical implementation of the ion-atom interface.

2.2.4 .Mixture

Up to this point, we have assumed a simple model in which the initial state $\rho_{\text{lon}}^{\text{ini}}$ is pure (up to a vacuum component). However, the actual state emitted by an ion is rather described by a continuous mixture of non-orthogonal pure components. Here, we generalize the previous calculation to this mixed-state scenario.

Concretely, the initial state generated by Alice's ion $\rho_{\text{lon}}^{\text{ini,A}}$ reads

$$\rho_{\text{lon}}^{\text{ini,A}} = \int P_A(\gamma) \left[\frac{a_H^\dagger \tilde{b}^\dagger(\gamma) + a_V^\dagger \tilde{c}^\dagger(\gamma)}{\sqrt{2}} \right] |0\rangle \langle 0| \left[\frac{a_H \tilde{b}(\gamma) + a_V \tilde{c}(\gamma)}{\sqrt{2}} \right] d\gamma \quad (2.50)$$

with P_A the mode emission probability distribution, such that $\int P_A = 1$, and γ the parameter indexing the family of non-orthogonal temporal modes for the spatial modes b and c . Thus, neglecting the vacuum component for readability, the initial state is a




Figure 2.6: Distribution time as a function of the overlap x for different target fidelities, over a 500 km repeater link.

parametrized mixture over γ_A and γ_B

$$\rho_{\text{ini}} = \rho_{\text{ion}}^A \otimes |\Psi^+\rangle \langle \Psi^+| \otimes |\Psi^+\rangle \langle \Psi^+| \otimes \rho_{\text{ion}}^B \quad (2.51)$$

$$= \iint \rho^A(\gamma_A) |\Psi^+\rangle \langle \Psi^+| \otimes |\Psi^+\rangle \langle \Psi^+| \otimes \rho^B(\gamma_B) d\gamma_A d\gamma_B \quad (2.52)$$

$$= \iint P_{AB} \bar{\rho}(\gamma_A, \gamma_B) d\gamma_A d\gamma_B. \quad (2.53)$$

By construction, the ion-swap protocol relies on a non-trace-preserving quantum channel $\mathcal{E}_{\text{ion}}$. Thus linearity of the channel naturally implies that the protocol applied to a mixture is the mixture of the protocol. However, $\mathcal{E}_{\text{ion}}$ has been built in the scope of a single temporal mode, implicitly fixing a specific shape γ_0 for each photon, so one needs to generalize the channel to temporal mixtures. As a basis-invariant operation, the partial trace extends directly to the temporal mixture, but the POVM element requires more care. A naive extension such as $\Pi_{\text{ion}} = \int \Pi_{\text{ion}}^\gamma d\gamma$ fails, as the modes $\{b_\gamma\}_\gamma$ are not orthogonal and this expression does not define a valid POVM on the full mode family. A shape-independent POVM, valid for any photon shape γ of a given spatial port, is therefore required to apply the linearity argument.

We first build a no-click POVM element Π_0^b acting uniformly on every shape of a spatial mode $b^\dagger$. From appendix C.5, given an orthonormal basis $\{\beta_\nu\}_\nu$ of the spatial mode b , the operator

$$\Pi_0^b = R_{\text{ion}}^\dagger \beta_\nu^\dagger \beta_\nu d\nu \quad (2.54)$$

is diagonal in the one-excitation subspace with respect to the family $\{b_\gamma^\dagger |0\rangle\}_\gamma$

$$\Pi_0^b |0\rangle = |0\rangle, \quad (2.55)$$

$$\Pi_0^b b_\gamma^\dagger |0\rangle = R_{\text{ion}} b_\gamma^\dagger |0\rangle. \quad (2.56)$$

Therefore, this POVM is suitable for detection on any shape of the spatial mode $b^\dagger$. Extending the construction to the remaining spatial modes c, d, e, f, g, h, i yields a heralding POVM Π_{ion} robust to the mixture of the ion's emission.

With this unified POVM, the partial trace and the heralding measurement both act linearly on the mixture, and the protocol of the mixture reduces to the mixture of the protocol. Plugging this argument into the pure-state result of the previous section, the heralded fidelity and probability of the full link read

$$\mathcal{F}_{\text{ion}} = \iint P_A(\gamma_A) P_B(\gamma_B) \mathcal{F}_{\text{ion}}^{\gamma_A, \gamma_B} d\gamma_A d\gamma_B, \quad (2.57)$$

$$\mathbb{P}_{\text{ion}} = \iint P_A(\gamma_A) P_B(\gamma_B) \mathbb{P}_{\text{ion}}^{\gamma_A, \gamma_B} d\gamma_A d\gamma_B, \quad (2.58)$$

where $\mathcal{F}_{\text{ion}}^{\gamma_A, \gamma_B}$ and $\mathbb{P}_{\text{ion}}^{\gamma_A, \gamma_B}$ are the pure-state expressions evaluated at the shapes (γ_A, γ_B) .

2.3 .Multiplexing

2.3.1 .Stretching and multiplexing

To boost entanglement distribution rate, we exploit the multiplexing capability of hubs' atomic-ensemble memories: several Bell-pair attempts can be stored in parallel while waiting for classical feedback from the central station. In this picture, the multi-mode capacity acts as a buffer: the cooldown time is used to attempt $M = \frac{L}{c \tau_{\text{bin}}}$ trials per cycle, where τ_{bin} is the time-bin duration. The resulting trial rate is $\nu = \frac{M}{L/c}$, i.e. an improvement by a factor M

$$\mathbb{P}_{\text{Lnk}} \rightarrow \frac{1}{M} \mathbb{P}_{\text{Lnk}}. \quad (2.59)$$

Photon stretching modifies this picture: a stretched photon of duration $\tau_{\text{str}} > \tau_{\text{bin}}$ spans

$$M_{\text{str}} = \frac{\tau_{\text{str}}}{\tau_{\text{bin}}}. \quad (2.60)$$

The time bin serves as the temporal reference, in practice, since $\tau_{\text{atom}} \leq \tau_{\text{str}} \leq \tau_{\text{ion}}$, we take the atomic-ensemble memory photon duration as reference, $\tau_{\text{bin}} = \tau_{\text{atom}}$. Ideally, the stretched photon should not overlap a subsequent time bin that would have been filled by a second photon (seef figure 2.7). That is, the signal photon should be followed by $M_{\text{str}} - 1$ empty bins, a requirement we refer to as the *isolated photon constraint*. While photon stretching is essential to achieve photon matching and improve fidelity, its use in a multiplexed architecture also imposes constraints that may limit repeater performance, which we aim to evaluate here.

If the *isolated photon* constraint is not satisfied, we may get *detectable unwanted events*: two successful heraldings within the same M_{str} window. Such events may be either discard, reducing the distribution rate but preserving high fidelity, or accepted, maintaining high distribution rate, but lowering the fidelity because of imperfection. This highlight the usual trade-off between distribution rate and entanglement quality. In this work, we adopt the first strategy where the detectable unwanted events are filtered to preserve high fidelity. We thus distinguish the heralding probability $\mathbb{P}_{\text{Chn}}$, announcing the presence of a delocalized excitation at the two edge hubs of the chain, from the loading probability $\mathbb{P}_{\text{Load}}$, announcing such an excitation followed by $M_{\text{str}} - 1$ empty bins, with $\mathbb{P}_{\text{Load}} \leq \mathbb{P}_{\text{Chn}}$ by construction. See figure 2.7 for illustration.

Actually, not all unwanted events can be detected. For example, if the SPDC source generates a photon pair with a lost idler photon, the signal photon may still be stored in the memory without any heralding. Such missed-heralded events are invisible to the central station and cannot be filtered out. If they occur within the M_{str} window, they lead to what called *undetectable unwanted event*. They generate an *imperfect state* ρ_{imp} reducing the ion-ion fidelity. These events cannot be neglected: since $\eta_{\text{trs}} \ll 1$ a signal photon may reach the local hub memory while its idler is lost on the way to the long-distance central station with non-negligible probability.

This degradation cannot be left unaddressed and must be quantified, photon stretching method may undermine the multiplexing advantage and thus reduce the overall performance of the repeater. To evaluate the impact of stretching on the multiplexing capability and repeater performance, we consider two realistic scenarios with and without stretching, highlighting their respective advantages and limitations. We define $M_{\text{ion}} = \tau_{\text{ion}}/\tau_{\text{bin}}$ as the number of bins spanned by the ion waveform, so that the shape overlap relates to the bin counts with $x = \frac{\tau_{\text{str}}}{\tau_{\text{ion}}} = \frac{M_{\text{str}}}{M_{\text{ion}}}$.

- **Scenario 1:** we consider no stretching, exploiting the full potential of the multiplexing, concretely $M_{\text{str}} = 1$ and

$$M_{\text{str}} = 1 \quad \text{and} \quad x = \frac{1}{M_{\text{ion}}} < 1. \quad (2.61)$$

Mainly, the photon shape mismatch x introduces distinguishability and suppresses part of the interference, resulting in a lower fidelity. Additionally, the photon shape mismatch x may slightly affect the ion-swap probability, depending on whether $\eta^\perp = \eta$ or not.

- **Scenario 2:** we assume a near-perfect stretching and study its impact on the multiplexing capabilities, concretely

$$M_{\text{str}} = M_{\text{ion}} \quad \text{and} \quad x = \frac{M_{\text{str}}}{M_{\text{ion}}} = 1. \quad (2.62)$$

In this context, the isolated photon constraint imposes to discard detectable un-




Figure 2.7: Illustration of the isolated photon constraint. The stretched photon spans M_{str} time bins, while the native AFC photon spans a single time bin as a reference. The constraint requires the $M_{\text{str}} - 1$ time bins following the photon of interest to be empty, so that no extra SPDC emission overlaps with the stretched photon window.

wanted events, which decreases the entanglement distribution rate. Additionally, the contribution of missed-heralded events that populate the $M_{\text{str}} - 1$ following windows may reduce the ion-ion fidelity.

2.3.2 . Qualitative study

We first focus on a qualitative study, aiming to evaluate the performance of the quantum repeater in the two scenarios. For simplicity, we consider minimal architecture where each chain is composed of a single elementary link, $\mathbb{P}_{\text{chn}} = \mathbb{P}_{\text{Lnk}}$, of length $L_{\text{chn}} = L_{\text{Lnk}} = 250$ km. Nevertheless, as it will be seen in the quantitative study, the general behavior extends to a 2-links chain. For clarity, we assume a symmetric parameters on the two back-bone edges, as well as identical detection efficiencies for the reference and orthogonal modes $\eta_{\text{ion}} = \eta_{\perp}$.

The two relevant metrics are the ion-ion fidelity $\mathcal{F}_{\text{ion}}$ and the distribution time T . The modelization of the two scenarios involves three quantities through which $\mathcal{F}_{\text{ion}}$ and T may be affected. First, the repeater chain loading probability $\mathbb{P}_{\text{Load}}$ may be affected by the discarding of unwanted events, since we consider the chain loaded when a single photon populates the M_{str} bins. Second, the ion-swap probability $\mathbb{P}_{\text{ion}}(x)$ may be affected by the mismatch x . Finally, the fidelity $\mathcal{F}_{\text{ion}}(x)$ is affected in two ways: by the mismatch x and the missed-heralded events, which cannot be filtered and may produce imperfect state ρ_{imp} . We therefore focus on $\mathbb{P}_{\text{Load}}$, $\mathbb{P}_{\text{ion}}(x)$, and $\mathcal{F}_{\text{ion}}(x)$.

Scenario 1 In this scenario, we assume no stretching and all heralded events are kept, thus the loading probability reduces to the heralding probability of the chain

$$\mathbb{P}_{\text{Load}} = \mathbb{P}_{\text{chn}} = \mathbb{P}_{\text{Lnk}} = 2 \eta_{\text{Lnk}} p + \eta_{\text{Lnk}} (1 - 2\eta_{\text{Lnk}}) p^2 + o(p^2), \quad (2.63)$$

where the last equality comes from the fact that the chain is composed of a single elementary link. Additionally, from equation (2.42) and $\delta = \eta_{\text{lon}} - \eta_{\perp} = 0$, the ion-swap probability $\mathbb{P}_{\text{lon}}$ is not affected

$$\mathbb{P}_{\text{lon}}(x) = \frac{\eta_{\text{lon}} \cdot (\eta_{\perp} + x \delta)^2}{8} = \frac{\eta_{\text{lon}}^4}{8}. \quad (2.64)$$

Thus, the overall distribution rate is not affected by the mismatch. The fidelity reads

$$\mathcal{F}_{\text{lon}}(x) = \tilde{\mathcal{F}}_{\text{lon}}^{(1)}(x) + \tilde{\mathcal{F}}_{\text{lon}}^{(2)}(x) p + o(p), \quad (2.65)$$

with $\tilde{\mathcal{F}}_{\text{lon}}^{(1)}(x)$ given in equation (2.47). The 2-photon sector does not contribute to the generation of Bell pairs, and only degrades the fidelity with the dilution correction (c.f. equation (1.29)). Concretely, $\tilde{\mathcal{F}}_{\text{lon}}^{(2)}(x) < 0$, and $\tilde{\mathcal{F}}_{\text{lon}}^{(1)}(x)$ represents the maximal achievable fidelity, reached in the ideal limit $p \rightarrow 0$. Therefore, setting $x = 1/M_{\text{str}}$, the ion-ion fidelity is bounded by

$$\mathcal{F}_{\text{lon}}(x) < \frac{1 + x^2}{2} = \frac{M_{\text{str}}^2 + 1}{2 M_{\text{str}}^2}. \quad (2.66)$$

In practice $M_{\text{str}} = \frac{\tau_{\text{str}}}{\tau_{\text{bin}}} = 10$, thus the bound stands

$$\mathcal{F}_{\text{lon}}(1/M_{\text{str}}) < 0.505, \quad (2.67)$$

showing the unfeasibility of distributing high quality entangled with photon mismatch.

Scenario 2 In this scenario, the AFC photon is stretched and the isolated photon constraint applies. Requiring a single photon followed by $M_{\text{str}} - 1$ empty bins, the loading probability reads

$$\mathbb{P}_{\text{Load}} = \mathbb{P}_{\text{Chn}} (1 - \mathbb{P}_{\text{Chn}})^{M_{\text{str}}-1}. \quad (2.68)$$

For a single-link chain, $\mathbb{P}_{\text{Chn}} = \mathbb{P}_{\text{Lnk}} = 2 \eta_{\text{Lnk}} p$. Since $\eta_{\text{Lnk}} p \ll 1$, the loading probability reduces to

$$\mathbb{P}_{\text{Load}} = (1 - 2(M_{\text{str}} - 1) \eta_{\text{Lnk}} p) 2 \eta_{\text{Lnk}} p, \quad (2.69)$$

so filtering unwanted events lowers the probability by a factor $\alpha_{\mathbb{P}}^{\downarrow} = 1 - 2(M_{\text{str}} - 1) \eta_{\text{Lnk}} p$. Realistic numbers $p = 10^{-2}$, $\eta_{\text{Lnk}} = 3 \cdot 10^{-3}$ and $M_{\text{str}} = 10$ give $\alpha_{\mathbb{P}}^{\downarrow} = 1 - 5.4 \cdot 10^{-4} = \mathcal{O}(10^{-3})$, which is not perceptible. Regarding the fidelity, the distinguishability vanishes since $x = 1$, and the only source of degradation comes from the missed-heralded events that may populate the $M_{\text{str}} - 1$ bins. For a single-link chain, the probability of a missed-heralded photon reads

$$\mathbb{P}_{\text{missed}} = (1 - \eta_{\text{Lnk}}) p. \quad (2.70)$$

The probability of generating an imperfect state within the $M_{\text{str}} - 1$ window follows

$$\mathbb{P}_{\text{imp}} = 1 - (1 - \mathbb{P}_{\text{missed}})^{M_{\text{str}}-1} = (M_{\text{str}} - 1) R_{\text{Lnk}} p, \quad (2.71)$$

yielding a heralded state

$$\rho_{\text{Load}} = (1 - \mathbb{P}_{\text{imp}}) \rho_{\text{Chn}} + \mathbb{P}_{\text{imp}} \rho_{\text{imp}}. \quad (2.72)$$

In a restricted scenario, we assume ρ_{imp} does not contribute to the fidelity (i.e. we set its fidelity to zero), giving a conservative lower bound. For two rails, the resulting fidelity is $\mathcal{F}_{\text{ion}}^{\downarrow} = (1 - \mathbb{P}_{\text{imp}})^2 \mathcal{F}_{\text{ion}}$, which defines the reduction factor

$$\alpha_{\mathcal{F}}^{\downarrow} = (1 - \mathbb{P}_{\text{imp}})^2 \approx 1 - 2 (M_{\text{str}} - 1) R_{\text{Lnk}} p, \quad (2.73)$$

requiring $M_{\text{str}} \cdot p \ll 1$ to get $\alpha_{\mathcal{F}}^{\downarrow}$ close to 1. Numerically, a targeted fidelity $\mathcal{F}_{\text{Target}} = 0.9$ with $M_{\text{str}} = 10$ bins leads to $\alpha_{\mathcal{F}}^{\downarrow} = 0.9$, yielding a final fidelity $\mathcal{F}_{\text{ion}} \approx 0.8$. This may trigger a reactive behavior where a lower p is required to restore $\mathcal{F}_{\text{ion}} = 0.9$, lowering the ion-swap probability $\mathbb{P}_{\text{ion}}$. This coupling needs a sharper numerical study which accounts for all joint effects.

This qualitative study exhibits two down factors for Scenario 2: a probability and a fidelity denoted $\alpha_{\mathbb{P}}^{\downarrow}$ and $\alpha_{\mathcal{F}}^{\downarrow}$ respectively. The probability down factor, as a consequence of the filtering, calls for $M_{\text{str}} \eta_{\text{Lnk}} p \ll 1$, seeking for low unwanted heralded photons (governed by $\eta_{\text{Lnk}} p$) within the empty-bin constraint (governed by M_{str}). On the contrary, the fidelity down factor comes from missed-heralded events and calls for $M_{\text{str}} R_{\text{Lnk}} p \ll 1$, a regime of high transmission probability (governed at first order by $R_{\text{Lnk}} p$) with low empty-bin constraint (governed by M_{str}). Thus, while sharing the same base $M_{\text{str}} p$, the two factors compete: $\alpha_{\mathbb{P}}^{\downarrow}$ favors a short repeater ($R_{\text{Lnk}} \ll 1$), whereas $\alpha_{\mathcal{F}}^{\downarrow}$ favors a long one ($\eta_{\text{Lnk}} \ll 1$). This competition is balanced in the next section with a numerical study.

In summary, this qualitative study showed that keeping the native photon shape does not affect the entanglement distribution time, but the fidelity is bounded by $\mathcal{F}_{\text{ion}} < 0.505$, which makes the entanglement unusable. A quick estimation of the isolated photon constraint impact on the fidelity in the stretched scenario shows a reduction of about 10% due to missed-heralded photons, while discarding events polluted by extra photons in the M_{str} window is almost invisible. As expected, the impact of the isolation constraint on the fidelity cannot be neglected and must be taken into account in the model. This qualitative study also exhibits a common quantity $M_{\text{str}} p$ as a figure of merit both for probability and fidelity down factor. This comes from the physical interpretation that the probability of generating a photon pair across M_{str} bins reads $M_{\text{str}} p$.

2.3.3 . Numerical study

We now turn to a numerical study of the discussion above. We extend the implementation of section 2.2 with the isolation constraint developed in the previous section,


which accounts for all the effects described above. Concretely, we give numbers for the implementation a 500 km dual-rail full repeater link over a 2-link chain, up to the contribution of the 2-photon sector, with shaping mismatch modelization x and its influence on the multiplexing capabilities.

First, we compare Scenario 1 and Scenario 2 in figure 2.8a, plotting the ion-swap probability $\mathbb{P}_{\text{ion}}$ as a function of the targeted ion-pair fidelity $\mathcal{F}_{\text{ion}}$, for several values of the mode constraint M_{str} . As $\tau_{\text{ion}}/\tau_{\text{bin}} = M_{\text{str}}$, we set $x = 1/M_{\text{str}}$ for Scenario 1. The ideal case $x = 1$, $M_{\text{str}} = 1$ is shown as a reference. As expected, in Scenario 1 most fidelities are out of reach, a direct consequence of


$$\mathcal{F}_{\text{ion}}(x) < \frac{1+x^2}{2} = \frac{M_{\text{str}}^2 + 1}{2M_{\text{str}}^2}. \quad (2.74)$$

Scenario 2, on the contrary, stays robust with respect to M_{str} : at a targeted fidelity $\mathcal{F}_{\text{ion}} = 0.8$ with $M_{\text{str}} = 10$, the ion-swap probability only drops from $7.5 \cdot 10^{-5}$ to $4 \cdot 10^{-5}$. The probability is also less sensitive to the constraint as the targeted fidelity approaches 1, which works in our favor.

The qualitative study brought out two down factors, $\alpha_{\mathcal{F}}^{\downarrow} \approx 1 - 2M_{\text{str}}R_{\text{Lnk}}p$ and $\alpha_{\mathbb{P}}^{\downarrow} \approx 1 - 2M_{\text{str}}\eta_{\text{Lnk}}p$, degrading respectively the ion-pair fidelity $\mathcal{F}_{\text{ion}}$ and the ion-swap probability $\mathbb{P}_{\text{ion}}$. We thus look for a regime which need to be balance: setting $M_{\text{str}}R_{\text{Lnk}}p \ll 1$ to suppress missed-herald events (i.e. $\alpha_{\mathcal{F}}^{\downarrow}$ calls for short-distance link), while keeping $M_{\text{str}}\eta_{\text{Lnk}}p \ll 1$, to suppress heralded unwanted events (i.e. $\alpha_{\mathbb{P}}^{\downarrow}$ calls for long-distance link). Obviously, both cannot hold at once since $R_{\text{Lnk}} = 1 - \eta_{\text{Lnk}}$, and the qualitative study cannot conclude the trade-off. We thus provides a numerical study, with figure 2.8b plotting the two down factors as a function of the repeater length, for several values of the mode constraint M_{str} . We first notice the plot shows the expected behavior: a short repeater favors the fidelity, while a long repeater favors the probability. The overall balance, however, turns to long-distance repeaters, with a transition for $L_{\text{chn}} \approx 50$ km. For $L_{\text{chn}} > 100$ km, $\alpha_{\mathbb{P}}^{\downarrow}$ improves sharply from 0.85 to nearly 1, consistent with $\alpha_{\mathbb{P}}^{\downarrow} = \mathcal{O}(10^{-3})$ in the qualitative analysis, while $\alpha_{\mathcal{F}}^{\downarrow}$ degrades weakly, only from 0.98 to 0.95. This agrees with the qualitative study, which identified, at $L_{\text{chn}} = 250$ km, the fidelity loss from undetected unwanted events as the main source of degradation.



(a) Ion-swap probability $\mathbb{P}_{\text{ion}}$ as a function of the targeted ion-pair fidelity $\mathcal{F}_{\text{ion}}$, for several values of the mode constraint M_{str} , over a 500 km 2-links chain. Dashed lines show Scenario 1 (no stretching with $x = 1/M_{\text{str}}$), solid lines Scenario 2 (near-perfect stretching $x \approx 1$ and $M_{\text{str}} > 1$), and the dash-dotted line the ideal case $x = 1$, $M_{\text{str}} = 1$. In Scenario 1 the fidelity is bounded by $(1 + x^2)/2$, leaving most targets out of reach, and Scenario 2 stays robust with respect to M_{str} .



(b) Probability and fidelity down factor, $\alpha_{\mathbb{P}}^{\downarrow}$ and $\alpha_{\mathcal{F}}^{\downarrow}$, as a function of the repeater length L_{chn} , for several values of the mode constraint M_{str} , over a 500 km chain and a targeted fidelity of $\mathcal{F}_{\text{ion}} = 0.9$. A short repeater favors the fidelity, suppressing missed-herald events, while a long repeater favors the probability, suppressing heralded unwanted events. The balance favor long distances: for $L_{\text{chn}} > 100$ km, $\alpha_{\mathbb{P}}^{\downarrow}$ rises sharply to nearly 1 while $\alpha_{\mathcal{F}}^{\downarrow}$ degrades only weakly.

3 -Exact modeling of an elementary link

The previous chapter introduced a model of the complete quantum repeater architecture, from the description of an elementary link up to the interface between the long-distance backbone and metropolitan networks. This model relies on a perturbative expansion of the source state, truncated at second order in the pump parameter p . While this expansion captures the essential behavior of the architecture, it does not take into account the contributions beyond the two-photon sector and therefore remains valid only in the low-pump regime. In this chapter, we lift this restriction and derive an exact model of the elementary link, accounting for the full multi-photon contribution of the source. From this exact treatment, we derive the link state in the Fock basis, the entanglement distribution rate and fidelity, the local detection statistics, the coherences from single-photon interferometry, and an entanglement witness.

3.1 .Exact computation

3.1.1 .State shared by Alice and Bob

Consider the set-up from Figure 3.1. Two SPDC sources feed a 50/50 beam-splitter mixing the modes $c_A^\dagger$ and $c_B^\dagger$, followed by a *click/no-click* heralding measurement. The efficiencies of the central station, Alice's hub, and Bob's hub are denoted η_{Lnk} , η_A , and η_B , with associated losses $R_{\text{Lnk}} = 1 - \eta_{\text{Lnk}}$, R_A , and R_B . The dark-counts are dc_{Lnk} , dc_A , dc_B , and the pump parameters of the two SPDC sources are p_A and p_B . The state shared by Alice and Bob reads

$$\rho_{\text{Lnk}} = \frac{1}{\mathcal{N}} \text{Tr}_{c_A, c_B} \left[\sqrt{\Pi_{\text{Lnk}}} \cdot U_{\text{BS}} \cdot \rho_{\text{SPDC}}^A \otimes \rho_{\text{SPDC}}^B \cdot U_{\text{BS}}^\dagger \cdot \sqrt{\Pi_{\text{Lnk}}} \right], \quad (3.1)$$

where ρ_{SPDC}^A and ρ_{SPDC}^B are the states emitted by the two SPDC sources, U_{BS} is the beam-splitter evolution, Π_{Lnk} is the POVM associated with the *click/no-click* heralding event at the central station, and $\mathcal{N}$ the normalized constant.

The full derivation of ρ_{Lnk} is detailed in appendix D.1. In a nutshell, this derivation relies on building a non-unitary matrix M that encodes both the mode mixing induced by the beam-splitter and the photon losses along the process. Based on singular value decomposition (SVD) of this matrix M , we define two rotated mode bases, $m_A, m_B \rightarrow m_\sigma^+, m_\sigma^-$ on the memory side and $c_A, c_B \rightarrow c_\sigma^+, c_\sigma^-$ on the central-station side:

$$\exp \left[\begin{pmatrix} c_A^\dagger & c_B^\dagger \end{pmatrix} M \begin{pmatrix} m_A^\dagger \\ m_B^\dagger \end{pmatrix} \right] = \exp \left[\sigma_+ \cdot c_\sigma^{+\dagger} m_\sigma^{-\dagger} + \sigma_- \cdot c_\sigma^- m_\sigma^{--} \right]. \quad (3.2)$$




Figure 3.1: Optical setup of an elementary link. Each hub hosts an SPDC source and an atomic-ensemble memory. The source, parameterized by a pump parameter p_A (resp. p_B), generates signal-idler photon pairs in the memory mode m_A (resp. m_B) and the long-distance channel mode c_A (resp. c_B). The central station mixes the modes c_A and c_B on a 50 – 50 beam-splitter, and a single click projects the memory modes onto an entangled state.

Using the basis-invariant of the trace $\text{Tr}_{c_A, c_B} = \text{Tr}_{c_\sigma^+, c_\sigma^-}$ and the identity

$$|\psi\rangle = \exp \left[\sigma_\pm c_\sigma^{\pm\dagger} m_\sigma^{\pm\dagger} \right] |00\rangle \implies \text{Tr}_{c_\sigma^\pm} [|\psi\rangle \langle\psi|] = \rho_{\text{Therm}}^{m_\sigma^\pm}(\sigma_\pm), \quad (3.3)$$

with $\rho_{\text{Therm}}^{m_\sigma^\pm}(\sigma_\pm)$ the thermal state as defined below, the calculation reduces to the SVD computation of the matrix M . The resulting state is a thermal state ρ_{Therm} expressed in the rotated modes m_σ and m_ω , in which the entanglement is carried. The procedure is equivalent to a standard Gaussian-optics treatment, here we keep the Fock-space representation as it offers a clearer physical interpretation in the low-pump regime and matches the diagonal structure of the heralding POVM in the photon-number basis.

The derivation, detailed in appendix D.1 and appendix D.2, yields

$$\rho_{\text{Lnk}} = \frac{1}{1 - \xi} \cdot \rho_{\text{Therm}}^{m_\sigma^+}(\sigma_+) \otimes \rho_{\text{Therm}}^{m_\sigma^-}(\sigma_-) - \frac{\xi}{1 - \xi} \cdot \rho_{\text{Therm}}^{m_\omega^+}(\omega_+) \otimes \rho_{\text{Therm}}^{m_\omega^-}(\omega_-), \quad (3.4)$$

carrying the fidelity

$$\mathcal{F}_{\text{Lnk}} = \frac{C_\sigma C_\omega}{C_\omega - (1 - \text{dc}_{\text{Lnk}}) C_\sigma} \left[S_+ \sigma_+^2 + S_- \sigma_-^2 - (1 - \text{dc}_{\text{Lnk}}) (W_+ \omega_+^2 + W_- \omega_-^2) \right]. \quad (3.5)$$

heralded with probability

$$\mathbb{P}_{\text{Lnk}} = (1 - \text{dc}_{\text{Lnk}})(1 - p_A)(1 - p_B) \left[\frac{1}{(1 - \sigma_+^2)(1 - \sigma_-^2)} - \frac{1 - \text{dc}_{\text{Lnk}}}{(1 - \omega_+^2)(1 - \omega_-^2)} \right]. \quad (3.6)$$

The thermal state $\rho_{\text{Therm}}^a(x)$ on mode a is defined as

$$\rho_{\text{Therm}}^a(x) = \frac{1}{1 + \bar{n}(x)} \sum_n \left(\frac{\bar{n}(x)}{1 + \bar{n}(x)} \right)^n |n\rangle \langle n|_a \quad \text{with:} \quad \bar{n}(x) = \frac{x^2}{1 - x^2}. \quad (3.7)$$

The factor ξ is defined as

$$\xi = (1 - \text{dc}_{\text{Lnk}}) \cdot \frac{(1 - \sigma_+^2)(1 - \sigma_-^2)}{(1 - \omega_+^2)(1 - \omega_-^2)}. \quad (3.8)$$

The quantities $\sigma_{\pm}$ (resp. $\omega_{\pm}$) are the singular values of the 2×2 matrix M_{σ} (resp. M_{ω}), and the modes $m_{\sigma}^+, m_{\sigma}^-$ (resp. $m_{\omega}^+, m_{\omega}^-$) are the linear combinations of m_A and m_B defined by the singular decomposition of M_{σ} (resp. M_{ω}). Both matrices characterize the optical set-up, and are derived in section D.1 as

$$M_{\sigma} = \frac{1}{\sqrt{2}} \begin{pmatrix} \sqrt{p_A} & \sqrt{p_A R_{\text{Lnk}}} \\ \sqrt{p_B} & -\sqrt{p_B R_{\text{Lnk}}} \end{pmatrix} \quad M_{\omega} = \frac{\sqrt{R_{\text{Lnk}}}}{\sqrt{2}} \begin{pmatrix} \sqrt{p_A} & \sqrt{p_A} \\ \sqrt{p_B} & -\sqrt{p_B} \end{pmatrix}. \quad (3.9)$$

The singular values $\sigma_{\pm}, \omega_{\pm}$ and the rotated modes $m_{\sigma}^+, m_{\sigma}^-, m_{\omega}^+, m_{\omega}^-$ are defined through

$$M_{\sigma} = U_{\sigma}^{\dagger} \cdot \begin{pmatrix} \sigma_+ & 0 \\ 0 & \sigma_- \end{pmatrix} \cdot V_{\sigma} \quad M_{\omega} = U_{\omega}^{\dagger} \cdot \begin{pmatrix} \omega_+ & 0 \\ 0 & \omega_- \end{pmatrix} \cdot V_{\omega}, \quad (3.10)$$

and

$$\begin{pmatrix} m_{\sigma}^+ \\ m_{\sigma}^- \end{pmatrix} = U_{\sigma} \begin{pmatrix} m_A \\ m_B \end{pmatrix} = \begin{pmatrix} \Sigma_A^+ & \Sigma_B^+ \\ \Sigma_A^- & \Sigma_B^- \end{pmatrix} \begin{pmatrix} m_A \\ m_B \end{pmatrix}, \quad (3.11)$$

$$\begin{pmatrix} m_{\omega}^+ \\ m_{\omega}^- \end{pmatrix} = U_{\omega} \begin{pmatrix} m_A \\ m_B \end{pmatrix} = \begin{pmatrix} \Omega_A^+ & \Omega_B^+ \\ \Omega_A^- & \Omega_B^- \end{pmatrix} \begin{pmatrix} m_A \\ m_B \end{pmatrix}.$$

From this SVD decomposition, we introduced the thermal envelopes

$$C_{\sigma} = (1 - \sigma_+^2)(1 - \sigma_-^2), \quad C_{\omega} = (1 - \omega_+^2)(1 - \omega_-^2), \quad (3.12)$$

so that $\xi = (1 - \text{dc}_{\text{Lnk}}) C_{\sigma}/C_{\omega}$, together with the Bell-pair overlaps onto the collective modes

$$S_{\pm} = \frac{(\Sigma_A^{\pm} + \Sigma_B^{\pm})^2}{2}, \quad W_{\pm} = \frac{(\Omega_A^{\pm} + \Omega_B^{\pm})^2}{2}. \quad (3.13)$$

A straightforward computer algebra program gives

$$\sigma_{\pm} = \frac{1}{2} \sqrt{(p_A + p_B)(1 + R_{\text{Lnk}}) \pm \sqrt{(p_A + p_B)^2(1 + R_{\text{Lnk}})^2 - 16 p_A p_B R_{\text{Lnk}}}} \quad (3.14)$$

$$\omega_+ = \sqrt{p_A R_{\text{Lnk}}} \quad \omega_- = \sqrt{p_B R_{\text{Lnk}}}.$$

We introduce

$$\Gamma_{\pm}(x, y) = \frac{1}{4} (p_A x + p_B y)(1 + R_{\text{Lnk}}) \pm \frac{1}{4} \sqrt{(p_A x + p_B y)^2(1 + R_{\text{Lnk}})^2 - 16 R_{\text{Lnk}} p_A x p_B y}, \quad (3.15)$$

so that $\sigma_{\pm}^2 = \Gamma_{\pm}(1, 1)$.

Symmetric case: The general case $p_A \neq p_B$ produces lengthy expressions for $\sigma_{\pm}$, $\omega_{\pm}$, U_{σ} , and U_{ω} . The repeater typically operates in a symmetric regime $p_A \approx p_B$, so the assumption $p_A = p_B = p$ is a reasonable working hypothesis. This yields

$$\begin{aligned} \sigma_+ &= \sqrt{p} & \sigma_- &= \sqrt{p R_{\text{Lnk}}} & m_{\sigma}^+ &= \frac{m_A + m_B}{\sqrt{2}} & m_{\sigma}^- &= \frac{m_A - m_B}{\sqrt{2}} \\ \omega_+ &= \sqrt{p R_{\text{Lnk}}} & \omega_- &= \sqrt{p R_{\text{Lnk}}} & m_{\omega}^+ &= \frac{m_A + m_B}{\sqrt{2}} & m_{\omega}^- &= \frac{-m_A + m_B}{\sqrt{2}} \end{aligned} \quad (3.16)$$

The heralding probability simplifies to

$$\mathbb{P}_{\text{Lnk}} = \frac{(1 - \text{dc}_{\text{Lnk}})(1 - p)}{(1 - R_{\text{Lnk}} p)^2} \cdot \left[p \cdot (1 - R_{\text{Lnk}}) + \text{dc}_{\text{Lnk}} \cdot (1 - p) \right], \quad (3.17)$$

and the shared state becomes

$$\begin{aligned} \rho_{\text{Lnk}} &= \frac{1 - p R_{\text{Lnk}}}{p(1 - R_{\text{Lnk}}) + \text{dc}_{\text{Lnk}}(1 - p)} \cdot \rho_{\text{Therm}}^{\frac{m_A + m_B}{\sqrt{2}}} \left(\sqrt{p} \right) \otimes \rho_{\text{Therm}}^{\frac{m_A - m_B}{\sqrt{2}}} \left(\sqrt{p R_{\text{Lnk}}} \right) \\ &\quad - \frac{(1 - \text{dc}_{\text{Lnk}})(1 - p)}{p(1 - R_{\text{Lnk}}) + \text{dc}_{\text{Lnk}}(1 - p)} \cdot \rho_{\text{Therm}}^{\frac{m_A + m_B}{\sqrt{2}}} \left(\sqrt{p R_{\text{Lnk}}} \right) \otimes \rho_{\text{Therm}}^{\frac{-m_A + m_B}{\sqrt{2}}} \left(\sqrt{p R_{\text{Lnk}}} \right). \end{aligned} \quad (3.18)$$

Physical interpretation and numerical results

At first glance, the expression of ρ_{Lnk} in equation (3.4) appears to be separable: it is written as a linear combination of tensor products of thermal states, with no obvious entanglement between the two sides. However, the thermal states do not live on the memory modes m_A and m_B themselves, but on the collective modes $m_{\sigma}^{\pm}$ and $m_{\omega}^{\pm}$ defined as linear combinations of m_A and m_B via M_{σ} and M_{ω} SVDs. The entanglement between Alice and Bob is therefore hidden in this collective basis: the two hubs do not share a product state on local modes, but a delocalized excitation distributed across $m_{\sigma}^{\pm}$ and $m_{\omega}^{\pm}$. In a symmetric ideal limit $p_A = p_B = p$, $\text{dc}_A = \text{dc}_B = 0$, $\eta_A = \eta_B = 1$, the rotated modes

reduce to $m_\sigma^+ = (m_A + m_B)/\sqrt{2}$ and $m_\sigma^- = (m_A - m_B)/\sqrt{2}$, and the state simplifies to

$$\rho_{\text{Lnk}} = (1-p) \sum_{k \geq 0} p^k |k+1\rangle \langle k+1|_{m_\sigma^+} \otimes |0\rangle \langle 0|_{m_\sigma^-}. \quad (3.19)$$

The entanglement is now manifest: the symmetric mode m_σ^+ carries a delocalized excitation between m_A and m_B , and in the perturbative limit $p \rightarrow 0$, one recovers $\rho_{\text{Lnk}} = |\Psi^+\rangle \langle \Psi^+|$.

Despite being introduced as a compact notation for the squared singular values of M_σ and M_ω , $\Gamma_\pm(x, y)$ admits a direct physical interpretation. The quantities $\Gamma_+(x, y)$ and $\Gamma_-(x, y)$ actually play the role of *effective pump parameters* for the two collective modes, representing the thermal occupation of the optical modes. The arguments x and y account for additional losses on Alice's and Bob's memory modes: they are set to 1 for the heralded state, and to R_A or R_B when computing the local detection probabilities, see equation (3.28) in the next section. Two limits clarify the role of this quantity. In the lossless limit $\eta_{\text{Lnk}} \rightarrow 1$, the pump concentrates on the symmetric collective mode m_σ^+

$$\Gamma_+(x, y) \xrightarrow{\eta_{\text{Lnk}} \rightarrow 1} \frac{x p_A + y p_B}{4}, \quad \Gamma_-(x, y) \xrightarrow{\eta_{\text{Lnk}} \rightarrow 1} 0, \quad (3.20)$$

recovering the Schmidt structure of a pure beam-splitter operation on two inputs. In the opposite limit $R_{\text{Lnk}} \rightarrow 1$, the beam-splitter is fully lossy and no mode mixing occurs: the two collective modes reduce to Alice's and Bob's individual modes, each carrying its own pump

$$\Gamma_+(1, 1) \xrightarrow{R_{\text{Lnk}} \rightarrow 1} \max(p_A, p_B), \quad \Gamma_-(1, 1) \xrightarrow{R_{\text{Lnk}} \rightarrow 1} \min(p_A, p_B). \quad (3.21)$$

For intermediate values of R_{Lnk} , the cross term $-16 R_{\text{Lnk}} p_A x p_B y$ under the square root reduces the gap between Γ_+ and Γ_- , reflecting the decoherence between the two collective modes as the beam-splitter losses increase. In this sense, $\Gamma_\pm$ naturally describes the joint Gaussian heralded state of Alice's and Bob's memories.

As a first numerical result, figure 3.2 shows the heralding probability $\mathbb{P}_{\text{Lnk}}$ as a function of the pump parameter p for different values of the elementary link length L (symmetric situation $p_A = p_B = p$). We first focus on figure 3.2a which displays the full theoretical range $p \in [0, 1]$. This plot highlights the competition between two effects: high pump parameter favors high photon emission and thus heralding clicks probability, but also produces multi-photon contributions that is evenly distributed across the detectors with the 50/50 beam-splitter, suppressing the single-click probability. This trade-off is clearly visible for short links $L \approx 1$ km: $\mathbb{P}_{\text{Lnk}}$ rises with p as single-photon emission becomes more likely, reaches a maximum, then decreases beyond $p \approx 0.5$ as multi-photon emission populates the two central station detectors. The maximum shifts to the right as the repeater length increases, longer links mean a lossier channel $\eta_{\text{Lnk}} = 10^{-\alpha L/20}$, suppressing multi-photon contributions and pushing back the threshold at which they start to degrade the heralding probability. For practical lengths $L \geq 100$ km, the maximum

is reached at pump values that are difficult to reach experimentally $p \geq 0.8$. Thus, in the operational regime, $\mathbb{P}_{\text{Lnk}}$ behaves as a monotonically increasing function of p . This is the regime shown in figure 3.2b, which zooms on $10^{-3} < p < 0.3$ and confirms the near-linear behavior of $\mathbb{P}_{\text{Lnk}}$ with p in log-log scale. This plot also provides quantitative orders of magnitude for practical implementations: typically $\mathbb{P}_{\text{Lnk}} \approx 10^{-4}$ for $L \approx 250$ km and $p \approx 0.05$, consistently with the rough estimate $\mathbb{P}_{\text{Lnk}} \sim p \eta_{\text{Lnk}} \approx 0.05 \times 0.003 \approx 10^{-4}$.

Figure 3.3 shows the elementary link fidelity $\mathcal{F}_{\text{Lnk}}$ as a function of the pump parameter for different values of the elementary length L (symmetric scenario $p_A = p_B = p$). To begin with, let us focus on figure 3.3a when no dark-count is considered. The behavior follows the analysis of section 1.1.2: the fidelity starts at $\mathcal{F}_{\text{Lnk}} = 1$ for $p \rightarrow 0$ and decreases as the pump parameter increases. As developed earlier, only the 1-photon sector contributes to the fidelity, with $\mathcal{F}_{\text{Lnk}}^{(1)} = 1$ and $\mathcal{F}_{\text{Lnk}}^{(n)} = 0$ for $n \neq 1$, thus raising the pump parameter feeds the multi-photon sectors, which dilutes this 1-photon contribution and lowers the fidelity. At fixed pump parameter the fidelity decreases with the repeater length, a behavior expected from the drop in transmission efficiency, and the $L \rightarrow 0$ and $L \rightarrow \infty$ bounds are shown for clarity. The plot also reports the local operation and classical communication (LOCC) bound $\mathcal{F}_{\text{LOCC}} = 0.5$, the maximum fidelity reachable through local operations and classical communication. Thus, reaching a fidelity $\mathcal{F}_{\text{Lnk}} < 0.5$ cannot witness any entanglement. Note that the converse does not hold: an entangled state can have $\mathcal{F}_{\text{Lnk}} < 0.5$.

We also provide the plot for a non-null dark-count, figure 3.3b shows this case for $\text{dc}_{\text{Lnk}} = 10^{-5}$. The result carries more physical structure and deserves closer attention. As p increases, the fidelity starts at 0, rises to a maximum or a plateau, then decreases. This shape defines three regimes along the p axis: the dark-count floor (1), the heralding window (2), and the multi-photon rolloff (3), described below.

- **Dark-count floor (1)** At low pump parameter, the heralding is governed by dark-counts: the pair-emission probability is small and an idler photon rarely reaches the central station, thus most of the heralding clicks are dark-counts rather than successful Bell measurements, which herald an empty state and do not contribute to entanglement. This regime is captured by the perturbative expansion of the fidelity derived in section 1.1.2. Keeping terms up to the 1-photon sector in equation (1.29) the fidelity reads

$$\mathcal{F}_{\text{Lnk}} = p \frac{\mathbb{P}_{\text{Lnk}}^{(1)}}{\mathbb{P}_{\text{Lnk}}^{(0)}} \mathcal{F}_{\text{Lnk}}^{(1)}, \quad (3.22)$$


since $\mathcal{F}_{\text{Lnk}}^{(0)} = 0$. In order of magnitude, the 0-photon sector is heralded only by dark-counts, $\mathbb{P}_{\text{Lnk}}^{(0)} \approx \text{dc}_{\text{Lnk}}$, while the 1-photon sector is heralded by the transmission of a single photon to the central station, $\mathbb{P}_{\text{Lnk}}^{(1)} \approx \eta_{\text{Lnk}}$. The fidelity is thus set by the signal-

to-noise ratio $\text{SNR} = p \eta_{\text{Lnk}} / \text{dc}_{\text{Lnk}}$ between real photon detection and dark-counts


$$\mathcal{F}_{\text{Lnk}} \xrightarrow{\text{SNR} \ll 1} 0. \quad (3.23)$$

The two contributions balance for $\text{SNR} = 1$, i.e. at the pump threshold $p_{1-2} = \text{dc}_{\text{Lnk}} / \eta_{\text{Lnk}}$, defining the frontier between the region (1) and region (2) with the set of points $\{(p_{1-2}, \mathcal{F}_{\text{Lnk}}(p_{1-2}))\}$. This threshold depends on the repeater length through the transmission efficiency, $p_{1-2}(L) = \text{dc}_{\text{Lnk}} / \eta_{\text{Lnk}} = \text{dc}_{\text{Lnk}} 10^{\alpha L / 10}$, and delimits the border between regimes (1) and (2) on the plot. We report it on figure 3.3b by evaluating $\mathcal{F}_{\text{Lnk}}(p_{1-2}(L))$ for varying L . At $p \approx p_{1-2}$, SPDC pairs and dark-counts contribute equally to the heralding, yielding an equal-weight mixture of the Bell pair and the empty sector, namely $\mathcal{F}_{\text{Lnk}}(p_{\text{th}}) = 0.5$. The threshold therefore tracks the LOCC bound over most of its range, however it shifts to higher pump parameters as the repeater length grows: longer links lead to lower single-photon transmission and larger p is needed to overcome the dark-count rate. Beyond a hundred kilometers losses are too high and the 1-photon sector contribution cannot match the dark-count rate without involving multi-photon contributions, and the threshold fidelity $\mathcal{F}_{\text{Lnk}}(p_{1-2})$ drops below 0.5 toward 0.

- **Heralding window (2)** Beyond the dark-count-driven regime, raising p above p_{1-2} makes the 1-photon sector contribution predominant over the dark-counts. The fidelity then rises with p until it reaches a maximum, where the balance between the growing 1-photon contribution and the multi-photon dilution is optimal for the entanglement quality. We thus define the boundary between (2) and (3) by the points $(p_{2-3}, \mathcal{F}_{\text{Lnk}}(p_{2-3}))$, where p_{2-3} is the pump parameter at which the 1-photon sector and the multi-photon contributions balance, $\tilde{\mathcal{F}}_{\text{Lnk}}^{(1)}(p_{2/3}) = 0.5$. This window, and more precisely its maximum, defines the actual working regime of the repeater that should be targeted with respect to the fidelity metric. For long quantum repeaters, however, this window vanishes: the plot shows that for $L > 500$ km the maximum fidelity never exceeds the LOCC bound. For such repeater length, the transmission is too lossy and the 1-photon contribution cannot overcome the dark-count rate without multi-photon dilution.
- **Multi-photon rolloff (3)** Once the dark-count contribution becomes negligible, the system gradually recovers a behavior similar to the dark-count-free case. Thus, pushing p further degrades the fidelity through the growing weight of the multi-photon sectors $i \geq 2$, which dilute the Bell pair with higher-order SPDC contributions and drives the fidelity to 0 in the theoretical limit $p \rightarrow 1$.




(a) Full theoretical range $p \in [0, 1]$.




(b) Operational range $10^{-3} < p < 0.3$ in log-log scale, for L between 0 and 500 km.

Figure 3.2: Heralding probability $\mathbb{P}_{\text{Lnk}}$ of the elementary link as a function of the pump parameter p , for several link lengths L (color), in the symmetric configuration $p_A = p_B = p$. (a): the heralding probability rises with p , reaches a maximum, then decreases as multi-photon emission populates both central-station detectors and suppresses the single-click pattern. The maximum shifts to the right with L as lossy transmission suppress multi-photon contributions and push the pump threshold at which multi-photons degrade $\mathbb{P}_{\text{Lnk}}$. (b): in this regime $\mathbb{P}_{\text{Lnk}}$ is a near-linear increasing function of p , with orders of magnitude $\mathbb{P}_{\text{Lnk}} \approx 10^{-4}$ for $L \approx 250$ km and $p \approx 0.05$.



(a) Ideal regime, $\text{dc}_{\text{Lnk}} = 0$.



(b) Realistic regime with dark-counts, $\text{dc}_{\text{Lnk}} = 10^{-5}$.

Figure 3.3: Elementary link fidelity $\mathcal{F}_{\text{Lnk}}$ as a function of the pump parameter p , for several link lengths L (color), in the symmetric configuration $p_A = p_B = p$.

(a): the fidelity is a monotonically decreasing function of p , starting from unity at $p \rightarrow 0$ and dropping as multi-photon contributions contaminate the heralded state.

(b): three regimes emerge along the p axis (see main text): the dark-count floor (1), the heralding window (2), and the multi-photon rolloff (3). Circle dots mark the (1)–(2) boundary, evaluated as $\mathcal{F}_{\text{Lnk}}(p_{1-2}(L))$ at the dark-count threshold $p_{1-2} \eta_{\text{Lnk}} = \text{dc}_{\text{Lnk}}$. Square dots mark the (2)–(3) boundary, evaluated as $\tilde{\mathcal{F}}_{\text{Lnk}}^{(1)} = 0.5$ where the 1-photon and multi-photon contributions balance each other.

3.1.2 . Local detection probabilities

The exact expression of ρ_{Lnk} derived in equation (3.4) gives access to all observable quantities at Alice's and Bob's nodes, in particular the local detection statistics. We consider the set-up of figure 3.4, in which two non-resolving photodetectors are placed on the memory modes m_A and m_B at Alice's and Bob's hubs. Each detector is characterized by its efficiency and dark-count rate, namely (η_A, dc_A) on Alice's side and (η_B, dc_B) on Bob's side, with associated losses $R_A = 1 - \eta_A$, $R_B = 1 - \eta_B$ and complements $\kappa_A = 1 - \text{dc}_A$, $\kappa_B = 1 - \text{dc}_B$. We denote $\mathbb{P}_{ij}$ the joint detection probability with $i, j \in \{0, 1\}$ standing for click (1) or no-click (0) at Alice's and Bob's detectors respectively.

The derivation of the joint detection probabilities follows the same structure as the exact computation of ρ_{Lnk} . The beam-splitter at the central station and the photon losses are encoded in a non-unitary mixing matrix, defining the collective memory modes m_σ and m_ω over which the excitations are delocalized. The difference lies in the content of this matrix: it now also encodes the losses at Alice's and Bob's hubs, with $\eta_A = \eta_{\text{Det}}^A \cdot \eta_{\text{Mem}}^A$ collecting both detection and memory efficiencies on Alice's side, and similarly for Bob. Since the trace is now global, framing in the rotated mode basis yields (see appendix A.2)

$$|\psi\rangle = \exp \left[\begin{pmatrix} c_\sigma^+ & c_\sigma^- \end{pmatrix} \begin{pmatrix} \sigma_- & 0 \\ 0 & \sigma_+ \end{pmatrix} \begin{pmatrix} m_\sigma^+ \\ m_\sigma^- \end{pmatrix} \right] |0\rangle \quad \longrightarrow \quad \text{Tr}[\rho] = \frac{1}{(1 - \sigma_+^2)(1 - \sigma_-^2)}, \quad (3.24)$$

with $\sigma^\pm$ the singular values of the matrix M_σ . The function $\Gamma_\pm(x, y)$ introduced earlier naturally reappears here: its arguments x and y now carry the asymmetric losses R_A and R_B , encoding how the SVD structure of the mode-mixing and setting the thermal population of the two collective modes.

From appendix A.1, the local POVM elements read

$$\Pi_0^a(\eta, \text{dc}) = \kappa \cdot R^{a^\dagger a} \quad (3.25)$$

$$\Pi_1^a(\eta, \text{dc}) = \mathbb{1} - \kappa \cdot R^{a^\dagger a}, \quad (3.26)$$

which, in our context, leads to the joint probabilities

$$\mathbb{P}_{ij} = \text{Tr}_{\text{tot}} [\Pi_i^{m_A}(\eta_A, \text{dc}_A) \otimes \Pi_i^{m_B}(\eta_B, \text{dc}_B) \cdot \rho_{\text{Lnk}}]. \quad (3.27)$$

Expanding the POVMs and applying the heralding constraint of the central station, this trace is reduced to a sum of expectation values of operators of the form $\prod_k R_k^{\alpha_k^\dagger \alpha_k}$ acting on the post-beam-splitter SPDC state, each of which can be computed via the SVD-based trace identity exposed earlier. Anticipating a general result, the local probabilities are given by setting $\alpha = \beta = 0$ in the displaced local probabilities $\mathbb{P}_{\text{Loc}}^{ij}(\alpha, \beta)$ derived in

appendix D.3:

$$\begin{aligned}
\frac{\mathbb{P}_{\text{Lnk}}}{(1-p_A)(1-p_B)} \mathbb{P}_{\text{Loc}}^{00} &= \frac{\kappa_A \kappa_B \kappa_{\text{Lnk}}}{(1 - \Gamma_+(R_A, R_B))(1 - \Gamma_-(R_A, R_B))} - \frac{\kappa_A \kappa_B \kappa_{\text{Lnk}}^2}{(1 - p_A R_A R_{\text{Lnk}})(1 - p_B R_B R_{\text{Lnk}})} \\
\frac{\mathbb{P}_{\text{Lnk}}}{(1-p_A)(1-p_B)} \mathbb{P}_{\text{Loc}}^{01} &= \frac{\kappa_A \kappa_{\text{Lnk}}}{(1 - \Gamma_+(R_A, 1))(1 - \Gamma_-(R_A, 1))} - \frac{\kappa_A \kappa_{\text{Lnk}}^2}{(1 - p_A R_A R_{\text{Lnk}})(1 - p_B R_B R_{\text{Lnk}})} \\
&\quad - \frac{\kappa_A \kappa_B \kappa_{\text{Lnk}}}{(1 - \Gamma_+(R_A, R_B))(1 - \Gamma_-(R_A, R_B))} + \frac{\kappa_A \kappa_B \kappa_{\text{Lnk}}^2}{(1 - p_A R_A R_{\text{Lnk}})(1 - p_B R_B R_{\text{Lnk}})} \\
\frac{\mathbb{P}_{\text{Lnk}}}{(1-p_A)(1-p_B)} \mathbb{P}^{10} &= \frac{\kappa_B \kappa_{\text{Lnk}}}{(1 - \Gamma_+(1, R_B))(1 - \Gamma_-(1, R_B))} - \frac{\kappa_B \kappa_{\text{Lnk}}^2}{(1 - p_A R_{\text{Lnk}})(1 - p_B R_B R_{\text{Lnk}})} \\
&\quad - \frac{\kappa_A \kappa_B \kappa_{\text{Lnk}}}{(1 - \Gamma_+(R_A, R_B))(1 - \Gamma_-(R_A, R_B))} + \frac{\kappa_A \kappa_B \kappa_{\text{Lnk}}^2}{(1 - p_A R_A R_{\text{Lnk}})(1 - p_B R_B R_{\text{Lnk}})} \\
\frac{\mathbb{P}_{\text{Lnk}}}{(1-p_A)(1-p_B)} \mathbb{P}^{11} &= \frac{\kappa_{\text{Lnk}}}{(1 - \Gamma_+(1, 1))(1 - \Gamma_-(1, 1))} - \frac{\kappa_A \kappa_{\text{Lnk}}}{(1 - \Gamma_+(R_A, 1))(1 - \Gamma_-(R_A, 1))} \\
&\quad - \frac{\kappa_B \kappa_{\text{Lnk}}}{(1 - \Gamma_+(1, R_B))(1 - \Gamma_-(1, R_B))} - \frac{\kappa_{\text{Lnk}}^2}{(1 - p_A R_{\text{Lnk}})(1 - p_B R_B R_{\text{Lnk}})} \\
&\quad + \frac{\kappa_A \kappa_{\text{Lnk}}^2}{(1 - p_A R_A R_{\text{Lnk}})(1 - p_B R_B R_{\text{Lnk}})} + \frac{\kappa_B \kappa_{\text{Lnk}}^2}{(1 - p_A R_{\text{Lnk}})(1 - p_B R_B R_{\text{Lnk}})} \\
&\quad + \frac{\kappa_A \kappa_B \kappa_{\text{Lnk}}}{(1 - \Gamma_+(R_A, R_B))(1 - \Gamma_-(R_A, R_B))} - \frac{\kappa_A \kappa_B \kappa_{\text{Lnk}}^2}{(1 - p_A R_A R_{\text{Lnk}})(1 - p_B R_B R_{\text{Lnk}})}
\end{aligned} \tag{3.28}$$

where $\mathbb{P}_{\text{Lnk}}$ is the heralding probability from the last computation given by

$$\mathbb{P}_{\text{Lnk}} = (1 - \text{dc}_{\text{Lnk}})(1 - p_A)(1 - p_B) \left[\frac{1}{(1 - \sigma_+^2)(1 - \sigma_-^2)} - \frac{1 - \text{dc}_{\text{Lnk}}}{(1 - \omega_+^2)(1 - \omega_-^2)} \right]. \tag{3.29}$$

and $\Gamma_{\pm}$ is the auxiliary function introduced in the previous section

$$\Gamma_{\pm}(x, y) = \frac{1}{4}(p_A x + p_B y)(1 + R_{\text{Lnk}}) \pm \frac{1}{4} \sqrt{(p_A x + p_B y)^2 (1 + R_{\text{Lnk}})^2 - 16 R_{\text{Lnk}} p_A x p_B y}. \tag{3.30}$$

These four expressions give the local detection statistics entirely in terms of the physical parameters $p_A, p_B, \eta_A, \eta_B, \eta_{\text{Lnk}}, \text{dc}_A, \text{dc}_B, \text{dc}_{\text{Lnk}}$. These probabilities will be used in subsection 3.2.1 for benchmarking the model against the data obtained experimentally by the De Riedmatten group.

As for the derivation of the exact state ρ_{Lnk} , the expressions above are general but analytically heavy, they are suitable for numerical evaluation rather than for analytical study. However, one must be aware of potential numerical errors: each $\mathbb{P}_{ij}$ is obtained as a difference of terms of comparable magnitude, a direct consequence of the POVM




Figure 3.4: Optical setup of the joint local probability measurement on an elementary link. Local detection at each node is described by the efficiencies $\eta_A = \eta_{\text{Mem}}^A \eta_{\text{Det}}^A$ for Alice and $\eta_B = \eta_{\text{Mem}}^B \eta_{\text{Det}}^B$ for Bob. In the symmetric scenario, we set $p_A = p_B = p$, $\eta_A = \eta_B = \eta_{AB}$, and likewise for the dark-count rate.

expansion. For instance, the joint probability $\mathbb{P}_{01}$ stems from

$$\Pi_{01} \Pi_{\text{Lnk}} = \Pi_0^{m_A} \Pi_1^{m_B} \Pi_0^{c_A} \Pi_0^{c_B} \quad (3.31)$$

$$= \kappa_A \kappa_{\text{Lnk}} \cdot R_A^{m_A^\dagger m_A} R_{\text{Lnk}}^{c_A^\dagger c_A} \quad (3.32)$$

$$- \kappa_A \kappa_{\text{Lnk}}^2 \cdot R_{\text{Lnk}}^{c_B^\dagger c_B} R_A^{m_A^\dagger m_A} R_{\text{Lnk}}^{c_A^\dagger c_A} \quad (3.33)$$

$$- \kappa_B \kappa_A \kappa_{\text{Lnk}} \cdot R_B^{m_B^\dagger m_B} R_A^{m_A^\dagger m_A} R_{\text{Lnk}}^{c_A^\dagger c_A} \quad (3.34)$$

$$+ \kappa_B \kappa_A \kappa_{\text{Lnk}}^2 \cdot R_A^{m_A^\dagger m_A} R_B^{m_B^\dagger m_B} R_{\text{Lnk}}^{c_B^\dagger c_B} R_A^{m_A^\dagger m_A} R_{\text{Lnk}}^{c_A^\dagger c_A} \quad (3.35)$$

which is structurally a sum of marginal contributions. In a situation of low joint probabilities, the marginal contributions may be close for cancellation, and a naive double-precision evaluation can miss realistic values.

3.1.3 . Coherence Evaluation

While characterizing the population of the elementary link state, the joint local detection statistics derived above do not capture the coherence of the delocalized excitations. A standard way to access this coherence is via single-photon interferometry [26]: the memory modes are mixed on a balanced beam-splitter with a tunable relative phase ϕ , and the single-click probability $\mathbb{P}^\vee(\phi)$ is measured as a function of ϕ . Under the Bell pair mixture assumption, $|\Psi^+\rangle$ and $|\Psi^-\rangle$, the resulting fringes contrast is directly related to the

off-diagonal coherence of the delocalized excitation. Concretely, for a state of the form

$$\rho = \frac{1+c}{2} |\Psi^+\rangle \langle \Psi^+| + \frac{1-c}{2} |\Psi^-\rangle \langle \Psi^-| = \frac{1}{2} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & c & 0 \\ 0 & c & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (3.36)$$

with $c \in [-1, 1]$, the single-click probability reads

$$\mathbb{P}^{\mathcal{V}}(\phi) = \frac{1 + c \cos \phi}{2}, \quad (3.37)$$

so that the visibility directly gives $|c|$. In our context, this model is relevant only in the low-pump regime, where Alice's and Bob's memories share a single delocalized excitation. Beyond this regime the derivation gives an exact visibility valid for all pump powers, providing to the experimentalist a fined model accounting for all i -photon contributions. It also offers a way to validate the model by comparing the predicted visibility with experimental data.

In the setup depicted in figure 3.5, the two memory modes m_A and m_B are recombined on a beam-splitter, and the single-click statistics are measured as a function of a relative phase ϕ . Thus, the visibility measurement adds a second beam-splitter, this time mixing the memory modes m_A and m_B , together with a relative phase ϕ . This stage is encoded by a new mixing matrix $\tilde{M}$, bringing an additional rotation on the memory modes. The probability is then obtained as a global trace over all modes, which becomes diagonal in the rotated basis defined by the SVD of $M_{\text{tot}} = \tilde{M}^\dagger M_{\sigma/\omega}$.

As algebraically heavy, we here provide the expression in the symmetric case $p_A = p_B = p$, $R_A = R_B = R_{AB}$, and without dark counts $\text{dc}_A = \text{dc}_B = \text{dc}_{\text{Lnk}} = 0$:

$$\mathcal{V} = \frac{(1 - R_{AB})(1 - pR_{\text{Lnk}})}{(1 - pR_{AB})(1 - pR_{AB}R_{\text{Lnk}})}. \quad (3.38)$$

Standard limits are recovered, for instance the low-pump expansion reads

$$\mathcal{V} = \eta_{AB} + \eta_{AB} (1 - (2 - \eta_{\text{Lnk}})\eta_{AB}) p + \text{o}(p), \quad (3.39)$$

recovering the lossy Bell-pair visibility with the first-order correction from the 2-photon sector, and full lossy memories $R_{AB} \rightarrow 1$ destroy the coherence.

Figure 3.6 reports the numerical implementation of the interferometry fringes $\mathbb{P}^{\mathcal{V}}$ and the resulting visibility $\mathcal{V}$ derived above. To begin with, figure 3.6a displays the fringes $\mathbb{P}^{\mathcal{V}}(\phi)$ for different pump values in the symmetric configuration. At low pump, the fringe closely follows the mixture of Bell-pair reference with a sinusoidal profile. At high pump, the contrast rises and the fringe shape turns into a narrow peak around $\phi = 0$. This is the expected behavior of multi-photon interference: the n -photon sectors contribute with a phase $e^{in\phi}$ and the resulting coherent sum narrows the fringe, similarly to a N -

slit Young experiment. However, such a peaked profile should not be understood as an improvement of the entanglement quality of the heralded state: the interpretation of visibility as a coherence witness only stands on the Bell-pair mixture assumption, which breaks when multi-photon sectors contribution are non-negligible. Additionally, the pump range $p \in [0, 1]$ displayed here is illustrative as values $p \geq 10^{-1}$ for SPDC sources are difficult to reach in practice. Figure 3.6b restricts the analysis to the experimentally relevant low-pump regime. It shows the visibility as a function of the pump parameter for different central-station dark-count values, for short length $L = 100$ m, aiming to match a laboratory interferometric setup¹. In this region, the visibility drops to zero as the dark-count rises, this trend matches the perturbative analysis of section 1.1: at low pump, the 1-photon sector dominates and the elementary-link state reduces to the Bell-pair mixture

$$\rho_{\text{Lnk}}^{(1)} = \frac{1}{1 + \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}})}{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}})}} |\Psi^+\rangle \langle \Psi^+| + \frac{1}{1 + \frac{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}})}{\kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}})}} |\Psi^-\rangle \langle \Psi^-|, \quad (3.40)$$

where dc_{Lnk} balances the weights of $|\Psi^+\rangle$ and $|\Psi^-\rangle$ and deteriorates the coherence of the heralded state.

3.2 .Experimental fit

In parallel to this theoretical work, the experimental group of H. de Riedmatten at ICFO (Barcelona) implemented an experimental DLCZ-type elementary link. At the time of writing, this implementation features two SPDC sources heralded by a central detection, with on-demand solid-state quantum memories separated by a few meters [35]. The group has characterized the two SPDC sources, measured the local detection probabilities $\mathbb{P}_{\text{Loc}}^{ij}$ heralded by a central click, and performed visibility measurements on the heralded state. This section confronts our exact model ρ_{Lnk} with two experimental data: an unpublished set of local probability measurements, and the published set of visibility measurements from [35].

3.2.1 .Local probabilities

We fit the experimental local probabilities statistics $\mathbb{P}_{\text{Loc}}^{ij, \text{Exp}}$ with the ones from equation (3.28). The fit is performed on the four probabilities $\mathbb{P}_{\text{Loc}}^{ij}$ as a function of the pump parameters p_A and p_B . It returns numerical values for the link parameters (efficiencies, dark counts), benchmarking the consistency of the model with experimental data. The remainder of the section presents the experimental setup, fitting method, and numerical results.

In this fit, we use unpublished data of ICFO from a previous version of the elementary link. As high pump parameter is hard to reach in practice, the measurements have been taken in the low-pump regime, where the heralded state is close to a Bell pair. The data

¹In the next section we fit this model with data from [35], data taken at short distance





Figure 3.5: Optical setup of the single-photon interference measurement. A phase ϕ is applied to the mode m_A , and the two memory modes are recombined on a 50-50 beam-splitter followed by a single-click measurement Π_V . The single-click statistics as a function of ϕ , namely $\mathbb{P}^V$, results in interference fringes where the contrast is directly related to the Bell pair coherence.



(a) Single-photon interferometry fringes $\mathbb{P}^\mathcal{V}(\phi)$ as a function of the relative phase ϕ , for different values of the pump parameter p .

(b) Visibility $\mathcal{V}$ extracted from the fringes as a function of the pump parameter p , for central-station dark-count values dc_{Lnk} .

Figure 3.6: Numerical evaluation of the interferometric fringes $\mathbb{P}^\mathcal{V}(\phi)$ and of the resulting visibility $\mathcal{V}$ derived above. Practical parameters are taken from table 1, in the symmetric configuration $p_A = p_B = p$ and $\eta_A = \eta_B = \eta_{AB}$. In both panels, the dash-dotted line marks the Bell-pair scenario from the perturbative study, setting asymptotic regime constraint by the interferometer efficiency.

Panel (a) shows how the fringe profile departs from a pure cosine as the pump rises, with multi-photon contributions sharpening the peak fringe.

Panel (b) shows the visibility for experimentally relevant low-pump regime, at link length $L = 100$ m, and quantifies how visibility collapse with the central-station dark-count.

set consists of five acquisition points spanning different pump powers, each built from approximately 10^5 heralding events from a pump configuration $p_k = (p_A^{(k)}, p_B^{(k)})$. For each point, the experimental team measured the pump powers and the local click counts at Alice's and Bob's side.

The theoretical model involves six fitting parameters: the efficiencies η_{Lnk} , η_A , η_B and the dark counts dc_{Lnk} , dc_A , dc_B . Let ξ denote the set of fitting parameters. We define a per-point p_k cost as the weighted root-mean-square deviation between theoretical and experimental probabilities:

$$\mathcal{L}_\xi(p_k) = \sqrt{\sum_{i,j} \alpha_{ij} \cdot \frac{\left(\mathbb{P}_{\text{Loc}}^{ij,\text{Exp}}(p_k) - \mathbb{P}_{\text{Loc}}^{ij}(p_k)\right)^2}{\left(\sum_{i,j} \mathbb{P}_{\text{Loc}}^{ij,\text{Exp}}(p_k)\right)^2}}, \quad (3.41)$$

where the coefficients $\{\alpha_{ij}\}$ weight the contribution of each event, discussed below. The total cost function combines the contributions of the five acquisition points $\{p_k\}$:

$$\text{MSE}_\xi = \sqrt{\sum_k \mathcal{L}_\xi^2(p_k)}. \quad (3.42)$$

We minimize MSE_ξ with a least squared routine from SciPy. To mitigate the dependence

on initialization which may lead to a local minima, we use a Monte Carlo method by repeating the optimization over N_{Fit} randomized initial parameters ξ_{ini} . The optimization is repeated over $N_{\text{Fit}} = 200$ random initial parameters drawn over a wide range, the distribution of the N solutions is then used to estimate optimal values and confidence intervals. The weights $\{\alpha_{ij}\}$ act as a fitting hyperparameters: in the low-pump regime, $\mathbb{P}_{\text{Loc}}^{11}$ is very small and its measurement is dominated by shot noise, so we set $\alpha_{11} = 0.1$ to reduce its influence, while keeping $\alpha_{ij} = 1$ for the other events.

Data preprocessing The model provides the local probabilities $\mathbb{P}_{\text{Loc}}^{ij}$ as a function of p_k , while the raw experimental data is expressed in pump powers and click counts. Two correspondences are therefore needed: one between the pump power P_{mW} in milliwatts and the pump parameter p_A and p_B , and one between the click counts N_{ij} and the local probabilities $\mathbb{P}_{\text{Loc}}^{ij, \text{Exp}}$.

- **Pump parameter calibration** The unheralded second-order autocorrelation of a single SPDC arm follows the statistics

$$g^{(2)} = 1 + \frac{1}{\langle n \rangle}, \quad (3.43)$$

with $\langle n \rangle = p$ the mean photon-pair number [128]. Assuming a linear relation between the pump power and the pump parameter $P_{\text{mW}} = a \cdot p$, valid for low pump parameter, we have

$$g^{(2)} = 1 + \frac{a}{P_{\text{mW}}}. \quad (3.44)$$

Fitting this expression against the experimental $g^{(2)}$ and pump-power measurements for each source yields

$$\begin{aligned} \text{Alice:} \quad & a = 17.8, \quad R^2 = 0.99, \\ \text{Bob:} \quad & a = 13.84, \quad R^2 = 0.95, \end{aligned} \quad (3.45)$$

with R the correlation coefficient, establishing the calibration between pump power and pump parameter for each SPDC source.

- **Local probabilities** Let N_{ij} be the click count for event (i, j) and $N_{\text{tot}} = \sum_{ij} N_{ij}$ the total number of heralded events. The empirical local probability is

$$\mathbb{P}_{\text{Loc}}^{ij, \text{Exp}} = \frac{N_{ij}}{N_{\text{tot}}}. \quad (3.46)$$

Each of the five acquisition points thus provides four experimental probabilities $\mathbb{P}_{\text{Loc}}^{ij, \text{Exp}}$ together with their associated pump parameters p_A and p_B .

The Fit As a preliminary study, we performed the procedure described above over $N_{\text{Fit}} = 200$ random initializations, considering a symmetric model where Alice's and Bob's hubs share the same parameters: $\eta_A = \eta_B = \eta_{AB}$ and $dc_A = dc_B = dc_{AB}$, with $p_A = p_B = p$. Starting from an initial $\text{MSE}_{\text{ini}} = 3 \pm 2$, nearly all fits converge to a common minimum $\text{MSE}_{\text{fit}} \approx 1.4 \cdot 10^{-2} \pm 2 \cdot 10^{-3}$, indicating a well-defined optimum $\theta^{\text{Opt}} = (\eta_{\text{Lnk}}^{\text{Opt}}, dc_{\text{Lnk}}^{\text{Opt}}, \eta_{AB}^{\text{Opt}}, dc_{AB}^{\text{Opt}})$. This symmetric model captures the correct order of magnitude, as reported in table 3.1, validating the modeling approach, but it cannot resolve the imbalance between $\mathbb{P}_{\text{Loc}}^{01}$ and $\mathbb{P}_{\text{Loc}}^{10}$. To address this, we relax the symmetry constraint and let $\eta_A \neq \eta_B$ and $dc_A \neq dc_B$. The plots below reports such a fit, with $N_{\text{Fit}} = 200$ random initializations, the same cost function and the same optimizer. figure 3.7 shows the fitted parameters across the N_{Fit} runs: all fits converge to a common MSE minimum, from $\text{MSE}_{\text{ini}} \approx 3 \pm 1$ to $\text{MSE}_{\text{fit}} \approx 0.033 \pm 0.005$, indicating an optimum θ^{Opt} summarized in table 3.1. We plot the fitted local probabilities in figure 3.8, with nearly all experimental data points captured within the 2σ envelope.

- **Efficiency.** The local efficiencies converge to $\eta_A \approx 0.138 \pm 0.002$ and $\eta_B \approx 0.045 \pm 0.001$, while the central-station efficiency converges to $\eta_{\text{Lnk}} \approx 0.86 \pm 0.05$. The local efficiencies are tightly constrained, with a relative dispersion $\sigma/\mu \approx 1.5\%$, whereas the central-station efficiency is markedly more volatile, with $\sigma/\mu \approx 5\%$, indicating a lower sensitivity of the fit to η_{Lnk} . This is consistent with the low-pump regime: the heralded state is close to the pure Bell pair $\rho_{\text{Lnk}} \approx |\Psi^+\rangle \langle \Psi^+|$, so local probabilities are weakly affected by η_{Lnk} .
- **Dark-counts.** A logarithmic parametrization was used for these parameters during the fit, their values span several orders of magnitude and the log-space fit brought better convergence. The fitted values $\log_{10}(dc_{\text{Lnk}}) \approx -3.4 \pm 0.5$, $\log_{10}(dc_A) \approx -3.8 \pm 0.5$ and $\log_{10}(dc_B) \approx -4.0 \pm 0.4$ correspond to dark-count probabilities of order 10^{-4} , with a relative dispersion on the exponent of $\sigma/\mu \approx 13\%$. The dark-counts are therefore less tightly constrained than the efficiencies and have a weaker impact on the cost function, as they are mainly driven by the coincidence local probability $\mathbb{P}_{\text{Loc}}^{11}$, whose contribution is marginal, typically $\mathbb{P}_{\text{Loc}}^{11}/\mathbb{P}_{\text{Loc}}^{01} \approx 1\%$.

3.2.2 .Single-photon interferometry

As an additional validation of the model, we fit the visibility data from [35], reporting a complete elementary link with on-demand spin-wave memories. In this experiment the authors witness the generation and storage of a coherent Bell state via single-photon interferometry statistics. We thus use these statistics as dataset and confront these fringes with the model derived in section 3.1.3, aiming to recover the parameters reported in [35].

The authors acquired the statistics at a single pump value in the low-pump regime, where multi-photon contributions to the fringe shape are negligible. Thus the robustness




Figure 3.7: Initial parameters (black crosses) and fitted parameters (orange crosses) across the $N_{\text{Fit}} = 200$ runs. All fits converge to a common optimum θ^{Opt} , demonstrating the robustness of the procedure to the randomization of the initial parameters.




Figure 3.8: Fitted local probabilities $\mathbb{P}_{\text{Loc}}^{ij}$ as a function of the pump parameter p . Most of the experimental data points lie within the 2σ envelope of the fitted model.

	η_{Lnk}	η_{A}	η_{B}	dc_{Lnk}	dc_{A}	dc_{B}
Symmetric	0.8 ± 0.15	0.062 ± 0.003		$2.0 \cdot 10^{-7} \pm 0.8$	$3.2 \cdot 10^{-4} \pm 0.4$	
Asymmetric	0.86 ± 0.04	0.138 ± 0.002	0.045 ± 0.001	$3.6 \cdot 10^{-4} \pm 0.5$	$1.8 \cdot 10^{-4} \pm 0.5$	$1.1 \cdot 10^{-4} \pm 0.4$

Table 3.1: Fitted parameters and their uncertainties for the symmetric and asymmetric models.

	η_{Lnk}	η_{A}	η_{B}	dc_{Lnk}	dc_{A}	dc_{B}	MSE
Fitted Parameters	0.099 ± 0.050	0.13 ± 0.067	0.12 ± 0.066	$6.2 \cdot 10^{-3} \pm 0.3$	$9.0 \cdot 10^{-3} \pm 0.3$	$9.4 \cdot 10^{-3} \pm 0.3$	$1.41 \pm 2 \cdot 10^{-6}$
Parameter [35]	0.20	$4.4 \cdot 10^{-3}$	$6.1 \cdot 10^{-3}$	n.p.	n.p.	n.p.	—

Table 3.2: Fitted parameters and their uncertainties for the visibility model.

of the model against multi-photon effects cannot be assessed here, and the pump parameter is treated as a fit variable rather than a known input. Concretely, the fitting parameters are $\xi = (\phi, p, \eta_{\text{Lnk}}, \text{dc}_{\text{Lnk}}, \eta_A, \text{dc}_A, \eta_B, \text{dc}_B)$, where ϕ is a residual phase offset of the interferometer. The data actually report the coincidences between a single-click at the central-station (on c_A or c_B) and at the interferometer (on m_A or m_B), yielding four configurations $(m_A, m_B) \times (c_A, c_B)$. The central-station click-side fixes the sign of the heralded Bell pair $|\Psi^\pm\rangle$, adding a phase 0 or π to the fringe. The memory click-side similarly selects one of the two modes of the interferometer, bringing an additional phase 0 or π . Combining the two effects, the configurations (m_A, c_A) and (m_B, c_B) share a phase 0, while (m_A, c_B) and (m_B, c_A) share a phase π .

The four fringes are fitted simultaneously following the procedure described previously: $N_{\text{Fit}} = 200$ random initializations, with the same optimizer, and statistics aggregated over the N_{Fit} to build an optimal set of parameters ξ^{Opt} with the associated fluctuation σ . The numerical results are summarized in table 3.2, together with the values reported by the authors of [35]. Figure 3.9 and figure 3.10 respectively report the parameter distributions over the N_{Fit} runs and the fitted fringes with the experimental data. As shown in figure 3.9, the cost function converges to a common minimum across all fits, decreasing from $\text{MSE}_{\text{ini}} \approx 5 \cdot 10^{-4} \pm 8 \cdot 10^{-4}$ to $\text{MSE}_{\text{fit}} \approx 5 \cdot 10^{-9} \pm 10^{-14}$. This suggest the optimization reliably converges to well defined set of optimal parameters, ξ^{Opt} . Additionally the corresponding fringes match the experimental data within the 1σ envelope, as shown in figure 3.10. However, the fitted parameters are less stable than for the local probabilities: the optimized efficiencies, for instance, show a relative fluctuation $\sigma/\mu \approx 50\%$ across the N_{Fit} runs, which contradicts the prior conclusion where the fits converged to a unique ξ^{Opt} . This is a consequence of the single-pump setup: the fringe shape depends only on effective combinations of the link parameters, leaving residual degrees of freedom in the full parameter space. Thus, lifting this degeneracy would require interferometric measurements at several pump values, as done for the local probabilities.

3.3 .Entanglement Witness

As explained previously, under the assumption of a Bell state mixture, the visibility measurement acts as an entanglement witness, quantifying the coherence of the delocalized excitation across the two considered modes. While providing a convenient way of witnessing the coherence, this approach holds only under restrictive assumptions on the state, which is not granted in the general description of an elementary link beyond the ideal low-pump regime. Additionally, it requires a non-local operation in which the two modes interfere, not practicable in the scenario of a 100 km elementary link. Therefore, we develop in this section to use an entanglement witness based on local operations and classical communication (LOCC), witnessing entanglement in the qubit subspace of two modes filled with arbitrary photon number, and working under arbitrary memory




Figure 3.9: Initial parameters (black crosses) and fitted parameters (orange crosses) across the $N_{\text{Fit}} = 200$ runs. All fits converge to a low MSE value $5 \times 10^{-9} \pm 10^{-14}$, while the fitted parameters fluctuate across the N_{Fit} , e.g. $\sigma/\mu \approx 50\%$ for the efficiency. This reflects the dependence of the fringe on parameter combinations, which leaves residual degrees of freedom with degeneracies in the fit.




Figure 3.10: Fit of the interference fringes as a function of the phase θ . All experimental data points lie within the 1σ envelope of the fitted model.




Figure 3.11: Optical setup for the measurement of the coincidence probability $\mathbb{P}_{\text{Coinc}}^A$ at Alice's node. The mode m_A is split on a 50-50 beam-splitter and probed by a joint double-click measurement Π_{Coinc}^A .

losses [122].

3.3.1 .Non-classicality

To begin with, let us focus on a LOCC-based non-classicality witness. While an entanglement witness focuses on the non-separability of a state, a non-classicality witness certifies the state cannot be described as a statistical mixture of classical states. The notion of non-classicality, however, depends on the chosen criterion: in quantum optics, the two main definitions rely on the negativity of the Wigner function [129] and on the non-positivity of the Glauber-Sudarshan P -representation. We focus here on the latter, which fit the framework described below and capture minimal non classically². The criteria states that any classical state admits a P -representation [130, 131]

$$\rho = \int_{\alpha} P(\alpha) |\alpha\rangle \langle \alpha| d\alpha, \quad \text{with } P \geq 0. \quad (3.47)$$

Called *quasi-probability distribution*, P lies at the heart of the classical character of the state. When positive, P turns to a probability distribution, describing ρ as a mixture of coherent states $|\alpha\rangle$, where the stochasticity carried by P reflects a lack of classical information rather than quantum indeterminacy³. Conversely, when P is non-positive, the state cannot be described as a classical mixture of coherent states, and a quantum basis such as the Fock basis $\{|n\rangle\}$ is required for its diagonalization, witnessing the non-classicality of the state.

²The Wigner function is the convolution of the Glauber-Sudarshan P -representation with a gaussian function, thus $W \leq 0 \implies P \leq 0$.

³Concretely, such a state can be produced from two lasers combined with classical communication.

We exploit this P -representation to build a non-classicality witness. While [132] relies on a measure-theoretic argument, we propose here a geometric formulation based on scalar product. For a two-mode state ρ in the modes $m_A^\dagger$ and $m_B^\dagger$, we define the bilinear form

$$(f, g) = \left\langle : f(m_A^\dagger m_A) g(m_B^\dagger m_B) : \right\rangle_\rho = \text{Tr} \left[: f(m_A^\dagger m_A) g(m_B^\dagger m_B) : \rho \right], \quad (3.48)$$

where $::$ denotes the normal-ordering operator. Under the assumption that P is positive, this bilinear form defines a scalar product, hence

$$P \text{ positive} \implies (f, g) \text{ scalar product} \implies \frac{|(f, g)|^2}{(f, f)(g, g)} \leq 1, \quad (3.49)$$

where the last implication is the Cauchy-Schwarz inequality. By contraposition,

$$\frac{\left| \left\langle : f(m_A^\dagger m_A) g(m_B^\dagger m_B) : \right\rangle_\rho \right|^2}{\left\langle : f(m_A^\dagger m_A)^2 : \right\rangle_\rho \left\langle : g(m_B^\dagger m_B)^2 : \right\rangle_\rho} > 1 \implies P \text{ not positive}, \quad (3.50)$$

yielding a criterion that witnesses non-classicality.

For our purpose, we use functions based on local click statistics

$$f(m_A^\dagger m_A) = \Pi_1^{m_A}(\eta_A/2) \quad \text{and} \quad g(m_B^\dagger m_B) = \Pi_1^{m_B}(\eta_B/2). \quad (3.51)$$

The Cauchy-Schwarz inequality then reads

$$\Xi = \frac{|\mathbb{P}_{\text{Loc}}^{11}(\eta_A/2, \eta_B/2)|^2}{\mathbb{P}_{\text{Coinc}}^A(\eta_A) \mathbb{P}_{\text{Coinc}}^B(\eta_B)} \leq 1, \quad (3.52)$$

where $\mathbb{P}_{\text{Loc}}^{11}(\eta_A/2, \eta_B/2)$ is the click-coincidence probability between Alice's mode m_A and Bob's mode m_B with halved efficiencies, and $\mathbb{P}_{\text{Coinc}}^A(\eta_A)$ is the probability of detecting at least two photons in m_A , measured via a twofold coincidence at Alice's location onto two detectors located after a beam-splitter, as depicted in figure 3.11 (and symmetrically for Bob with $\mathbb{P}_{\text{Coinc}}^B(\eta_B)$). For clarity, we simply write $\mathbb{P}_{\text{Coinc}}^{A/B}$ and $\mathbb{P}_{\text{Loc}}^{11}$, keeping in mind that $\mathbb{P}_{\text{Loc}}^{11}$ refers to halved efficiencies $\eta_A/2$ and $\eta_B/2$.

We first evaluate the Cauchy-Schwarz coefficient Ξ in the low-pump regime. Considering the 1-photon contribution, the perturbative expansion of section 1.1.1 shows that both the numerator and the denominator of Ξ vanish, leaving the witness undefined. Concretely, $\mathbb{P}_{\text{Loc}}^{11}$ requires a click coincidence between Alice and Bob, and $\mathbb{P}_{\text{Coinc}}^A$ (resp. $\mathbb{P}_{\text{Coinc}}^B$) requires at least two excitations in m_A (respectively m_B), which is forbidden in the 1-photon sector as the heralded state carries a single excitation delocalized over the two memories

$$\begin{aligned} \mathbb{P}_{\text{Coinc}}^{A/B} &= o(p) \\ \mathbb{P}_{\text{Loc}}^{11} &= o(p) \end{aligned} \quad (3.53)$$




Figure 3.12: Cauchy-Schwarz coefficient Ξ according to the repeater length for different values of the pump parameter.

The witness must therefore be evaluated at least up to the 2-photon sector. This contribution is computed in appendix D.5, giving

$$\Xi = \left(\frac{2 - \eta_{\text{Lnk}}}{4 - 3\eta_{\text{Lnk}}} \right)^2, \quad (3.54)$$

which is not η_A nor η_B depends. For $x \in [0, 1]$, the function $x \mapsto \left(\frac{2-x}{4-3x} \right)^2$ is monotonically increasing from $f(0) = 1/4$ to $f(1) = 1$. The Cauchy-Schwarz inequality is therefore never violated, and the non-classical nature of ρ_{Lnk} cannot be witnessed in the low-pump regime.

We now extend the study to the full pump range by deriving an exact expression of Ξ . The exact derivation leads to a heavy expression not suitable for analytical study, we thus give numerical results. See appendix D.3 and appendix D.6 for the full derivation of $\mathbb{P}_{\text{Loc}}^{11}$ and $\mathbb{P}_{\text{Coinc}}$. Figure 3.12 shows Ξ as a function of the repeater length L_{Lnk} for several values of the pump parameter p , and we conclude, as expected, that no (η_{AB}, p) pair violates the classical bound, and the low-pump limit derived earlier is recovered for small p .

The numerical plot shows a counter-intuitive trend where Ξ approaches the classical bound as p increases. This is opposite to what one would expect, as ρ_{Lnk} tends to a pure Bell pair where a strong non-classical signature should be observed. In fact, this reflects the fact that the Cauchy-Schwarz inequality is a sufficient but not necessary condition for non-classicality, which fails to witness the non-positivity of P for these POVM statistics: ρ_{Lnk} is strongly non-classical and entangled, but the chosen click statistics admit a classical-like structure leaving the witness silent.

3.3.2 . Constructing the Witness




Figure 3.13: Optical setup for the measurement of the displaced local probabilities. This joint measurement is equivalent to a projection onto the Bloch sphere along the axes $\vec{n}_\alpha$ and $\vec{n}_\beta$ in the qubit subspace.

While the non-classical witnesses described above focus on witnessing the non-positivity of P , entanglement witnesses are more restrictive and instead focus on the non-separability of a state. In the following paragraph we develop an entanglement witness suitable for the certification of entanglement in an elementary link, i.e. LOCC-based and robust to losses. As the links carry entanglement through the Bell pair from the 1-photon sector, we focus on the qubit space while accounting for the contribution of higher photon-number components. Concretely, the target is a LOCC-based witness, working in a lossy environment, witnessing entanglement in the qubit subspace $\{|0\rangle, |1\rangle\}$ and accounting for multi-photon contributions.

Building non-commuting measurement

Let's recall that witnessing entanglement via local measurement requires statistics from two non-commuting measurement settings [133]. The statistics produced within a single local measurement setting can always be reproduced by a separable state. Formally, for an arbitrary bipartite state $\rho_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$ and two POVMs $\{M_i^A\}, \{M_j^B\}$ each corresponding to a single measurement setting:

$$\exists \sigma_{AB} \text{ separable s.t. } \forall (i, j), \quad \text{Tr} [M_i^A \otimes M_j^B \rho_{AB}] = \text{Tr} [M_i^A \otimes M_j^B \sigma_{AB}]. \quad (3.55)$$

This conclusion extends to any collection of mutually commuting settings: commuting POVMs are simultaneously diagonalizable, thus their statistics arise as classical post-processings of a single common measurement, to which the above argument applies. This establishes the necessity of at least two non-commuting measurement settings for

the witness.

In a qubit context, the two settings are typically chosen as the Pauli operators X and Z , or any rotation, the only requirement being that the chosen settings carry non-commuting components. When considering optical modes with arbitrary photon-number excitations, a suitable procedure is required to realize such non-commuting operations in the qubit subspace $\text{Span}\{|0\rangle, |1\rangle\}$. In [134], the authors propose to use displacement operators to simulate rotations in the qubit space. Concretely, as depicted in figure 3.13, a displacement is applied before the photon detection, leading to the set of POVMs

$$\Pi_0^{a,\alpha}(\eta, \text{dc}) = \mathcal{D}^\dagger(\alpha) \Pi_0^a(\eta, \text{dc}) \mathcal{D}(\alpha), \quad (3.56)$$

$$\Pi_1^{a,\alpha}(\eta, \text{dc}) = \mathcal{D}^\dagger(\alpha) \Pi_1^a(\eta, \text{dc}) \mathcal{D}(\alpha), \quad (3.57)$$

where the upper index a, α denotes detection on mode a displaced by α , so that $\alpha = 0$ reduces to detection on mode $a, \alpha \rightarrow a$ itself. Defining the observable for the setting α in mode $a^\dagger$

$$\hat{o}(a : \alpha) := \Pi_0^{a,\alpha} - \Pi_1^{a,\alpha}, \quad (3.58)$$

that we will simply write $\hat{o}_\alpha$, the resulting statistics, restricted to the qubit-space basis $\{|0\rangle, |1\rangle\}$, is equivalent to a projection on the Bloch sphere along the axis


$$\vec{n}_\alpha \propto \begin{pmatrix} -e^{-\eta|\alpha|^2} |\alpha| \eta \cos(\theta) \\ e^{-\eta|\alpha|^2} |\alpha| \eta \sin(\theta) \\ \frac{1}{2} e^{-\eta|\alpha|^2} \eta (1 - |\alpha|^2 \eta) \end{pmatrix}, \quad (3.59)$$

with probability

$$\mu_\alpha = \sqrt{\eta^2 e^{-2|\alpha|^2 \eta} \left(|\alpha|^2 \left(|\alpha|^2 \eta^2 - 2\eta + 4 \right) + 1 \right)}. \quad (3.60)$$

Beyond providing a geometric intuition, this construction unlocks parametrized measurement on the Bloch sphere, allowing the construction of two POVMs associated with non-collinear projections $\vec{n}_{\alpha_1}$ and $\vec{n}_{\alpha_2}$, constructed from well chosen displacements α_1 and α_2 .

From this representation, figure 3.14a shows the evolution of $\vec{n}_\alpha$ as a function of $|\alpha|$ for several values of the phase $\theta = \arg(\alpha)$. This representation shows upper hemisphere of the Bloch sphere is reachable with a vector norm that decreases with $|\alpha|$, as the displacement populates the mode with coherent photons which reduces state discriminability. This reveals a trade-off in the choice of α_1 and α_2 , increasing the angle between $\vec{n}_{\alpha_1}$ and $\vec{n}_{\alpha_2}$ by setting $\theta_1 - \theta_2 = \frac{\pi}{2}$ favours entanglement detection, but the displacement amplitudes required also degrade the measurement contrast. Figure 3.14b illustrates this trade-off: the green arrows target orthogonal measurements at the cost of a weak contrast due to the arrow collinear to the X axis, while the blue arrows target high contrast at the cost of orthogonality. Given this trade-off, the optimal choice is obtained by numerical optimization of the targeted witness.



(a) Evolution of $\vec{n}_\alpha$ as a function of $|\alpha|$ for different values of the phase $\arg(\alpha) = \theta$.

(b) Evolution of $\vec{n}_\alpha$ as a function of $|\alpha|$ for different values of the photon detection efficiency η , θ is fixed.

Figure 3.14: Evolution of the projection vector $\vec{n}_\alpha$ on the Bloch sphere, in 3D and 2D. The amplitude $|\alpha|$ controls the meridional displacement, while the phase $\arg(\alpha)$ rotates the vector around the z -axis, enabling the construction of non-commuting measurements. The measurement efficiency η degrades the detection, which is here reflected in a reduction of the norm of $\vec{n}_\alpha$.

CHSH Inequality

To begin with, we propose the study of CHSH inequality [135]. Considering a bipartite system a and b with respectively the settings $\alpha \in \{0, 1\}$ and $\beta \in \{0, 1\}$, and outcome in $\{-1, 1\}$, the CHSH score stands

$$S = |\mathbb{E}[0, 0] + \mathbb{E}[0, 1] + \mathbb{E}[1, 0] - \mathbb{E}[1, 1]|, \quad (3.61)$$

$$\mathbb{E}[i, j] = \sum_{a, b \in \{-1, 1\}} a \cdot b \cdot \mathbb{P}(a, b | i, j), \quad (3.62)$$

with the classical bound 2 and maximum quantum bound $2\sqrt{2}$. In our context, we define the two settings $\alpha \in \{\alpha_1, \alpha_2\}$ and $\beta \in \{\beta_1, \beta_2\}$ for the observable

$$\mathcal{O}(a : \alpha, b : \beta) = \hat{o}_\alpha \otimes \hat{o}_\beta, \quad (3.63)$$

simply written $\mathcal{O}_{\alpha, \beta}$. The CHSH score reads

$$S = |\mathbb{E}[\alpha_1, \beta_1] + \mathbb{E}[\alpha_1, \beta_2] + \mathbb{E}[\alpha_2, \beta_1] - \mathbb{E}[\alpha_2, \beta_2]|, \quad (3.64)$$

$$(3.65)$$

with the expected value defined from the observable

$$\mathbb{E}[\alpha, \beta] = \langle \hat{\alpha} \otimes \hat{\beta} \rangle_{\rho_{\text{Lnk}}} \quad (3.66)$$

$$= \left\langle \left(\Pi_0^{m_A, \alpha} - \Pi_1^{m_A, \alpha} \right) \otimes \left(\Pi_0^{m_B, \beta} - \Pi_1^{m_B, \beta} \right) \right\rangle \quad (3.67)$$

$$= \mathbb{P}_{\text{Loc}}^{00}(\alpha, \beta) - \mathbb{P}_{\text{Loc}}^{01}(\alpha, \beta) - \mathbb{P}_{\text{Loc}}^{10}(\alpha, \beta) + \mathbb{P}_{\text{Loc}}^{11}(\alpha, \beta) \quad (3.68)$$

with $\mathbb{P}_{\text{Loc}}^{ij}(\alpha, \beta)$ the joint probability of having no-click (0) or click (1) at Alice's side (i) with displacement α and Bob's side (j) with displacement β .

The CHSH score S is expressed from the joint probabilities $\mathbb{P}_{\text{Loc}}^{ij}$, and the values of (α_1, β_1) and (α_2, β_2) that maximize S are found by numerical optimization on those probabilities. The full derivation and expression of $\mathbb{P}_{\text{Loc}}^{ij}(\alpha, \beta)$ is given in appendix D.3, here we focus on numerical results. Figure 3.15 shows the optimized CHSH score as a function of the local efficiency $\eta_A = \eta_B = \eta_{\text{AB}}$, with $\eta_{\text{AB}} = \eta_{\text{Mem}} \eta_{\text{Det}}$, for several values of the pump parameter. The best violations are reached at low pump parameter, in agreement with the perturbative analysis. For high pump parameter, the contributions from higher photon-number sectors dilute the Bell pair carried by the 1-photon sector, reducing the score S as maximal violation is reached with a pure Bell state. The threshold $\eta_0 \approx 0.83$ and the saturation value $S \approx 2.7$ at $\eta = 1$ discussed in [134] are recovered. This figure also illustrates that the CHSH test is not suited to witness the elementary link state ρ_{Lnk} entanglement. The plot shows CHSH violation is highly sensitive to the efficiency of the local modes, even in the ideal Bell-pair scenario the violation requires $\eta \geq 0.83$, while realistic memory efficiencies reach $\eta_{\text{Mem}} \approx 0.5$ in optimistic estimation. Despite the advantage of requiring no prior assumption on the state, the CHSH test trades this generality for a strong sensitivity to mode losses, motivating the need for a loss-aware witness suited to ρ_{Lnk} .

Witness derivation

Based on [122], we develop a witness robust to all noise sources. The witness is based on the Peres-Horodecki criterion, which states that separable two-qubit states have a positive partial transpose (PPT) [136, 137]. While the CHSH score treats the state as a black box and relies only on measurement statistics, this criterion operates at the level of the density matrix itself, allowing assumptions on the state for finer bound. Concretely, for a setting (α, β) , we define the score $w_{\rho_{\text{Lnk}}}(\alpha, \beta)$ witnessing the non-separability of ρ_{Lnk} when negative, such that

$$w_{\rho_{\text{Lnk}}}(\alpha, \beta) = w_{\text{PPT}}^{\rho_{\text{Lnk}}}(\alpha, \beta) - \langle \bar{\mathcal{O}}_{\alpha, \beta} \rangle_{\rho_{\text{Lnk}}} < 0 \quad \implies \quad \begin{array}{l} \rho_{\text{Lnk}} \text{ non-separable} \\ \text{in the qubit space} \end{array} \quad (3.69)$$




Figure 3.15: Optimized CHSH score S of the exact elementary-link state ρ_{Lnk} as a function of the memory efficiency η_{Mem} , for several values of the pump parameter p . The dashed black line corresponds to the optimized CHSH score of an ideal Bell pair. The CHSH inequality is violated by the Bell pair for $\eta_{\text{Mem}} > 0.84$, recovering the result of [134] up to optimization fluctuations. The elementary-link CHSH score approaches the Bell-pair score as $p \rightarrow 0$, consistent with the low-pump regime heralding a pure Bell pair. This witness is not suitable for our scenario: it requires both a high memory efficiency ($\eta_{\text{Mem}} > 0.84$) and a very low pump regime ($p < 10^{-2}$), which is incompatible with achieving a high distribution rate.

where $\overline{\mathcal{O}}$ is the phase-averaged observable

$$\overline{\mathcal{O}}_{\alpha,\beta} = \frac{1}{2\pi} \int_0^{2\pi} e^{i\phi(m_A^\dagger m_A + m_B^\dagger m_B)} \mathcal{O}_{\alpha,\beta} e^{-i\phi(m_A^\dagger m_A + m_B^\dagger m_B)} d\phi, \quad (3.70)$$

with $\mathcal{O}$ defined in equation (3.63). The separability bound $w_{\text{PPT}}^{\rho_{\text{Lnk}}}$ is defined as

$$w_{\text{PPT}}^{\rho_{\text{Lnk}}}(\alpha, \beta) = \max_{\sigma} \langle \overline{\mathcal{O}}_{\alpha,\beta} \rangle_{\sigma} \quad (3.71)$$

$$\text{s.t.} \quad \begin{aligned} \text{(a)} \quad & \sigma^{TA} \succeq 0 \\ \text{(b)} \quad & \sigma_{i,i} = \mathbb{P}_{\text{Loc}}^{i,i}, \end{aligned} \quad (3.72)$$

with (a) the PPT criterion, and (b) an additional assumption stating that the diagonal terms of the state variable σ are fixed to the measured local probabilities. Note that the bounds depends on ρ_{Lnk} with (b), for clarity we will drop the ρ_{Lnk} and write w_{PPT} , keeping in mind it is defines for ρ_{Lnk} . For clarity, we write w and w_{PPT} in the following, keeping in mind that these quantities are defined for a setting (α, β) on the state ρ_{Lnk} . This last assumption lies at the heart of the robustness against losses: it injects extra information on losses into the optimization, making the witness “loss-aware”.

Let us give a concrete illustration. Consider the ideal low pump scenario where a Bell pair has been generated in the elementary link, with no dark counts ($\text{dc}_A = \text{dc}_B = 0$) and symmetric losses ($\eta_A = \eta_B = \eta_{AB}$). In this context, the calculation from appendix D.8 leads to

$$w^{\text{In}} = 4\eta_{AB}^2 \cdot \text{Re}[\alpha^* \beta] \cdot e^{-\eta_{AB}(|\alpha|^2 + |\beta|^2)}, \quad (3.73)$$

the notation w^{In} refers to the witness in the qubit space, neglecting the multi-photon contributions: as we assume low pump regime, the state is living in the subspace $\mathcal{H}^{\text{In}} = \text{Span}\{|ij\rangle\}$ with $i, j = 0, 1$. From this derivation the witness is negative whenever $\text{Re}[\alpha^* \beta] < 0$, for any value of the local efficiency η_{AB} , as depicted in figure 3.16. The witness therefore detects entanglement for arbitrary losses, reaching its minimum for

$$\alpha_{\text{Opt}} = \pm \frac{1}{\sqrt{2\eta_{AB}}} e^{i\phi} \quad \text{and} \quad \beta_{\text{Opt}} = \mp \frac{1}{\sqrt{2\eta_{AB}}} e^{i(\phi \pm \pi)}. \quad (3.74)$$

This motivates the development of such a witness for detecting entanglement in the elementary link with lossy memories.

The previous derivation holds in the qubit subspace, i.e. assuming ρ_{Lnk} lives in $\mathcal{H}^{\text{In}}$, which is valid only in the low-pump limit. In realistic scenario higher n -photon sectors contribute to the state and the bound must be corrected by considering the full Fock space $\mathcal{H} = \mathcal{H}^{\text{In}} \oplus \mathcal{H}^{\text{Out}}$. Note that the PPT criterion is necessary and sufficient for two qubit systems, but is only sufficient for higher dimensional systems. Thus, by considering the multi-photon contribution some entangled states will never be detect by the witness.

Based on the calculation from appendix D.9, for a setting (α, β) the total witness reads

$$w = (\text{In} w_{\text{PPT}} + \text{Out} w_{\text{PPT}}) - \langle \overline{\mathcal{O}} \rangle_{\rho_{\text{Lnk}}} . \quad (3.75)$$

The expectation value of the observable is computed from the displaced joint local probabilities $\mathbb{P}_{\text{Loc}}^{ij}(\alpha, \beta)$ derived in appendix D.3:

$$\langle \overline{\mathcal{O}} \rangle_{\rho_{\text{Lnk}}} = \mathbb{P}_{\text{Loc}}^{00}(\alpha, \beta) - \mathbb{P}_{\text{Loc}}^{01}(\alpha, \beta) - \mathbb{P}_{\text{Loc}}^{10}(\alpha, \beta) + \mathbb{P}_{\text{Loc}}^{11}(\alpha, \beta) . \quad (3.76)$$

The in and out qubit-subspace contributions are

$$\text{In} w_{\text{PPT}} = \sum_{ij} o_{ij} \mathbb{P}_{\text{Loc}}^{ij} + 2|o'_{01}| \cdot \min \left(\sqrt{\mathbb{P}_{\text{Loc}}^{00} \mathbb{P}_{\text{Loc}}^{11}}, \sqrt{\mathbb{P}_{\text{Loc}}^{01} \mathbb{P}_{\text{Loc}}^{10}} \right) , \quad (3.77)$$

$$\text{Out} w_{\text{PPT}} = 2(\mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}}) + 4S_{\text{max}} \sqrt{\mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}} \left(\frac{1}{2} - \mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}} \right)} ,$$

where $\mathbb{P}_{\text{Coinc}}^{\text{A}}$ and $\mathbb{P}_{\text{Coinc}}^{\text{B}}$ respectively denote the local coincidence detection probabilities of Alice and Bob, i.e. the joint click after a 50/50 beam splitter on Alice's or Bob's local mode m_{A} or m_{B} . From appendix D.7, the o_{ij} coefficients are given by the observable matrix elements

$$\begin{aligned} \langle 0 | \hat{o}_{\alpha} | 0 \rangle &= 2\kappa \cdot e^{-\eta_o |\alpha|^2} - 1, \\ \langle 0 | \hat{o}_{\alpha} | 1 \rangle &= -2\kappa \cdot \eta_o \alpha^* e^{-\eta_o |\alpha|^2}, \\ \langle 1 | \hat{o}_{\alpha} | 1 \rangle &= 2\kappa \cdot \left[1 + \eta_o \left(\eta_o |\alpha|^2 - 1 \right) \right] e^{-\eta_o |\alpha|^2} - 1, \end{aligned} \quad (3.78)$$

with the convention $o_{ij} = \langle i | \hat{o}_{\alpha} | j \rangle \langle i | \hat{o}_{\beta} | j \rangle$ on the diagonal and $o'_{ij} = \langle i | \hat{o}_{\alpha} | j \rangle \langle j | \hat{o}_{\beta} | i \rangle$ off the diagonal. The witness is thus expressed as a function of the displaced joint local probabilities $\mathbb{P}_{\text{Loc}}^{ij}(\alpha, \beta)$ and the coincidence probabilities $\mathbb{P}_{\text{Coinc}}^{\text{A/B}}$.

Constraint (b) of equation (3.72) carries the physical content of the witness and deserves a closer look. The constraint imposes $\mathbb{P}_{\text{Loc}}^{ij} = \rho_{\text{Lnk}}^{ij}$, fixing the diagonal elements of the state to the measured local joint probabilities, which is precisely what makes the witness loss-aware. However, our model assumes a noise-less state with all losses pushed into the detectors, so in fact $\mathbb{P}_{\text{Loc}}^{ij} \leq \rho_{\text{Lnk}}^{ij}$. Therefore, from the witness perspective, the considered state is an effective state $\tilde{\rho}_{\text{Lnk}}$ that absorbs all losses, both detector and memory losses, into the state itself. The witness operates on this effective state, since LOCC operations cannot create entanglement, non-separability of $\tilde{\rho}_{\text{Lnk}}$ implies non-separability of ρ_{Lnk} :

$$\rho_{\text{Lnk}}^{\sim} \text{ non-separable} \quad \implies \quad \rho_{\text{Lnk}} \text{ non-separable.} \quad (3.79)$$




Figure 3.16: In-qubit entanglement witness $w^{\text{in}} = \langle \overline{\mathcal{O}} \rangle - \text{In} w_{\text{pPT}}$ of an ideal Bell pair $|\Psi^+\rangle \langle \Psi^+|$ as a function of the settings (α, β) . Symmetric configuration, no dark counts, with realistic experimental parameters from table 1 of the introduction. The dashed line marks the zero contour. Two disjoint witnessing regions are visible, each defined by $\alpha \cdot \beta < 0$. The optimal value $w_{\text{opt}} \approx -0.1$ is reached at $\alpha_{\text{opt}} = -\beta_{\text{opt}} = \pm \frac{1}{\sqrt{\eta_{\text{AB}}}} \approx 1.6$.

Numerical results

The witness derived above provides a way of certifying the entanglement of ρ_{Lnk} from experimental quantities: the joint local probabilities $\mathbb{P}_{\text{Loc}}^{ij}$, the displaced joint local probabilities $\mathbb{P}_{\text{Loc}}^{ij}(\alpha, \beta)$, and Alice's and Bob's coincidence probabilities $\mathbb{P}_{\text{Coinc}}^{\text{A}}$ and $\mathbb{P}_{\text{Coinc}}^{\text{B}}$. All these quantities have already been derived in previous calculations. The displaced joint local probabilities have been derived when studying the CHSH score in appendix D.3, the joint local probabilities correspond to the displaced ones evaluated at $\alpha = \beta = 0$, and the coincidence probabilities were derived in appendix D.6 when studying the Cauchy-Schwarz inequality. Due to the heavy form of these expressions, we do not provide analytical results and focus on their numerical implementation.

Before diving into the numerical implementation, the efficiencies require a bit more attention. As done in this manuscript, detection losses are conventionally modelled as imperfections in the detector, leaving ρ_{Lnk} untouched. For undisplaced measurements, the two pictures provides identical click statistics. For displaced measurements, they do not commute and pushing the losses to the detection rescales the displacement amplitude. Quantitatively, the click probability satisfies the identity [122]

$$\text{Tr} \left[\mathcal{D}(\alpha)^\dagger R_{\text{A}}^{m_{\text{A}}^\dagger m_{\text{A}}} \mathcal{D}(\alpha) \cdot \rho_{\text{Lnk}} \right] = \text{Tr} \left[\mathcal{D}(\sqrt{\eta} \alpha)^\dagger |0\rangle \langle 0| \mathcal{D}(\sqrt{\eta} \alpha) \cdot U_\eta \rho_{\text{Lnk}} U_\eta^\dagger \right], \quad (3.80)$$

where $U_\eta = e^{\phi(m_{\text{A}}^\dagger e - m_{\text{A}} e^\dagger)}$ is the loss channel of transmission η acting on the state. The left side describes a noisy detector preceded by a displacement α acting on the state ρ_{Lnk} . The

right side describes an ideal detector preceded by a rescaled displacement $\sqrt{\eta} \alpha$ acting on the lossy state $U_{\eta} \rho_{\text{Lnk}} U_{\eta}^{\dagger}$. In our context, calling $\eta_{\mathcal{O}}$ the efficiency of the observable $\overline{\mathcal{O}}$ (related to the matrix elements derived in equation (3.78)), and η_{ρ} the efficiency intrinsic to the state ρ_{Lnk} , the equivalence stands

	Modelization	$\longleftrightarrow$	Targeted Scenario	
	All losses in the detection		Loss in the detection and loss in the state	
Detection:	$(R_{\rho} R_{\mathcal{O}})^{a^{\dagger} a}$	$\longleftrightarrow$	$R_{\mathcal{O}}^{a^{\dagger} a}$	(3.81)
Displacement:	$\frac{\alpha}{\sqrt{\eta_{\rho}}}$	$\longleftrightarrow$	α	
State:	ρ_{Lnk}	$\longleftrightarrow$	$U_{\eta_{\rho}} \rho_{\text{Lnk}} U_{\eta_{\rho}}^{\dagger}$	

allowing the modelization of a lossy state η_{ρ} under the detection of a lossy observable $\eta_{\mathcal{O}}$. As the witness has been set under the constraint $\rho_{\text{Lnk}}^{ij} = \mathbb{P}_{\text{Loc}}^{ij}$, and the local probability is under detection and memory losses η_{Det} and η_{Mem} , the witness sees an effective state $\tilde{\rho}_{\text{Lnk}}$ carrying both the losses from detection and memories, thus $\eta_{\rho} = \eta_{\text{Det}} \cdot \eta_{\text{Mem}}$. The matrix elements of $\eta_{\mathcal{O}}$ therefore only carries the losses from the displacement $\eta_{\mathcal{O}} = \eta_{\text{Disp}}$. The table 3.3 shows the effective efficiencies for all the considered quantities.

Quantity	Symbol	Associated efficiency
Local probabilities (no displacement)	$\mathbb{P}_{\text{Loc}}^{ij}$	$\eta_{\text{Mem}} \cdot \eta_{\text{Det}}$
Local probabilities (with displacement)	$\mathbb{P}_{\text{Loc}}^{ij}(\alpha, \beta)$	$\eta_{\text{Mem}} \cdot \eta_{\text{Disp}} \cdot \eta_{\text{Det}}$
Coincidence probabilities	$\mathbb{P}_{\text{Coinc}}^{A,B}$	$\eta_{\text{Mem}} \cdot \eta_{\text{Coinc}} \cdot \eta_{\text{Det}}$
Matrix elements of $\overline{\mathcal{O}}$	$\mathcal{O}_{ji}$	η_{Disp}
Effective displacement	$\tilde{\alpha}$	$\alpha / \sqrt{\eta_{\text{Mem}} \cdot \eta_{\text{Det}}}$

Table 3.3: Effective efficiencies carried by each quantity entering the entanglement witness. The amplitude of the displacement is rescaled to $\tilde{\alpha}$ to absorb the upstream losses $\eta_{\text{Mem}} \cdot \eta_{\text{Det}}$ via the transfer identity (3.80).

We start by studying the witness as a function of α and β under symmetric assumptions $p_A = p_B = p$, $\eta_A = \eta_B$, and no dark counts $\text{dc}_A = \text{dc}_B = 0$. The remaining experimental parameters are taken from table 1, chosen to reflect a realistic experimental scenario. Figure 3.17 displays the witness value over the (α, β) plane for three values of the pump parameter p , highlighting with dashed border the region where entanglement is detected. For low p , the witness recovers the shape of the in-qubit approximation (cf. figure 3.16), consistent with the idealized low-pump regime where $\rho_{\text{Lnk}} = |\Psi^+\rangle \langle \Psi^+| + o(p)$ and ${}^{\text{Out}}w_{\text{pPT}} = o(p) \rightarrow 0$, we provide a quantitative analysis of this comparison in figure 3.18. The entanglement-witnessing region splits into two symmetric lobes, referred to as *witnessing lobes*, with an optimal value $w_{\text{opt}} \approx -0.1$ reached at $\alpha_{\text{opt}} = -\beta_{\text{opt}} = \pm \frac{1}{\sqrt{\eta_{\text{AB}}}} \approx 1.6$. The




Figure 3.17: Entanglement witness $w = w_{\text{PPT}} - \langle \overline{\mathcal{O}} \rangle_{\rho_{\text{Lnk}}}$ of the realistic elementary link ρ_{Lnk} as a function of the settings (α, β) , for three values of the pump parameter p . Symmetric configuration, no dark counts, with realistic experimental parameters from table 1. The dashed line marks the zero contour. Panels (a) and (b) show two disjoint witnessing regions, while panel (c) shows no witnessing region, indicating a pump parameter cutoff $10^{-2} < p_{\text{cut}} < 10^{-1}$.

intermediate plot at $p = 10^{-2}$ shows that increasing the pump parameter the minimum witness value ($w_{\text{Opt}} \approx -0.05$) and contract the witnessing lobes. Thus higher pump values make entanglement harder to certify, simultaneously degrading the contrast (from -0.1 to -0.05) and reducing the area of the witnessing region. This clarifies also the in-qubit picture where the whole $\alpha \cdot \beta < 0$ region witnesses entanglement: it is actually the asymptotic limit $p \rightarrow 0$ in which the lobes extend to the entire top-left and bottom-right quarter. At $p = 10^{-1}$, the witness is out of its working regime: the witnessing lobes vanish, indicating a pump cutoff p_{cut} between 10^{-2} and 10^{-1} .

Focusing on this cutoff, we plot in figure 3.19 the optimized witness w_{Opt} over (α, β) as a function of the pump parameter p and the memory efficiency η_{Mem} , still under symmetric assumptions, with realistic parameters fixed from table 1. The plot reveals a cutoff at $p_{\text{cut}} \approx 4 \cdot 10^{-2}$, above which the witness no longer detects entanglement. This threshold is a strong constraint in practice: since the heralding rate scales with p , restricting the operating regime to $p < p_{\text{cut}}$ limit practical achievable distribution rate. It also confirms the in-qubit result of equation (3.73) where, below p_{cut} , the witness certifies entanglement for any value of the memory efficiency η_{Mem} . The same study performed with $\text{dc}_A = \text{dc}_B = 10^{-4}$, it shows that dark counts strongly degrade the witness contrast ($w_{\text{Opt}} \approx -0.2 \rightarrow -0.05$) without significantly shifting p_{cut} .

This numerical study demonstrates the ability of the witness to certify entanglement of an elementary link under realistic experimental conditions. The witness remains robust below the pump threshold $p_{\text{cut}} \approx 4 \cdot 10^{-2}$ for arbitrary memory losses, and remains operational in the presence of dark counts. A complementary study of the relative phase between α and β shows that the detuning rotates the overall witnessing pattern in the (α, β) plane without altering the cutoff.




Figure 3.18: Comparison of the entanglement witness in the low-pump regime for the exact elementary-link state ρ_{Lnk} (a) and the idealized Bell pair $|\Psi^+\rangle\langle\Psi^+|$ (b). Symmetric configuration, no dark counts, with realistic experimental parameters from table 1. The dashed line marks the zero contour. Panel (c) shows the difference between (a) and (b), with a relative error of order 10^{-2} concentrated at the center of the plane where numerical errors dominate.




Figure 3.19: Optimized entanglement witness w_{Opt} of the elementary link ρ_{Lnk} as a function of the pump parameter p and memory efficiency η_{Mem} , without dark counts (a) and with $dc = 10^{-4}$ (b), symmetric configuration, with experimental parameters from table 1. The dashed line marks the zero contour, revealing a clear pump cutoff at $p_{\text{Cut}} \approx 4 \cdot 10^{-2}$ independent of η_{Mem} . The witness certifies entanglement for arbitrary memory losses below p_{Cut} . Panel (b) further shows that dark counts degrade the witness contrast without significantly shifting p_{Cut} .

Part II

Quantum Circuit Synthesis

Introduction

The quantum unitary synthesis problem

As discussed in the general introduction of this manuscript, running a quantum algorithm on physical hardware requires translating abstract operations into concrete sequences of elementary operations. Given a universal gate set, *quantum circuit synthesis* aims to find a gate sequence implementing a target operation. In the circuit picture the relevant abstraction is the unitary evolution U acting on the logical qubits, so this task is also referred to as *unitary synthesis*.

This task faces an inherently hard combinatorial landscape: calling respectively n and g the number of qubits and the gate-count, at each layer $n \cdot g$ gate placements are available, thus the number of candidate circuits of depth d scales exponentially as $O((n \cdot g)^d)$ [83]. Meanwhile, the number of circuits implementing the target algorithm is sparse. This exponential growth, together with the solution sparsity, makes the synthesis of non-trivial unitaries rapidly intractable, even for moderate system sizes. Full-algorithm synthesis is therefore out of reach: useful algorithms require between 10^7 and 10^{11} logical gates [5, 71]. Thus, we focus instead on the decomposition of small subroutines for which exact synthesis remains tractable and can therefore be compiled independently.

In this second part of the manuscript, we address this combinatorial problem by framing synthesis as a sequential decision problem and tackle it with AlphaZero, a reinforcement learning agent originally developed for the game of Go, which excels at exploring large spaces with sparse solutions [115, 116]. We train agents to synthesize arbitrary subroutines under a fixed architecture with gate-set and connectivity constraints. Inspired by [118–120] for T -gate optimization, we explore dynamic circuit architectures, which involve mid-circuit measurements and feed-forward operations. This architecture yields lower T -counts than static circuits with all-to-all connectivity. As a proof of concept, we demonstrate agents capable of synthesizing standard subroutines such as the Toffoli and CCZ gates.

Outline

In Chapter 4, we frame quantum circuit synthesis as a single-player game, where an agent searches for a gate sequence to reach the targeted unitary. We then introduce Markov decision process and the reinforcement learning framework itself, and use it to formalize synthesis as a reinforcement learning problem⁴.

⁴For conciseness, the conceptual background of Markov decision process and reinforcement learning is deferred to the appendix.

The Chapter 5, we motivate the use of AlphaZero: synthesis requires long-term planning to maintain global coherence across the whole gate sequence, a regime where greedy strategies fail. We make this need explicit before introducing the AlphaZero agent, from the exploration-exploitation dilemma to the neural-network-guided Monte Carlo tree search.

Finally, Chapter 6 apply AlphaZero agents to quantum circuit synthesis. We detail the curriculum learning pipeline used for training, then benchmark agents synthesizing subroutines such as the Toffoli or CCZ gate under various connectivity constraints. We close with a critical discussion of the approach and perspectives for further development.

4 -A Game of Unitary Synthesis

4.1 .Framing synthesis as a single-player game

Approximate unitary synthesis can naturally be formulated as a continuous optimisation problem: from a given parameterized circuit $U(\{\theta_i\})$, the set of parameters is optimized to reach an optimal value $\theta^* = \{\theta_i^*\}$ that approximates the target unitary $\|U(\theta^*) - U_{\text{targ}}\| < \epsilon$. Exact synthesis is, by contrast, a discrete problem: it aims to find a gate sequence $\{g_i\}$ from a discrete gate set $\mathcal{G}$ that implements the target unitary exactly $\|\prod_i g_i - U_{\text{targ}}\| = 0$. Thus, exact quantum circuit synthesis can be framed as a constrained combinatorial problem: for a given gate set, one aims to find a sequence of gates that implements a target unitary, subject to hardware constraints such as qubit connectivity.

From this perspective, the problem naturally takes the form of a single-player game: from an allowed set of moves $\mathcal{M} = \{m_i\}$ the player aims to find an optimal sequence $\mathbf{m}^*$ that brings the game to a target configuration M_{targ} , possibly under additional constraints. An illustrative example is the Rubik's cube: starting from a scrambled configuration C_{rand} , the player must find a sequence of rotations $\mathbf{r} = \{r_k\}$, drawn from a discrete rotation set $\mathcal{R} = \{r_i\}$, that returns the cube to its solved state C_{ld} . In quantum circuit synthesis, starting with a configuration $U_{\text{targ}}^\dagger$ the compiler must find the sequence of gates $\mathbf{g} = \{g_i\}$ that leads to unitary U_{ld} .

The analogy can actually be pushed a bit further, helping to capture the main concepts of the framework.

- **Discrete:** both problems consist in finding a sequence of elementary operations drawn from a discrete set, respectively denoted $\mathcal{R} = \{r_i\}_i$ for the rotations of the Rubik's cube and $\mathcal{G} = \{g_i\}_i$ for the quantum gates.
- **Non-commutative:** the allowed cube rotations and gates both generate a non-abelian group, finite for the Rubik's cube and continuous for quantum circuit synthesis.
- **Universality:** in both cases, the group covers all cube configurations or unitaries (up to ϵ), reflecting the notion of universality of the set [89].
- **Two representations:** both games admit two complementary representations. A constructive one, listing the sequence of operations applied so far, $\mathbf{r} = \{r_i\}_n$ for the cube and $\mathbf{g} = \{g_i\}_n$ for the circuit, and, most importantly, an immediate one, capturing the current configuration of the cube C_n or the unitary of the circuit U_n . This second representation actually defines the state of the game: it contains all the information needed to finish it. Said otherwise, for any intermediate step, not knowing what has been done before does not make the resolution harder. This is

precisely the memoryless, or Markovian, property of the game, which we will exploit later in this part.

Within this game framework, synthesis reduces to a pathfinding problem in the Cayley graph of the group generated by the allowed moves, i.e. find a sequence of moves that maps $U_{\text{Id}} \rightarrow U_{\text{targ}}$. The discrete moves dynamic and memoryless representation naturally fits with the Markov Decision Processes formalism, a mathematical framework describing memory-less sequential decision problems. The next section briefly introduces the Markov Decision Process formalism and shows how reinforcement learning can be deployed as a solver.

4.2 .Reinforcement Learning

The conceptual framework behind reinforcement learning, from the origins of Markov Decision Processes to the agent-environment interaction loop, is presented in detail in appendix D.11. In this short section, we only recall the key notions and notations needed to formalize quantum circuit synthesis as a reinforcement learning problem, and refer to the appendix for a deeper understanding.

A *Reinforcement Learning* (RL) problem is a *Markov Decision Process* (MDP), a memory-less model where a *decision maker* acts on a system to optimize its earnings over time. Let us first give a brief overview of the MDP framework. A MDP is defined by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, R, \gamma \rangle$. The *state space* $\mathcal{S}$ is the set of all possible states, and the *action space* $\mathcal{A}$ the set of all possible actions. The *transition function* $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ encodes the game dynamics: for a state-action pair (s, a) it returns the distribution $\mathbb{P}(\cdot | s, a)$ over the next state, with $\mathbb{P}(s' | s, a)$ the probability of transitioning to s' while picking the action a . In this manuscript we only use deterministic transitions, i.e. an action a from a state s leads to a unique state s' . The *reward function* R defines the immediate feedback $r = R(s, a, s')$ along a transition $s \xrightarrow{a} s'$. In the context of deterministic transition, we have $R(s, a, s') = R(s')$. Finally, the *discount factor* $\gamma \in [0, 1)$ ensuring the convergence of the cumulative reward and weighting the importance of future rewards, with $\gamma \rightarrow 0$ favoring immediate rewards and $\gamma \rightarrow 1$ favoring long-term returns.

To picture these notions, consider the toy model of a frog jumping on a set of lily pads [138]. The state s is the pad currently occupied, and an action a is a jump toward a neighboring pad. Thus the state space $\mathcal{S}$ is the set of pads reachable by the frog, and the action space the set of jumps available from each pad. The transition function $\mathcal{P}$ encodes where the frog lands: deterministically on the targeted pad, or stochastically if jumps may miss. Assigning a reward to each landing, for instance a positive value on a pad carrying food and a negative one on a pad close to a predator, the frog must find a policy π to avoid the predator and reach the food, maximizing its cumulative reward.

By providing local signals, the reward function actually sets the global objective of the process. Calling t the time step of the interactive loop, the goal of the decision maker

is to find a strategy that maximizes the expected cumulative reward, or gain, denoted $G_{t:\infty} = \sum_{k=t}^{\infty} \gamma^k \cdot r_k$. With $r_t = R(s_{t-1}, a_{t-1}, s_t) = R(s_t)$, where the last equality stands for deterministic transition. The strategy itself is encoded in a *policy* $\pi(a | s)$, the probability distribution of picking an action a from a state s . To quantify how good a policy is, one introduces two value functions. The *state value function* measures the expected cumulative reward obtained from a state s when following the policy π

$$V^\pi(s) = \mathbb{E}_\pi \left[\sum_{k=0}^{\infty} \gamma^k \cdot r_{k+1} \mid s_0 = s \right]. \quad (4.1)$$

The *action value function* refines this notion by fixing the first action a taken from s before following π

$$Q^\pi(s, a) = \mathbb{E}_\pi \left[\sum_{k=0}^{\infty} \gamma^k \cdot r_{k+1} \mid s_0 = s, a_0 = a \right]. \quad (4.2)$$

Thus, solving an MDP amounts to finding an *optimal policy* π^* that maximizes these value functions over all states

$$\pi^* = \arg \max_{\pi} V^\pi(s), \quad \forall s \in \mathcal{S}. \quad (4.3)$$

Note that this notation is abusive and used for conciseness, a proper definition is given in the appendix.

While the MDP formalism sets the rules and objective of the game, it does not specify how to solve it, at this point the policy π must be explicitly provided. This is precisely the role of *Reinforcement Learning* (RL): rather than relying on a prior knowledge of the dynamics, an RL *agent* approximates the optimal policy by iteratively interacting with the *environment*, defined by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, R, \gamma \rangle$. From a certain perspective, the environment encodes the game dynamics and the agent the smartness. At each step t , the agent observes the state s_t , picks an action $a_t \in \mathcal{A}_{s_t}$ according to its current policy π_t , and receives in return the next state s_{t+1} together with a reward r_t . The agent uses this feedback to update its policy $\pi_t \rightarrow \pi_{t+1}$, progressively converging toward π^* . The loop repeats until either a terminal state s_{end} is reached or a maximum number of steps t_{max} is exceeded. This mechanism is illustrated in figure 4.1.

4.3 .Reinforcement Learning for Quantum Circuit Synthesis

At this point, all the conceptual ingredients are in place to formalize quantum circuit synthesis as a reinforcement learning problem. The agent builds the desired circuit by incrementally appending gates from a predefined universal gate set. Each appended gate leads to a new circuit, corresponding to a unitary evolution. A reward is then computed, for example by quantifying the proximity of the circuit with the target unitary. The game ends when the agent either synthesizes the target unitary or reaches a maximum number of steps. In this context, the environment is the quantum circuit representation equipped




Figure 4.1: Minimalist illustration of the reinforcement learning interaction mechanics. At each time step t , the agent picks an action a_t from $\mathcal{A}$ according to the policy $\pi(\cdot | s_t)$. This action drives the environment transition $s_t \xrightarrow{a_t} s_{t+1}$, producing a reward feedback $R(s_{t+1}) = r_{t+1}$. The interaction loop stops when the environment reaches the target state s_{targ} or the maximum step count t_{max} .

with a simulator able to compute unitaries from a gate sequence.

More formally, we denote by $\mathcal{G}$ the set of available gates compatible with hardware constraints such as qubit connectivity, by t_{max} the maximum number of steps, which also corresponds to the maximum allowed circuit depth, and by n the number of qubits. The gate set $\mathcal{G}$ is chosen universal. At each step t , the agent picks an action a_t corresponding to a gate $g \in \mathcal{G}$. This gate is appended to the circuit, from which the environment computes the associated unitary U_{t+1} , updating the state $s_t \rightarrow s_{t+1} = U_{t+1}$. A reward r_{t+1} is then derived from s_{t+1} according to the reward function. The agent receives both the updated state s_{t+1} and the associated reward r_{t+1} , and picks a new action a_{t+1} from $\mathcal{G}$. Through this loop, the agent incrementally builds the circuit one gate at a time.

Regarding the initial state, a natural choice is to start from the empty circuit, $s_0 = U_{\text{Id}}$, and let the agent append gates until it reaches the target $s_{\text{end}} = U_{\text{targ}}$. With this formulation each game has its own target U_{targ} , meaning the agent's objective switches from one game to another. This may complicate the learning process, as the agent must generalize across a wide range of objectives. To address this, we reverse the construction: instead of building the circuit implementing U_{targ} , the agent uncomputes $U_{\text{targ}}^\dagger$ back to the identity U_{Id} . Concretely, the circuit starts with the targeted unitary, meaning $s_0 = U_{\text{targ}}^\dagger$, and the agent aims to reach the identity by appending gates $g \in \mathcal{G}$ on the circuit. Thus, at each step t , the environment is in the state $s_t = U_t \cdot U_{\text{targ}}^\dagger$, where $U_t = \prod_{k=0}^t g_k$ is the unitary built from the appended gates. Finally, the gate sequence resulting in a successful synthesis implements U_{targ} , as $U_{\text{end}} \cdot U_{\text{targ}}^\dagger = U_{\text{Id}}$. With this formulation, all games share a common goal, reaching the identity, while differing only by their initial state $s_0 = U_{\text{targ}}^\dagger$. The figure 4.2 illustrate such an interactive loop.

To avoid trivial actions, such as choosing a gate that would cancel a previous one, the

action set available at step t is masked $\mathcal{A} \rightarrow \mathcal{A}_t$, allowing only non-trivial gates. Concretely, every gate $g \in \mathcal{G}$ is pushed through the current circuit using commutation relations. When g meets a gate g' that does not commute with, the resulting product $g \cdot g'$ is checked: if it reduces to the identity the corresponding action is discarded, else the gate is kept. This produces a simple non-triviality mask on $\mathcal{A}$, preventing the agent from picking trivial actions.

To formally define the game, we set the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, R, \gamma \rangle$ as follows. For clarity, the presentation does not follow the tuple order.

- **Action space $\mathcal{A}$:** at each step t , the agent picks a gate $g \in \mathcal{G}$ and the qubits it acts on: $a_t = \{g + \text{qubit index}\}$. The action space is therefore built by enumerating each gate of $\mathcal{G}$ with the qubit indices and allowed connectivity. In fact, this enumeration accommodates hardware constraints, such as restricted qubit connectivity or specific gates forbidden on certain qubits. For instance, a 3-qubit circuit with $\mathcal{G} = \{H, T, \text{CNOT}\}$ with nearest-neighbour connectivity leads to the action space $\mathcal{A} = \{H_1, H_2, H_3, T_1, T_2, T_3, \text{CNOT}_{12}, \text{CNOT}_{23}\}$. This enumeration is called *native gate set*, denoted $\mathcal{G}_n$, as it depends on the physical architecture. We thus have $\mathcal{A} = \mathcal{G}_n$.
- **Transition probability $\mathcal{P}$:** as previously announced the transitions are deterministic, meaning a state-action pair (s_t, a_t) leads to a unique next state s_{t+1} . Consequently, at each step, an action $a_t \in \mathcal{A}$ is applied to the left of the current state, yielding to the transition $s_t \rightarrow s_{t+1} = a_t \cdot s_t$. Let us notice the quantum circuit simulator only requires one multiplication per step: $s_{t+1} = g_t \cdot s_t$. Combined with $U_{\text{targ}}^\dagger$ as initial state, the state evolves as

$$s_{t+1} = g_t \cdot U_t \cdot U_{\text{targ}}^\dagger \quad \text{with} \quad U_t = \prod_{k=0}^{t-1} g_k. \quad (4.4)$$

- **State space $\mathcal{S}$:** formally the state space contains all unitaries reachable from the native gate set. Provided that $\mathcal{G}_n$ is universal, the set of unitaries generated is dense in $\text{SU}(2^n)$, so that $\mathcal{S} \subseteq \text{SU}(2^n)$. In practice, the agent operates under a finite budget of t_{max} steps, restricting $\mathcal{S}$ to the unitaries reachable within at most t_{max} gates. This finite-depth restriction makes the state space combinatorially large but discrete and finite, in contrast with the continuous group $\text{SU}(2^n)$. Thus, denoting by $\langle \mathcal{H} \rangle_N$ the subgroup generated by sequences of at most N elements of $\mathcal{H}$, we have $\mathcal{S} = \langle \mathcal{G}_n \rangle_{t_{\text{max}}}$.
- **Reward function R :** the MDP formalism allows intermediate rewards, which should, in principle, guide the agent throughout its exploration. A naive choice would be a metric quantifying the closeness between the current state s_t and the target. For instance, one could use the fidelity between $s_t = U_t \cdot U_{\text{targ}}^\dagger$ and the

identity U_{Id} :

$$R(s_t) = \mathcal{F}(U_t, U_{\text{targ}}^\dagger) = \frac{|\text{Tr}[U_t \cdot U_{\text{targ}}^\dagger]|}{2^n}. \quad (4.5)$$

However, intermediate rewards can mislead the agent: successful trajectory may pass through intermediate circuits with orthogonal state to the target, yielding to a low fidelity score and penalizing the agent, despite being on a valid solving path. This behavior is further explored in section 5.1. To avoid such a situation, the environment instead returns a single strong end-reward when the target is reached and zero otherwise:

$$R(s_t) = \mathcal{F}_\delta(U_t, U_{\text{targ}}^\dagger) = \begin{cases} 1 & \text{if } s_t = U_{\text{Id}} \\ 0 & \text{else} \end{cases} \quad (4.6)$$

with $\mathcal{F}_\delta$ the binary fidelity. This sends a strong signal to the agent only when it achieves an exact decomposition, at the cost of a sparse reward: the reward vanishes on every state except the target. Additionally, this leaves the agent without any feedback on its progress along the way, raising a credit-assignment problem: the agent must infer the desired sequence of actions from a single end-signal. As detailed later, this settings fits naturally with AlphaZero agents.

- **Discount factor** γ : in the context of a single-end-reward, the discount factor plays the role of optimizer for the circuit depth. Since the only nonzero reward is delivered at the end of the game, the cumulative reward of a successful game reads

$$G_{0:\infty} = \sum_{t=0}^{t_{\text{end}}} \gamma^t \cdot r_t = \gamma^{t_{\text{end}}}, \quad (4.7)$$

where t_{end} is the number of steps to reach the desired unitary, thus the circuit depth d . Thus, setting $\gamma = 1 - \epsilon$ penalizes long circuits while preserving the reward for finding an exact decomposition, encouraging the agent to find short gate sequence.

At this point, we fully defined quantum circuit synthesis as an RL game : rules, rewards, and game dynamics are set. What remains is the agent itself, that is, the learning algorithm guiding the MDP resolution. This choice cannot be generic, it must match the specificities of quantum circuit synthesis: a large combinatorial state space, with sparse rewards, with long-term strategy capabilities since the first gates of a circuit impact all subsequent choices. This motivates the use of AlphaZero, an agent specifically designed for sparse-reward and long-horizon decision-making problems.




Figure 4.2: Illustration of the quantum circuit synthesis game in a reinforcement learning framework. Targeting the unitary U_{targ} , at each step t the agent picks a gate g from a gate set $\mathcal{G}_n$ with its corresponding qubit connectivity. The gate is appended to the ongoing circuit, the environment computes the updated circuit U_t , resulting in a new state $s_t = U_{\text{targ}}^\dagger \cdot U_t$. The binary fidelity $\mathcal{F}_\delta(U_{\text{targ}}, U_t)$ is returned to the agent as local feedback. The interaction loop stops when the environment reaches the target state s_{targ} or the maximum step count t_{max} .

5 -AlphaZero as a long-term strategy agent

In the previous chapter, we frame quantum circuit synthesis as a decision game in the Markov Decision Problem and Reinforcement Learning frameworks, setting the theoretical background for its resolution. However, we left the question of solvability open. In this chapter, we bridge this gap by introducing a suitable agent for the synthesis game: AlphaZero. We first motivate the need for long-term planning for global coherence across gate sequence. Then, we introduce a dilemma at the heart of tree exploration: the exploration-exploitation dilemma. Finally, we present AlphaZero, a suitable agent for long-term strategy games.

5.1 .Long-term strategies for quantum-circuit synthesis

Quantum circuit synthesis fundamentally requires lookahead decision-making as the impact of each gate choice constrains all subsequent ones. Indeed, the gates selected early in a decomposition must remain consistent with the target unitary throughout construction, requiring global coherence across the whole sequence. Yet, the synthesis process is fundamentally iterative: gates are appended one at a time, with no possibility to revisit earlier choices as the circuit grows. This challenge between sequential actions and global consistency is well known in combinatorial games: a chess player may sacrifice a piece to reach a better position later, or a Rubik's cube solver passes through intermediate configurations that seem more disordered than the initial one. In general, local moves must be chosen for global rewarding outcomes, and this is precisely the regime where greedy strategies fail: by focusing on short-term promising gates, greedy search gets trapped in regions of the decision tree that may never lead to the target unitary. Quantum circuit synthesis therefore calls for a RL agent able to plan several moves ahead, balancing immediate progress with long-term strategic value.

Let us take a simple example: synthesizing the rotation of $\pi/4$ around the X axis, $R_x(\pi/4)$. Using the universal gate set $\{H, T, X\}$, the optimal decomposition of such a rotation is

$$V = e^{i\pi/8} H T H. \quad (5.1)$$

A standard way to quantify the closeness of two unitaries U and V is the fidelity (c.f. equation (4.5)), defined as

$$\mathcal{F}(U, V) = \frac{|\text{Tr}[U \cdot V^\dagger]|}{2^n}, \quad (5.2)$$

with n the number of qubits. As a scalar product on $\text{SU}(2^n)$, the fidelity gives the projection of the two unitaries

$$\mathcal{F}(U, V) = |\cos \theta|, \quad (5.3)$$

where θ is the angle between U and V . In our example, applying the decomposition gate by gate yields the intermediate fidelities

$$\mathcal{F}(H, V) \approx 0.27, \quad \mathcal{F}(HT, V) = 0, \quad \mathcal{F}(HTH, V) = 1. \quad (5.4)$$

In this example the intermediate step HT is exactly orthogonal to the target V , dropping the local reward to 0. This illustrates that intermediate metrics may not provide reliable information on the relevance of the current gate sequence for long-term gain. A greedy strategy, picking at each step the gate that maximizes the fidelity, would have selected a different intermediate gate and missed the optimal decomposition. More fundamentally, this highlights that greedy approaches optimize a continuous variable, while quantum circuit synthesis is inherently a discrete combinatorial problem based on arithmetic invariants that continuous variables cannot capture [139].

An other argument can be drawn from the Cayley graph representation of the gate set. In its general form, the Cayley graph can be defined as below. Given a group $(G, \star)$ with a set of generators S , the Cayley graph $\Gamma = \Gamma(G, S)$ is defined as the graph where nodes v_g are the group elements $g \in G$, and edges encode the action of the generators $s \in S$, thus an edge represents a transition $v_g \rightarrow v_{g \star s}$. In the context of quantum circuit synthesis, the group is the set of reachable unitaries $G \in \text{SU}(2^n)$ and the set of generators is the gate set $S = \mathcal{G}$, thus each node of the graph is a unitary generated by a gate sequence. The Cayley graph captures the structure of the synthesis problem: finding the decomposition of U is actually finding a path from the identity node $v_{U_{\text{id}}}$ to the target unitary node v_U . The structure of this Cayley graph has been characterized by Matsumoto and Amano for 1-qubit circuits. In [140] the authors developed a synthesis algorithm based on the syllables $\{T, HT, SHT\}$. Concretely, they showed that any 1-qubit unitary U generated by the gate set $\langle H, T \rangle$ admits a unique decomposition of the form

$$U = \omega^l \cdot C \cdot \prod_{i=1}^N g_i, \quad (5.5)$$

where $g_i \in \{T, HT, SHT\}$, C a Clifford operator, and $\omega^l = e^{il\frac{\pi}{4}}$ a global phase. Thus, up to a global phase and a Clifford operator C , U is uniquely defined by the syllable word $(g_1, g_2, \dots, g_N)$. This uniqueness shows that the quotient $\langle H, T \rangle / \mathcal{C}$, with $\mathcal{C}$ the Clifford group, is in bijection with the free monoid on $\{T, HT, SHT\}$, meaning the syllable word $(g_1, \dots, g_N)$ admits no non-trivial simplification. Equivalently, the Cayley graph of $\langle H, T \rangle / \mathcal{C}$ is a tree: the path from U_{id} to the target unitary U is unique. Thus, up to Clifford simplification, the first syllables of the decomposition entirely determine whether the target unitary can be reached. Concretely, a wrong early choice cannot be compensated by later gates without backtracking, this motivate the need for an agent able to plan several moves ahead rather than greedily optimizing local fidelity.

The Matsumoto-Amano result holds in a specific framework: single-qubit unitaries

in $U(2)$, fixed gate set $\langle H, T \rangle$, and specific syllables $\{T, HT, SHT\}$. Yet, this case already captures the core tension of quantum circuit synthesis: an iterative gate-by-gate construction must achieve long-term coherence in a discrete combinatorial space. This motivates the use of an agent with long-term strategic capabilities, able to plan several moves ahead rather than greedily optimizing local fidelity. Long-term strategic planning in sparse and large combinatorial space is a well-known regime in reinforcement learning, typically benchmarked on board games such as Go or chess, where local moves reveal their value at the end of the game. This is precisely the purpose of *Monte Carlo Tree Search* based RL agents such as *AlphaZero*. Originally developed for Go, this agent handles long-term strategy particularly well via exploration-exploitation balance and deep look-ahead planning. To gradually introduce AlphaZero, the following section presents the Monte Carlo Tree Search algorithm as a policy provider for long-term planning agents, building from its historical roots in the exploration-exploitation dilemma.

5.2 .Monte Carlo Tree Search

The previous section established the need for an agent able to plan over the long term in a sparse and discrete combinatorial space. This challenge is well known in large combinatorial problems, and has driven the community to develop dedicated algorithms. At the heart of these approaches lies a fundamental tradeoff between exploitation and exploration, which lead to the construction of the Monte Carlo Tree Search (MCTS) algorithm which lies at the core of AlphaZero. The following sections aim to provide theoretical background of MCTS.

5.2.1 .The exploration-exploitation dilemma

Framing reinforcement learning as a decision tree problem, the agent faces a fundamental dilemma during the exploration: choosing whether to *exploit* known high-reward branches, or to *explore* unvisited ones that may yield higher long-term returns. This exploration-exploitation tradeoff first took root in the work of Thompson [141, 142]. Considering two independent Bernoulli processes with unknown success probabilities, Thompson aimed to maximize the expected number of successes over T trials. This minimal sampling problem already illustrates the dilemma, facing the possibility of spending trials to aggregate samples and identify the rewarded Bernoulli process, or committing early to the most promising with the current known information. Motivated by advances in industrial quality control [143], the community started to study the dilemma through a minimalist toy-model: the *multi-armed bandit problem* (MABP) [144]. The setup is simple: a player faces a row of n slot machines, each with an unknown reward distribution of expected value μ_i . With a budget of T rounds, and without any prior information, the player must find a sequence of actions maximizing the cumulative

reward

$$G_T = \sum_{t=1}^T r_t,$$

where r_t is the immediate reward collected at step t . At each round, the player faces the core dilemma: exploit its knowledge by pulling a known reward machine, or explore an underestimate arm with the hope of finding a better reward, at the cost of potentially wasting a round. This elementary setting already captures the core challenge of Reinforcement Learning. In the RL paradigm, at any given state s the agent faces many possible actions, each yielding an unknown potential reward. From state s , and with limited resources, the agent faces the choice of picking an under-explored action, or exploiting a known one with a reliable estimate of the reward. In this context, the RL action-value function $Q^\pi(s, a)$ mirrors the MABP reward distribution μ_i . Thus, RL and MABP share the challenge of selecting the optimal sequence of actions to maximize the cumulative reward, while simultaneously updating the knowledge on the system.

Based on the MABP toy model, the community began to derive formal performance bounds. A key result comes from [145] which introduced the concept of an *Upper Confidence Index*. The authors established that, for a restricted budget of T rounds, the cumulative regret is saturated by $\mathcal{O}(\log T)$ in an asymptotic regime $T \rightarrow \infty$. In other words, the cost of incomplete information, called regret, cannot vanish faster than $\log T$, regardless of the strategy employed. In this reference, the authors only established this result in a theoretical asymptotic regime, leaving open the question for practical strategy with finite-time guarantees.

The authors of [146] address this question by proposing a family of algorithms with finite-time regret bounds. Among these is the *Upper Confidence Bound 1* (UCB1), an algorithm providing a simple and practical arm selection rule for MABP. At each round, the player pulls the arm i^* that maximizes the quantity

$$i^* = \arg \max_i \bar{x}_i + \sqrt{\frac{2 \ln t}{t_i}}, \quad (5.6)$$

where $\bar{x}_i$ is the empirical average reward of the i^{th} arm computed from previous tries, t_i the number of times the i^{th} has been pulled, and t the overall number of rounds¹. The two terms play complementary roles. The first term $\bar{x}_i$ drives exploitation by favoring arms known to provide high rewards. The second term $\sqrt{2 \ln t / t_i}$ drives exploration by being large for arms that have been rarely visited relatively to t . When an arm i is pulled frequently, t_i grows, naturally shifting priority back toward exploitation. This adaptive structure optimizes the regret through experience, without any prior knowledge, and will reappear directly in the action selection rule at the heart of AlphaZero.

While being a practical toy-model for the exploration-exploitation trade-off, the multi-

¹We call T the total budget, thus $t < T$.

armed bandit fundamentally differs from conventional Markov Decision Problem. In the MABP paradigm, the player faces the same set of arms all game long: the game is a stationary process. In an MDP paradigm, an action leads to a new state-node, with a new set of possible actions, and new reward distributions. Additionally, the evaluation of new nodes during the exploration of the tree may retroactively affect older ones: exploring new nodes brings reward information, which updates the value estimates of previously visited nodes. Thus, the reward estimation of the nodes evolves continuously through the tree search.

Focusing on efficient decision-tree exploration for non-stationary process, in [147] Coulom proposed the *Monte Carlo Tree Search* (MCTS), an algorithm based on heuristic decision tree exploration with back-propagation node updates. The core idea is the following: from the environment current state s the agent performs n_{MCTS} rollouts of the decision tree, also called *virtual games*, each ending with a Monte Carlo simulation to an end-leaf nodes which returns a sample of the cumulative reward. This yields in n_{MCTS} samples of the expected cumulative reward of the state s , building an estimation of the action-value and state-value function Q and V as defined in appendix D.11. As we will detail in later sections, each rollout consists of four steps: selection, expansion, simulation, and back-propagation. From the resulting statistics, the algorithm picks an action via a majority vote, resulting to a new state $s \rightarrow s'$, and repeating the process until the end of the game.

From a certain perspective, a MDP can be seen as a succession of MABD: at each node s the player faces a local instance, the action-value function $Q^\pi(s, a)$ playing the role of the expected arm reward μ_i . This is precisely what the authors of [148] noticed, naming this framework as Bandit Based Monte-Carlo Planning, where they used the work from the MABP literature, benefiting from formal regret bounds for decision-tree exploration strategy. More precisely, they proposed to use UCB1 [146] of equation (5.6) to build a MCTS-based algorithm where the selection step is driven by the UCB1 rule. Named *Upper Confidence Bound applied to Tree* (UCT), the decision tree is explored via MCTS while the exploration-exploitation trade-off is managed by UCB1. Concretely, at each node s of the tree, the selected action a^* follows

$$a^* = \arg \max_{a \in \mathcal{A}} \hat{Q}(s, a) + c \cdot \sqrt{\frac{\ln N(s)}{N(s, a)}}, \quad (5.7)$$

with $N(s)$ the number of times the current node s has been visited, $N(s, a)$ the number of times the child node from the transition $s \xrightarrow{a} s'$ has been visited, and c a balance constant. Here $\hat{Q}(s, a)$ is an empirical estimate of the action value function from the previous explorations, quantifying how promising the action a is based on the current knowledge. This is a direct application of the UCB1 index presented earlier: in the degenerate case where every action leads back to the same state s , the problem reduces to a MABP and the formula recovers UCB1. The UCT demonstrated its effectiveness in [149, 150], where the

authors implemented it in MoGo, a UCT-powered Go engine winning multiple computer Go tournaments on a 9×9 board. The success of MoGo established UCT as a practical method for long-term strategic planning.

5.2.2 . The Monte Carlo Tree Search Algorithm

Monte Carlo Tree Search is a heuristic search algorithm for decision processes [148]. It iteratively explores a decision tree where nodes and edges respectively represent states and actions. Given an arbitrary state s_{ini} as input, MCTS builds a sampled decision tree rooted at s_{ini} , where valuable regions are unrolled, yielding an estimate of the policy and expected reward. Starting from a single-node tree s_{ini} , the algorithm runs n_{MCTS} rollouts, expanding the tree by one node at each iteration. Each node s stores two quantities: a visit counter $N(s)$, and an empirical estimation of the total cumulative reward $W(s)$. The purpose of the MCTS rollouts is to refine these quantities via tree exploration.

Before detailing the rollout process, let us fix the terminology and notation. We call *sampled tree* the tree output from the MCTS exploration (as a sample of the theoretical full decision tree). The subtree constructed during the MCTS process via the n_{MCTS} rollouts is called the *search tree* (since it is under a search process). In the search tree, each node represents a state s , associated with a set of available actions $\mathcal{A}_s$. Each edge corresponds to an action $a \in \mathcal{A}_s$ leading to a child state s' . Every node stores two quantities updated during the rollouts: a visit counter $N(s)$, and the total cumulative reward $W(s)$. Similarly, each edge (s, a) carries a visit counter $N(s, a)$ and the total cumulative reward $W(s, a)$, corresponding to the statistics of the child node reached by action a from s . Concretely, the resulting policy

$$\hat{\pi}(a | s) = \frac{N(s, a)}{N(s)} \quad (5.8)$$

and the empirical estimation of the action and state-value functions reads

$$\hat{Q}(s, a) = \frac{W(s, a)}{N(s, a)} \quad \hat{V}(s) = \sum_{a \in \mathcal{A}_s} \hat{\pi}(a | s) \hat{Q}(s, a). \quad (5.9)$$

In the scope of this thesis, we use deterministic environments where a state-action pair (s, a) leads to a unique state s' . Thus edge and node statistics coincide, that is²: $N(s, a) = N(s')$ and $W(s, a) = W(s')$.

With notation defined, we can now detail a single MCTS rollout. Each of the n_{MCTS} rollouts can be decomposed into four steps: selection, expansion, simulation, and back-propagation. Let us assume we are in the middle of an MCTS instance, more precisely at the n^{th} iteration with $1 < n < n_{\text{MCTS}}$, meaning the n search tree has n unrolled nodes. We call s_{ini} the state from which the MCTS instance is called.

- **Selection Step:** starting from the root of the search tree s_{ini} , the nodes of the

²In the general formulation of MCTS, statistics are attached to edges rather than nodes to account for probabilistic transitions.

partially unrolled search tree are successively selected until a leaf is reached (a node that has never been visited by a previous rollout), following the exploration-exploitation framework of [148]. The constructed path

$$\sigma = \{s_0 = s_{\text{ini}}, s_1, \dots, s_k = s_{\text{leaf}}\} \quad (5.10)$$

is built by consecutively choosing the action a^* that maximizes the UCT criterion

$$a^* = \arg \max_{a \in \mathcal{A}_f} \hat{Q}(s, a) + c \cdot \sqrt{\frac{\ln N(s)}{N(s, a)}} \quad (5.11)$$

where $N(s)$ is the number of times node s has been visited, $N(s, a)$ the number of times the child node reached by action a from s has been visited, $\hat{Q}(s, a)$ the empirical estimation of the cumulative reward collected through that child node, and c a balance constant. This step drives the exploration through known nodes of the tree until an unexplored node, i.e. a leaf, is reached.

- **Expansion Step:** once a leaf s_{leaf} is reached, a new child node s_{new} is created, expanding the search tree by one node.
- **Simulation Step:** a random game is played from this expanded node until a terminal state is reached. Calling σ_{rand} the path of the random game $s_{\text{new}} \rightarrow s_{\text{end}}$, we compute the cumulative reward

$$G_{\sigma_{\text{rand}}} = \sum_{s' \in \sigma_{\text{rand}}} R(s'). \quad (5.12)$$

- **Back-propagation Step:** finally, the node statistics along the selection path σ are updated. Each visited node $s \in \sigma$ collects a cumulative reward

$$G_{\sigma_s + \sigma_{\text{rand}}} = \sum_{s' \in \sigma_s + \sigma_{\text{rand}}} R(s'), \quad (5.13)$$

where $\sigma_s \subset \sigma$ denotes the sub-trajectory starting from s , $\sigma_s + \sigma_{\text{rand}}$ the concatenated path, and $r(s')$ is the reward associated with the s' transition state. For each visited state $s \in \sigma$, the empirical estimation of the total cumulative reward $W(s)$ is then updated via

$$W(s) \leftarrow W(s) + G_{\sigma_s + \sigma_{\text{rand}}} \quad (5.14)$$

and the visit counter is incremented by one

$$N(s) \leftarrow N(s) + 1. \quad (5.15)$$

Let us develop some intuition of the MCTS algorithm. The selection and expansion steps exploit the exploration-exploitation trade-off of UCB1 to reach a promising unvisited




Figure 5.1: Illustration of the MCTS iterative process. Assuming we are in the middle of an MCTS instance, the tree is first traversed until a leaf is reached (*selection*). From this leaf, a new child is expanded, growing the sampled tree by one node (*expansion*). A random game is then played out from this new child (*simulation*), and the selected nodes are updated according to the final reward of the simulated game (*back-propagation*).

region of the tree. The random simulation provides a sample of the cumulative reward attainable from the leaf. The back-propagation then refines the value estimates of the visited nodes by incorporating the outcome of the simulation, improving subsequent predictions. Thus, after n_{MCTS} rollouts, the algorithm returns a partially expanded sampled tree, where each visited node carries empirical statistics $W(s)$ and $N(s)$. The root statistics $W(s_{\text{ini}})$ and $N(s_{\text{ini}}, a)$ are the most reliable, as they aggregate information from every rollout. For better understanding, figure 5.1 gives a graphical representation of the MCTS algorithm, and algorithm 1 a formal pseudo-code implementation, using the functions

- $\text{InitTree}(s_{\text{ini}})$: initializes the search tree as a single root node s_{ini} with $N(s_{\text{ini}}) = 0$ and $W(s_{\text{ini}}) = 0$.
- $\text{Environment}(a, s)$: applies action a to state s and returns the resulting child state s' .
- $\text{Expand}(\text{tree}, s_{\text{leaf}})$: creates a new child node s_{new} from the leaf s_{leaf} , growing the tree by one node.
- $\text{RandomAction}(\mathcal{A}_{s'})$: draws an action uniformly at random among the available actions of s' .
- $r(s')$: reward associated with the transition reaching state s' .

Algorithm 1 Monte Carlo Tree Search

Require: root state s_{ini} , number of rollouts n_{MCTS}

```
1:  $tree \leftarrow \text{InitTree}(s_{\text{ini}})$ 
2: for  $i \leftarrow 1$  to  $n_{\text{MCTS}}$  do
  — Selection —
3:    $s \leftarrow s_{\text{ini}}$ 
4:    $\sigma \leftarrow \{s_{\text{ini}}\}$ 
5:   while  $s$  is not a leaf do
6:      $a^* \leftarrow \arg \max_{a \in \mathcal{A}_s} \hat{Q}(s, a) + c \cdot \sqrt{\ln N(s) / N(s, a)}$ 
7:      $s \leftarrow \text{Environment}(a^*, s)$ 
8:      $\sigma \leftarrow \sigma \cup \{s\}$ 
9:   end while
10:   $s_{\text{leaf}} \leftarrow s$ 
  — Expansion —
11:   $s_{\text{new}} \leftarrow \text{Expand}(tree, s_{\text{leaf}})$ 
  — Simulation —
12:   $s' \leftarrow s_{\text{new}}$ 
13:   $G \leftarrow 0$ 
14:  while  $s' \neq s_{\text{end}}$  do
15:     $a \leftarrow \text{RandomAction}(\mathcal{A}_{s'})$ 
16:     $s' \leftarrow \text{Environment}(a, s')$ 
17:     $G \leftarrow G + R(s')$ 
18:  end while
  — Back-propagation —
19:  for each  $s \in \sigma$  do
20:     $G_s \leftarrow \sum_{s'' \in \sigma_s} R(s'') + G$ 
21:     $W(s) \leftarrow W(s) + G_s$ 
22:     $N(s) \leftarrow N(s) + 1$ 
23:  end for
24: end for
  — Estimation from the root statistics —
25: for each  $a \in \mathcal{A}_{s_{\text{ini}}}$  do
26:    $\hat{\pi}(a \mid s_{\text{ini}}) \leftarrow N(s_{\text{ini}}, a) / N(s_{\text{ini}})$ 
27:    $\hat{Q}(s_{\text{ini}}, a) \leftarrow W(s_{\text{ini}}, a) / N(s_{\text{ini}}, a)$ 
28: end for
29: return  $\hat{\pi}, \hat{Q}$ 
```

5.2.3 .Monte Carlo Tree Search as a Reinforcement Learning agent

In a reinforcement learning framework, MCTS can naturally be used as a policy provider for the agent: at each step of the game, the agent calls an MCTS instance from its current state, explores the decision tree, collects statistics from n_{MCTS} simulated games, and computes an empirical policy as shown in equation (5.8). These n_{MCTS} are run in a sandboxed copy of the environment, and are referred to as *virtual games* since they do not affect the real game state. This places MCTS-based agents in the class of model-based algorithms: the agent has access to an exact simulation of the environment, allowing strategy refining via virtual games studies. Much like a chess player mentally exploring several moves ahead, the agent simulates future sequences of play and selects the most promising action in the real environment. See figure 5.2 for illustration.

Concretely, MCTS-based RL agents are structured in two nested loops. The outer loop corresponds to the game itself: for every node s , the agent picks an action a according to a policy $\pi(a | s)$, leading to a new node s' in the game tree. The inner loop corresponds to the MCTS procedure: before picking an action the agent plays n_{MCTS} virtual games from its current node s , getting a sampled tree from which a policy $\hat{\pi}(\cdot | s)$ can be estimated.

As an example, let us consider a time step t where the environment is in state s_t . Facing a multitude of choices, the agent calls an MCTS instance with root node s_t . Playing n_{MCTS} virtual games with $s_{\text{ini}} = s_t$ as the initial state, the agent gets data from the MCTS tree exploration. From these this data, the agent is able to build a policy $\hat{\pi}(\cdot | s_t)$, as it has been exposed in equation (5.8). For instance with the Boltzmann distribution

$$\hat{\pi}(a | s_t) = \begin{cases} \frac{N(s_t, a)^{1/\tau}}{\sum_{a'} N(s_t, a')^{1/\tau}} & \tau \in]0, 1] \\ 1 \text{ if } a = a^* \text{ else } 0 & \tau = 0 \end{cases} \quad (5.16)$$

with $N(\cdot)$ the visit counter from the MCTS instance, τ the temperature which balances the distribution, and a^* the most promising move

$$a^* = \arg \max_{a \in \mathcal{A}_s} N(s, a). \quad (5.17)$$

The temperature parameter τ directly refers to the Boltzmann distribution:

$$\hat{\pi}(a | s_t) = \frac{1}{Z} \exp\left(\frac{1}{\tau} \ln N(a, s_t)\right), \quad (5.18)$$

which balances the sharpness of the distribution.

By planning several moves ahead with end-rolling virtual games, MCTS agents excel at long-term strategic decisions, with the asymptotic guarantee that $\hat{V}(s_t), \hat{Q}(s_t, \cdot) \rightarrow V^{\pi^*}(s_t), Q^{\pi^*}(s_t, \cdot)$ for $n_{\text{MCTS}} \rightarrow \infty$ as demonstrated in [148]. This strength, however, comes at a significant computational cost: each move requires n_{MCTS} calls to the environment simulator, which becomes problematic when simulation is expensive. A

more fundamental limitation is the absence of a persistent memory across games. The MCTS is purely local, rebuilding its search tree from scratch every new games, without accumulating a global understanding of the environment across the games played. This lack of persistent knowledge motivates the integration of MCTS with heuristic learner such as neural networks, which lies at the heart of the AlphaZero framework.

5.3 .AlphaZero

Unlike conventional MCTS which explores the decision tree on the fly without any prior knowledge, AlphaZero exploits an intensive *training phase*, tuning the agent to the target problem before deployment. The core idea is to integrate a neural network trained from *self-play* games, embedding knowledge about the game, helping guide the MCTS. The agent benefits from the neural network integration in two ways. Firstly, the representation capabilities of the neural network make the tree search smarter. Secondly, the knowledge accumulated in the network reduces the number of iterations needed for accurate MCTS statistics, allowing deeper exploration within the same computational budget. Concretely, AlphaZero involves a network with two heads³: a policy head predicting the policy π , and a value head estimating V^π . In this section, we use the term policy network and value network to refer to the policy head and value head of the neural network, respectively. This section focuses on the integration of the neural networks in the MCTS pipeline, and, specifically, how the neural network guides the tree exploration for better policy evaluation.

5.3.1 .Guiding MCTS with a Neural Network

In standard UCT, the exploration bonus is managed with a single constant C_{UCT} , treating all actions uniformly regardless of their potential. Moreover, when a node is visited for the first time, the actions are selected randomly. [151] proposed an adaptive trade-off by weighting the exploration term with a prior $u(s, a)$, enabling guided exploration during the virtual games. This UCT variant, called PUCT for Predictor UCT, reads

$$\text{PUCT}(s, a) = \hat{Q}(s, a) + C_{PUCT} u(s, a) \cdot \frac{\sqrt{N(s)}}{N(s, a) + 1}. \quad (5.19)$$

AlphaZero adopts this criterion for its tree search. Rather than hand-crafting the prior $u(s, a)$, the policy network provides a policy, $\tilde{\pi}(a | s)$. Formally, the AlphaZero version of the PUCT is

$$\text{PUCT}(s, a) = \hat{Q}(s, a) + C_{PUCT} \tilde{\pi}(a | s) \cdot \frac{\sqrt{N(s)}}{N(s, a) + 1}, \quad (5.20)$$

³The original AlphaGo [115] uses two separate networks, while AlphaZero implementation [116] uses a single network with two heads, maximizing the shared representation of the policy and value.




Figure 5.2: Illustration of a reinforcement learning MCTS agent. At each time step t , the agent runs an MCTS instance from the current environment state s_t . This MCTS instance plays n_{MCTS} virtual games, using $s_{\text{ini}} = s_t$ as the root of the tree. A policy is then built from the sampled tree, from which the agent picks an action.

where $\tilde{\pi}(a | s)$ is the policy from the neural network, $\hat{Q}(s, a)$ is an estimation of the action-value function from the tree search, $N(s)$ the MCTS visited counter, and C_{PUCT} is the balancing constant. This heuristic method guides the exploration during the virtual games. It also removes the need for game-specific priors, allowing the network to capture structure that would be missed by a hand-crafted heuristic.

Moreover, by replacing the simulation step with a single evaluation of the network, the MCTS efficiency is drastically improved. Trained on a multitude of self-play games, the value network directly estimates the expected reward from a leaf node $\tilde{G}$ (replacing $G_{\sigma_{\text{rand}}}$), bypassing the need for a full simulation. While calling the neural network may be computationally more expensive than a rollout, AlphaZero compensates for this through an accurate estimation of the leaf-node value. Indeed, a single evaluation of the value network is worth an enormous number of standard guided simulation, offering a better trade-off between accuracy and computational cost. More precisely, in the original AlphaGo [115], evaluating a single move takes on the order of microseconds, while a neural network call requires a few milliseconds. Thus, for a game requiring roughly 100 moves to end, a single network call is about 10 times slower than simulating an entire game. In other words, one can simulate approximately 10 full games per neural network call. However, the authors showed that the value network requires way fewer evaluations to match the accuracy of a policy-guided simulation, reducing the overall computational cost per virtual game by a factor of 15000. This allows a deeper exploration and therefore smarter strategy planning for a limited computational budget.

5.3.2 .AlphaZero policy

During a game, AlphaZero follows the standard MCTS-based agent structure: an outer loop corresponding to the real game steps, and an inner loop running n_{MCTS} virtual games per step. The key difference with standard MCTS agent lies in the rollouts: the selection step is now guided with the policy network, and the simulation replaced by a value network call.

At each step, AlphaZero agent calls an MCTS instance from its current state s , which plays n_{MCTS} virtual games in a sandboxed environment. Then, based on the sampled tree statistics, the agent builds a policy $\hat{\pi}(\cdot | s)$ and picks an action. In the original AlphaZero paper, the policy is built via a Boltzmann distribution in the zero-temperature limit $\tau \rightarrow 0$. Thus, the most visited child-node

$$a^* = \arg \max_{a \in \mathcal{A}_s} N(s, a) \quad (5.21)$$

is picked. This greedy selection exploits the knowledge accumulated by the networks during the training phase. Unlike standard MCTS, which discovers the game tree on the fly, AlphaZero benefits from a persistent knowledge of the game acquired during a training phase prior to deployment, which we describe next.

5.3.3 .Training AlphaZero

The training pipeline of AlphaZero alternates between two phases: self-play and network update. During the first phase, the agent plays training games, called *self-play games*, with the guided version of MCTS described above, accumulating sampled tree and statistics along the way. Then, these statistics are used as training data for updating the network: given a state, the network learns to directly predict the policy and value that the tree search would have produced. Over successive iterations, the network produces more accurate estimation, refining the MCTS guidance while progressively focusing the exploration on promising regions of the tree.

More formally, the training pipeline is composed of three nested loops. The outermost loop is the network update: for each update, n_{Games} self-play games are played, aggregating data over these games to feed the next network update. The second loop runs these n_{Games} self-play games one by one. The innermost loop runs, at each step of a self-play game, an MCTS instance of n_{MCTS} virtual games. In the standard machine learning vocabulary, an iteration of the outer loop is called an *epoch*. Every epoch the updated agent competes with the previous one on few games, benchmarking the impact of the network update on the agent smartness, and the winner is kept for the next epoch. After n_{Epoch} , the training ends and the networks are ready to be used for real games. For clarity, the figure 5.3 illustrates the pipeline, and algorithm 2 gives a formal description.

The choice of the network architecture has a non-negligible impact on the agent's learning performance. This must actually match with the environment representation and the game mechanics. For instance, a game whose observations are captured as images naturally fits a convolutional network, which efficiently extracts local spatial patterns. Conversely, a linear observation with long-range correlations is better suited to a transformer architecture, which captures contextual dependencies attention layers. A more detailed discussion of the network architecture is given in section 6.4.

Usually, the overall learning strategy adopted follows an exploration-oriented setting during the training, as opposed to an exploitation-oriented setting for real games. During the training phase, we limit the number of virtual games per move n_{MCTS} , preventing the agent from deep exploration. Additionally, stochasticity is added to the exploration via a non-zero temperature $\tau > 0$ and Dirichlet noise at the initial step, the details of which are not discussed here. This setting boosts the exploration capabilities of the agent and the resulting training data spans diverse regions of the decision tree, exposing the networks to a broad range of situations for general learning.

Conversely, the agent switches to an exploitation setting during real games, leveraging the capabilities of the trained network. With a large number of simulation per moves and a greedy behavior with a temperature $\tau = 0$ (as depicted in equation (5.16)), the agent focuses on high-reward regions of the tree with a deep exploration. This setting exploits the knowledge accumulated during the training phase. For instance, in the AlphaZero original paper [116], the authors limited the number of virtual games by $n_{\text{MCTS}} = 800$ with

strong Dirichlet noise and high policy temperature, while the number of virtual games reached hundreds of thousands during real game, without any stochasticity.

Algorithm 2 Training AlphaZero

```
1: for  $i \leftarrow 1$  to  $n_{epoch}$  do
2:    $data \leftarrow ResetData()$ 
3:   for  $j \leftarrow 1$  to  $n_{game}$  do
4:      $s \leftarrow s_{start}$ 
5:     while  $s \neq s_{stop}$  do
6:        $\hat{\pi}, \hat{Q} \leftarrow MCTS(s, n_{MCTS})$ 
7:        $data \leftarrow AddStorage(\hat{\pi}, \hat{Q})$ 
8:        $a \leftarrow Pick(\hat{\pi})$ 
9:        $s \leftarrow Environment(a, s)$ 
10:    end while
11:  end for
12:   $UpdateNeuralNetwork(data)$ 
13:   $Competition()$ 
14: end for
```

- $MCTS(s, n_{MCTS})$: call of a guided MCTS instance. This function runs the networks-guided MCTS with n_{MCTS} iterations, from the current state s . Returns an estimation of the policy $\hat{\pi}$ and value function $\hat{V}$ and $\hat{Q}$ from the statistics.
- $AddStorage(\hat{\pi}, \hat{Q})$: add the guided MCTS estimated policy $\hat{\pi}$ and value function $\hat{Q}$ to the storage.
- $Pick(\hat{\pi})$: selects an action according to the policy $\hat{\pi}$.
- $Environment(a, s)$: plays the chosen move by calling the environment, update the state of the environment, and returns a reward and an observation.
- $UpdateNeuralNetwork(data)$: update the weights of the neural network with $data$.
- $Competition()$: run competition between the current and updated neural networks, the best one is kept.
- $ResetData()$: reset the data for a new epoch.




Figure 5.3: Framework of the training phase of AlphaZero.

6 -Results and Perspectives

The previous chapter introduced AlphaZero as an enhanced version of Monte Carlo Tree Search where the tree exploration is guided by a neural network. With this theoretical basis, we now apply AlphaZero to quantum circuit synthesis using the game framework developed in chapter 4. In this chapter, we first describe the learning strategy, detailing our implementation and parametrization. We then present the results, before proposing perspectives for further development.

6.1 .Learning pipeline

As seen in the previous sections, training an AlphaZero agent requires a computationally heavy phase where the agent plays self-play games, aggregates MCTS data, and updates the policy and value networks. In this pipeline, the agent is tasked with synthesizing a set of *training unitaries* $\{U_i^{\text{train}}\}$, each provided with a known decomposition. From a certain point of view, this set plays the role of a training dataset: it guides the learning process and shapes the final performance of the agent. Although AlphaZero is part of unsupervised learning, we will refer as *training set* the set of unitaries on which the agent learns to synthesize. The choice of training unitaries therefore deserves careful design. In this work, they are built from randomly sampled circuits of predefined depth. Concretely, each circuit is iteratively constructed from the agent's gate set $\mathcal{G}_n$, with each gate drawn from a uniform distribution. This procedure guarantees the existence of an exact decomposition of each training unitary U_i^{train} within depth d_i . Targeting short-depth solutions, we leverage this knowledge by constraining the agent to find a decomposition with at most d_i gates. Furthermore, to prevent trivial circuits that could mislead the learning process, gate redundancies are filtered out using the action mask detailed in section 5.1.

Curriculum learning

Even when a solution is guaranteed to exist, learning from random circuits remains challenging: rewards are sparse and the space of possible circuits grows exponentially with depth. We therefore adopt a structured training strategy, referred to as *curriculum learning*, which orchestrates how the agent is exposed to the training data [152]. To prevent the agent from behaving as a random walker during the first epochs, the difficulty of the game increased progressively as learning advances. Concretely, the depth of each random training circuit U_i^{train} is drawn from a Gaussian distribution with mean μ and variance σ . At every k epoch, the Gaussian center μ is incremented by one. Formally,

$$d_i \sim \mathcal{G}(\mu_i, \sigma), \quad (6.1)$$

where $\mathcal{G}(\cdot, \cdot)$ denotes the Gaussian probability distribution and the mean is updated as

$$\mu_i = \mu_{\text{ini}} + \left\lfloor \frac{i}{k \cdot n_{\text{game}}} \right\rfloor, \quad (6.2)$$

with μ_{ini} the initial mean depth, from which the curriculum starts. This significantly reduces the computational cost of training. Starting from low-depth circuits initially increases the number of won games which the agent learns from, since MCTS alone may already find solutions. As the average target depth grows, the policy is progressively refined to synthesize non-trivial target unitaries, mitigating reward sparsity. See figure 6.1.

Scheduled policy temperature

To further improve the learning process, we use a scheduled policy temperature which encourages exploration in the first steps of games and exploitation for the last steps. As seen previously, higher temperature makes the action selection more stochastic, following the probability distribution of the MCTS policy for $\tau = 1$, while a lower temperature makes it more deterministic, selecting deterministically the action with the highest probability for $\tau = 0$. This setting boosts the exploration capabilities of the agent and the resulting training data spans diverse regions of the decision tree, exposing the networks to a broad range of situations for general learning. Concretely, during the self-play games the agent policy follows

$$\pi(a | s_t) = \begin{cases} \frac{N(s_t, a)^{1/\tau}}{\sum_{a'} N(s_t, a')^{1/\tau}} & \tau \in]0, 1] \\ 1 \text{ if } a = a^* \text{ else } 0 & \tau = 0 \end{cases} \quad (6.3)$$

with $N(\cdot)$ the visit counter from the MCTS instance, τ the temperature which balance the distribution, and a^* the most promising move

$$a^* = \arg \max_{a \in \mathcal{A}_s} N(s, a). \quad (6.4)$$

In practice, we set a temperature of $\tau = 1$ for the first 5 steps before decreasing it to $\tau = 0.6$ for the next 10 actions and down to $\tau = 0.4$ for the remaining ones.

6.2 .Numerical implementation

Game parameters

Before describing the training procedure, let us fix the context variables of the environment on which AlphaZero is trained. As discussed previously, the reinforcement learning environment is fully defined by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, R, \gamma \rangle$, which has been set in chapter 4. However, *game-parameters*, such as the number of qubits n or qubit connectivity, directly impacts the instances of the tuple: the former impacts the dimension of the




Figure 6.1: Illustration of the curriculum learning strategy. A set of training unitaries $\{U_i^{\text{train}}\}$ is generated from random circuits, each admitting a depth- d_i decomposition by construction. The target depth is progressively increased during training, exposing the agent to increasingly difficult synthesis tasks.

state space $\mathcal{S}$, while the latter defines the action space $\mathcal{A}$. Therefore, an AlphaZero agent only applies to synthesis problem sharing the same set of metaparameters. Otherwise, each combination of these game parameters yields a distinct instance of the environment, on which the agent must be trained from scratch. In this manuscript, we explore four instances that can be divided into two main families: all-to-all connectivities and $n+1$ architecture. The *all-to-all connectivity architectures* correspond to the conventional architecture where the n data qubits carry gates and all qubits can talk to each other. We use the gateset $\mathcal{G} = \{H, S, T, \text{CNOT}\}$ for this architecture. Inspired by [118–120], the *dynamic circuit architectures* refer to a $n+1$ architecture with n data qubits and one additional ancilla qubit. In this case, the agent is constrained to apply single-qubit gates on the ancilla only, while CNOTs are shared between the ancilla and the data qubits. The gateset used for this architecture is $\mathcal{G} = \{H, T, \text{CNOT}\}$. Note that this connectivity maintains universality, swap operations can be implemented using CNOT operations. For each of those architectures, we train our networks for $n = 2$ and $n = 3$ qubits circuits, leading to a total of four agents with a dedicated state space $\mathcal{S}$ and action space $\mathcal{A}$.

Agent hyper-parameters

The agents are trained using the curriculum learning strategy described above, with an average depth ranging from 5 to 30. For each target depth d_i , we run $n_{\text{epoch}} = 5$ epochs, each consisting of $n_{\text{games}} = 2048$ games, resulting in $n_{\text{steps}} = 2048 \times d_i \times n_{\text{MCTS}}$ MCTS data for each epoch, with a batch size of 64. At the end of each epoch, the neural network weights are updated using Adam for the gradient descent with a learning rate of 10^{-3} [153]. For stability, we also include an L_2 -regularization term with a weight of 10^{-4} . The updated network π_{new} is then compared to the best-performing one π_{best} over 100 games. π_{best} is replaced by π_{new} if the latter synthesizes at least 5 more target unitaries

than the former. After the 5 epochs completed, the target depth is incremented by one $d_{i+1} = d_i + 1$. The table table 6.1 summarizes the numerical numbers.

Given the complexity of the unitary synthesis task, we opt for a two-head deep residual network architecture, or ResNet [154]. ResNets are known to be efficiently trainable and stable even with many deep layers, thanks to their residual connections. They are also compatible with fine-tuning, as required by the curriculum learning strategy where the network weights need to be progressively readjusted with the increasing difficulty. Our ResNet takes as input the game state U_t split into two flattened arrays, one for the real part and one for the imaginary part. The architecture consists of a main torso of 5 residual blocks, each with 64 channels for feature extraction and a kernel size of 3×3 . The network then splits into two distinct heads for policy and value predictions. The policy head first applies a convolutional layer with 32 channels and a kernel size of 1×1 , followed by a dense layer with a softmax activation function, producing a probability distribution over all actions $a \in \mathcal{A}$. Similarly, the value head applies a convolutional layer with 8 channels and a kernel size of 1×1 , followed by a dense layer with a tanh activation function, producing a scalar value in $[-1, 1]$. To stabilize and accelerate training, batch normalization layers are added after each convolutional operation.

Hardware implementation

To exploit parallel GPU acceleration on matrix multiplications, the game environment has been fully implemented in JAX [155]. We rely on the PGX game framework, a standard environment library supported by most machine learning toolkits [156]. Our implementation is publicly available on GitHub¹.

On the agent side, we exploited the parallel capabilities of MCTS through a vectorized implementation of AlphaZero in JAX. This enables the parallel execution of multiple virtual games on GPU during self-play, significantly accelerating the training phase. Moreover, keeping both environment and agent on the GPU keeps all data in VRAM during the training, avoiding the PCIe transfers that would bottleneck the learning loop. Our implementation is adapted from turbozero [git-valcarce, 157]. For training, the AlphaZero agent is parameterized with a simulation budget of $n_{\text{MCTS}} = 200$ per action, balancing performance and resource usage.

Regarding the hardware resources, we train the agents on two NVIDIA Titan Xp GPUs. This setup results in training time of ≈ 1 day 19 h for the architecture with 2 data qubits and 1 ancilla with a restricted qubit connectivity ($2 + 1$ architecture), ≈ 4 day 7 h for 3 data qubits and 1 ancilla ($3 + 1$) qubits and a restricted qubit connectivity, and ≈ 4 day 10 h for 3 qubits without ancilla and a all-to-all connectivity. These 3 cases are discussed further in the next section. We study the carbon impact of the training process in the next section.

6.3 .Results

¹https://github.com/xvalcarce/quantum_compilation

Quantity	Per iteration	Total
Target depth	$d_i = 5 \rightarrow 30$	$30 - 5 + 1 = 26$
Epochs	$n_{epoch} = 5$	$5 \times 26 = 130$
Real games	$n_{games} = 2048$	$2048 \times 130 \approx 2.7 \times 10^5$
Steps	$n_{steps} = d_i$	$(\sum_{d=5}^{30} d) \times (n_{epoch} \cdot n_{games}) \approx 4.7 \times 10^6$
MCTS virtual games	$n_{MCTS} = 200$	$200 \times 4.7 \times 10^6 \approx 9.3 \times 10^8$

Table 6.1: Training computational cost. Each line shows the number of operations per iteration of the parent loop, and the cumulative total over the full training. The total number of MCTS simulations reaches the order of 10^9 , illustrating the resource intensity of the training phase.

6.3.1 .Quantitative benchmark

To evaluate the performance of our trained agents, we propose a benchmark based on randomly sampled target unitaries. For each agent, we choose these unitaries from randomly constructed quantum circuits using the same gate set and qubit connectivity the agent has been trained on. A synthesis is considered successful if it implements the target unitary and if the designed circuit depth is at most the depth of the circuit used to produce the target unitary. For this benchmark, we set the simulation budget to $n_{MCTS} = 400$, balancing agent performance with runtime efficiency to enable exploration across a broad range of target unitaries in a reasonable time. Figure 6.2 illustrates the trade-off between agent performance and runtimes for different values of n_{MCTS} .

For each test unitary, we first attempt synthesis using a deterministic policy ($\tau = 0$). If no circuits were found, we switch to a stochastic policy ($\tau = 1$) and perform 10 independent runs, for an average total runtime of approximately 10 s, as it may help the agent to escape from local minima in early gates. Synthesis is considered successful if any of these runs produces a circuit implementing the target unitary.

To benchmark our trained agents we also run a Monte Carlo tree search without neural network guidance. We keep the simulation budget to be the same. Similarly, we first run a deterministic MCTS policy with $\tau = 0$, and, if that fails, we run 10 runs with a temperature $\tau = 1$.

We conduct this benchmark for agents trained for the 2+1 and 3+1 qubits architecture. The results are shown in figure 6.3. Our AlphaZero agents outperformed the MCTS baseline in every case, allowing to synthesize unitaries sampled from high depth circuits. Comparing the 2 + 1 and 3 + 1 architectures, we see the benefits of AlphaZero agents against MCTS agents as the number of qubit increases. Note that the average depth of the synthesized circuits is always lower than the quantum circuit used to produce the target unitaries, hence either demonstrating simplification capabilities of our agents, or circuits with trivial simplification from the random generation.

6.3.2 .Qualitative benchmark




Figure 6.2: Unitary synthesis success for 100 randomly sampled circuits of depth 10 on a 3+1 architecture, using AlphaZero agents with varying simulation budget, n_{MCTS} , and a fix temperature $\tau = 0$.




Figure 6.3: Performance of trained AlphaZero agents in synthesizing unitaries generated from random circuits of varying depth into circuits of similar or lower depth. For each depth d , we attempt to synthesize 100 random target unitaries. An attempt is considered successful if the synthesized circuit is of at most depth d . For each target unitary, we first run the AlphaZero agent with $\tau = 0$, if synthesis fails, we give it 10 runs at $\tau = 1$. As a baseline, we run an MCTS with the same two regimes.


As quantitative benchmarks might not be relevant for practical use cases, we compile unitaries that are fundamental elements of quantum operations. All the unitaries explored here are clearly defined in Ref. [158]. The circuits successfully synthesized by our agents are available in table 6.2.

2-qubits unitaries: For 2-qubits unitaries, we aim to compile Control-S (CS), Control-T (CT), Control-H (CH), Control- $\sqrt{\text{NOT}}$ (CV), and iSWAP gates. Except for CT, our agents successfully synthesized all of these gates, both in the all-to-all connectivity and clean ancilla architectures. These compilations have been done using the deterministic regime, $\tau = 0$, and therefore have a runtime of a few seconds. The produced circuits match known optimal decompositions both in terms of T-count and circuit depth. The CT unitary have been found only in the clean-ancilla regime. This is in accordance with the fact that the CT gate can not be implemented using a Clifford+ T gateset on 2 qubits without the use of an ancilla [118]. The circuit found matches the currently best known implementation proposed in [119].


3-qubits unitaries: For 3-qubits operations, we attempt to synthesize the Toffoli or CCX, CCZ, and Fredkin or Control-SWAP gates. In the all-to-all connectivity architecture, we recover known optimal circuits for the Toffoli and CCZ gates. While synthesizing the CCZ took $3.3s$ using the deterministic $\tau = 0$ regime, finding the Toffoli decomposition took $102.2s$ and took 100 runs using a temperature $\tau = 0.6$. Synthesizing the Fredkin operation took 19 gates instead of the known optimal of 17, and necessitated 1000 runs at $\tau = 0.6$ for a total runtime of $3042s$. Using clean-ancilla architecture, we find an implementation of the Toffoli and CCZ gate using $4T$ -gates in 200s runtimes from 200 runs at $\tau = 0.6$ (see the discussion in the next section for the synthesis of the Toffoli gate). This decomposition is consistent with the $4T$ -gates construction proposed in [118], showing the capabilities of the agent. However, our agent failed to synthesize the Fredkin gate in the given 1000 runs - or 3000s of maximum runtime.

6.3.3 .Human feed-back


Trained agents can be used in a human-machine interaction loop, i.e. from a first circuit draft provided by the agent, it is often possible to simplify and improved the circuit. To illustrate this interaction, we study the compilation of a Toffoli operation using the clean-ancilla architecture. In approximately 200s runtime, our agent provides the quantum circuit given in figure 6.4a. From this circuit, we manually discard the first 7 gates, which are equivalent to the identity channel. These unnecessary operations are likely caused by the high temperature τ used to sample the actions in the first few steps. We then replace the last two consecutive $T^\dagger$ gates with a single $S^\dagger$ operation. Note that these two improvements can be algorithmically addressed using identity checks and basic replacement rules.



(a) Circuit synthesized by the 3+1 agent. The first 7 gates reduce to the identity operation and the last two $T^\dagger$ can be replaced by a $S^\dagger$ gate.



(b) Synthesis of the correction unitary U_C using the 3 qubits all-to-all agent.



(c) Complete implementation of the Toffoli gate on 3 qubits and 1 clean ancilla, using $4T$ gates.

Figure 6.4: Using reinforcement learning to synthesize the Toffoli gate for conditionally clean ancilla architecture.

As the Toffoli implementation proposed in figure 6.4a is conditioned on the ancilla being measured in $|0\rangle$, we provide the correction operation U_C , to apply on the data qubit when the ancilla is measured in the $|1\rangle$ state. Demonstrating the complementarity of our agents, we use two AlphaZero agents trained on 3 qubits using the $\{H, S, T, \text{CNOT}\}$ and $\{X, Z, H, S, T, \text{CNOT}\}$ gatesets, respectively. In a runtime of less than a second, we find the corrections given in figure 6.4b. Together, this leads to a deterministic implementation of the Toffoli gate using 4 T gates as depicted in figure 6.4c. This matches the T -depth proven optimal in [85], which took a runtime of an hour.

Note that we were not able to retrieve a circuit with a similar T -count using Synthetiq, with a Clifford+ T gateset and one clean ancilla given a maximum runtime of 200s.

6.4 .Perspectives

6.4.1 .Limits

As shown in figure 6.3, the AlphaZero agent is able to synthesize random unitaries up to depth 30 across multiple architectures, providing a proof of concept for reinforcement-learning-based quantum compilation. The approach also synthesizes practical gates such as the Toffoli or the controlled-SWAP, including a new implementation of the Toffoli on a 3+1 architecture. However, several limitations of the current pipeline must be discussed.

First, synthesized circuits often contains trivial simplification and identity blocks. For instance, the dynamic Toffoli circuit produced by the 3+1 agent could be manually

simplified by recognizing that the first seven gates reduce to the identity. This highlights a limitation of the agent in optimal-path finding, which may actually originate from the quality of the training dataset. The target unitaries are built from random-circuit sampling, and despite the non-trivial action mask filtering redundancies, the resulting circuits often contain higher-order simplifications that the mask does not detect. The agent is therefore exposed to targets with overestimated effective depth, allowing it to learn longer than necessary paths while still being rewarded for exact synthesis.

While a deterministic policy with temperature $\tau = 0$ enables fast runtimes, it leads to lower success rates than multiple stochastic runs at $\tau > 0$. The quantitative analysis on the 3+1 architecture shows a non-negligible improvement in compilation success for higher temperatures, at the cost of a runtime increase to $\sim 10\text{ s}$ compared to a single deterministic shot. This prevents the agent from being used as an on-the-fly compiler for subroutines of large circuits, as is commonly done in classical compilation. The agent rather serves as a first draft for further improvements, either through a human-machine interaction loop or via downstream automated tooling.

Finally, the main limitation of this work is its lack of scalability to larger circuits. Using unitaries as the state representation leads to impractical size for the data on which the agent is trained. Indeed, a n -qubit environment carries states of size 2^{2n} , which quickly becomes out of hand even with sparse representations of unitaries. This has been noticed during training, where the VRAM of the GPUs saturates rapidly. This calls for more compact representations of the states.

These limitations motivate further development of the framework beyond the proof of concept presented here. The next section explores several directions of improvement, ranging from smarter sampling strategies for the training set to algorithm improvement, with the goal of pushing the agent toward smarter and larger circuits compilation.

6.4.2 .Improvements

Training set: As RL agent performances are strongly affected by the quality of the self-play game, an improvement lies in increasing the quality of the training set. First, we may focus on sampling less trivial training unitaires. The current redundancy mask introduced in section 4.3 only prevents two consecutive gates from canceling each other, it leaves higher-order simplifications such as three or four-gate blocks reducing to an identity. Quantum circuit simplification tools could be leveraged to build a stronger non-trivial training set with a same depth. For instance, [85] uses the ring representation of unitaries to develop formal synthesis via SAT solvers, based on this representation the authors also provide a lower bound on the optimal circuit depth. This bound could serve as a constraint during training circuit generation, filtering trivial circuits which bring small value to the training phase.

A second improvement concerns the non-Clifford content of the sampled circuits, which we aim to increase. Adding T gates to a sampled circuit does not guarantee an

increase in its non-Clifford content: commutation relations may simplify and vanish T gates. A more robust approach relies again on the ring representation of unitaries [159]. In [139], the authors observed that, for one-qubit circuits, the T -count, i.e. the minimal number of T gates required, is directly related to the period of the coefficients' unitary over the ring. Concretely, they introduced the *smallest denominator exponent* $\text{sde}(z, x)$ of a base $x \in \mathbb{Z}[\omega]$ with respect to $z \in \mathbb{Z}[1/\sqrt{2}, i]$, defined as the smallest k such that $zx^k \in \mathbb{Z}[\omega]$, quantifying the non-Cliffordness of single-qubit unitaries. In a subsequent paper [160], the same authors proved that $\text{sde}(\sqrt{2}, z) = \mathcal{T}(U)$, where $\mathcal{T}(U)$ denotes the T -count of the unitary U , $\hat{U} \in \mathcal{J}_1$ its quantum channel representation, and z a coefficient of $\hat{U}$. However, this result only holds for single-qubit unitaries, thus, further work is needed for the development of a lower bounds on the T -count of multi-qubit unitaries. Such bounds could be used to constrain the non-Cliffordness of the training circuits, ensuring the agent is exposed to sampled unitaries with genuinely non-trivial T -content. In this context, a possible training strategy would combine the two: a main curriculum progressively increasing the circuit depth, and an inner curriculum, fine-tuning the agent at each depth stage with controlled non-Clifford content to strengthen its handling of T gates.

Hyper-parameters: While AlphaZero provides a sophisticated framework for balancing exploration and exploitation with long-term strategic capabilities, the performance of the agent lies on a set of hyper-parameters whose optimal values are not intuitive. Bayesian optimization offers a way to address this issue, automatically exploring the hyper-parameter space by modeling the agent's performance as a probabilistic function of its configuration. This approach was used during the development of AlphaGo [115], which tuned the hyper-parameters of its search by Bayesian optimization. A similar tuning strategy could be applied to our agent.

Neural Network architecture: Another improvement lies in a deeper exploration of the neural network architecture, which may lead to a smarter agent. Different network architectures are known to suit different problem structures. For instance, convolutional networks excel at image recognition with the kernel capturing local patterns, while transformers dominate natural language processing using attention layers to capture long-range context across distant words. The choice of architecture should therefore reflect the structure of the data. In our implementation, we use a ResNet coupled with a convolutional input layer. This combination targets local pattern recognition in the unitary while benefiting from the stability and trainability of residual connections. However, since the state is encoded as the Kronecker representation of the unitary, correlations between qubits are spread across the whole matrix, often beyond the scope of the 3×3 kernel. For instance, let us consider a single CNOT acting on a 3-qubit register, with the third qubit

controlling the first. The corresponding Kronecker representation is

$$U_{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (6.5)$$

The Kronecker representation scatters the control and target sequence across the matrix, they cannot be captured simultaneously by a standard 3×3 convolutional kernel. The network sees them as disconnected features and must rely on the depth of the subsequent residual layers to catch the full CNOT pattern. A possible way to deal with long-range correlations of the Kronecker representation would be the use of a transformer architecture. Even though the Kronecker representation spreads correlations across the whole matrix, the attention layers of a transformer could capture them directly, bypassing the need for deep residual stacks.

State representation: As discussed earlier, using Kronecker for state representation may not be suitable for future scaling as it grows exponentially with the number of qubits n , $s \in \text{SU}(2^n)$. A potential alternative would be to switch to a ZX-calculus state representation, for instance using the encoding from the Tsim simulator [161]. Another possible improvement would be to add extra information to the state representation, for instance from a ring representation of unitaries. Building on the previous discussion, a natural candidate is the smallest-denominator exponent over $\mathbb{Z}[\omega]$ [160]. This result does not extend to multi-qubit unitaries, but a similar approach based on a ring representation of unitaries has been used in [83].

Hardware and computing power: On a more practical level, training the agent on recent hardware would directly speed up self-play and allow for a larger virtual game budget per move. We train the agents on two NVIDIA Titan Xp GPUs, with 12 GB of VRAM each, required to store the parallel MCTS rollouts and their associated statistics. However, the Titan Xp is based on the Pascal architecture, a hardware architecture old from several generation which are now standard features for deep learning workloads. In particular, the lack of Tensor Cores and efficient half-precision support makes the Titan Xp orders of magnitude slower than recent consumer GPUs on neural network workloads. To identify the bottleneck of the training computation, let us decompose the cost of a single MCTS rollout.

- **Selection:** the agent traverses the expanded tree until a leaf is reached, evaluating

the PUCT formula equation (5.19) from prior stored in the visited nodes, no network call is performed at this step.

- **Expansion and simulation:** the agent expands a new leaf, performing one unitary multiplication to get the leaf-state, and calling the network once to obtain both the prior policy and the value estimate of the leaf.
- **Back-propagation:** only node statistics are updated, which is essentially free.

A single MCTS rollout therefore costs one network inferences and one unitary multiplication. A 16×16 complex matrix multiplication requires $O(6 \cdot 16^3) = O(10^4)$ FLOPs. A single forward pass of our network costs $O(10^8)$ FLOPs, four orders of magnitude above the unitary multiplication. The dominant cost therefore lies in the network calls. The combination of BF16 half-precision arithmetic and Tensor Cores in recent NVIDIA architectures accelerates neural network inference by orders of magnitude over Pascal. This make each rollout proportionally faster and unlocking broader tree exploration within the same computational budget. Additionally, recent GPUs carry several tens of GB of VRAM with bandwidths reaching 1-3 TB/s, contrasting with Titan XP GPU limited at 500 GB/s, which enables larger networks and larger MCTS batches.

Qubits	Unitary	$n+1, \{H, T, CNOT\}$	all-to-all, $\{H, S, T, CNOT\}$
2	CS		
	CT		No circuit found
	CH		
	CV		
	iSWAP		
3	Toffoli	See figure 6.4c	
	CCZ		
	Fredkin	No circuit found	

Table 6.2: Synthesized quantum circuits for some fundamental $n = 2$ and $n = 3$ qubits unitaries. The $n+1$ column corresponds to clean-ancilla architecture circuits build from a $\{H, T, CNOT\}$ gateset. We manually replaced two consecutive T -gates by a S gate, labeled in red. The **all-to-all** column contains circuits build from a $\{H, S, T, CNOT\}$ gateset and using an all-to-all connectivity.

7 -Conclusion

From a broad perspective, this manuscript explores two protocols situated at opposite ends of the quantum information stack. At the network level, it studies the distribution of fragile entangled states between distant modes sharing a single excitation, a key ingredient for long-range quantum communication. At the circuit level, it proposes an efficient compiler using non-Cliffordness sparsely, as this is a costly resource whose consumption ultimately complicates the classical simulation of fault-tolerant quantum computation. While these settings differ considerably in scale and purpose, they are united by a common objective: the optimization of scarce quantum resources under realistic physical constraints. From the distribution of entanglement across large-scale networks to gate-level circuit compilation, this thesis illustrates how advances in low-level quantum information protocols may help shape the foundations of both future quantum communication and computing architectures.

Large Quantum Networks

This manuscript first provides theoretical support for the QIA consortium, whose objective is to develop the first large-scale networks in Europe. A concrete architecture for entangling remote metropolitan quantum networks through a quantum-repeater backbone is proposed. Within a perturbative framework accounting for up to two-photon contributions, we provide numerical evidence supporting the feasibility of this approach. Specifically, our results indicate that tens of entangled states with fidelities exceeding 90% can be distributed per hour over 500 km. This proposal was selected by QIA's Executive Committee as the architecture to be deployed during the project's second phase, which has recently begun.

Beyond this perturbative analysis, we develop an exact model of an elementary quantum network link by deriving explicit expressions for the heralding probability and the conditional state, fully accounting for all multiphoton contributions. These results enable optimization of the source parameters to achieve high entanglement distribution rates while maintaining high-quality states in future quantum networks. They also provide guidance for ongoing experimental implementations by yielding accurate predictions of local photon detection statistics and coherence measurements. As a concrete example, a collaboration with the De Riedmatten group demonstrates that our model accurately reproduces the observed local photon detection statistics and single-photon interference measurements, while also providing direct estimates of network efficiencies and noise levels.

Finally, we propose a witness that reliably certifies the presence of two-qubit en-

tanglement in an elementary network link while accounting for multiphoton events. Importantly, this witness is well suited to ongoing experimental implementations, as it is robust to losses, and is also relevant for future long-distance demonstrations, since it relies solely on local detection events.

The efforts reported in this thesis have already led to several results. The work carried out within the framework of the European project QIA has resulted in a manuscript published in a high-quality peer-reviewed physics journal. We have also released Python packages to characterize a dual-rail two-link repeater chain up to two-photon contributions [162], as well as an exact model of a DLCZ elementary link providing expected photon detection statistics and relevant quantities to verify its operation in the quantum regime via the proposed entanglement witness [163].

The most immediate perspective of this work is to support the near-term activities of the QIA consortium by first demonstrating how to compensate for the mismatch between the photon waveforms emitted by single trapped ions and those retrieved from rare-earth-doped crystal memories. A subsequent step will be to support the development of two-link quantum repeaters. With the exact theoretical models now in place, their adoption by QIA partners can significantly facilitate such implementations by enabling the optimization of source parameters and the accurate characterization of subtle sources of inefficiency and noise. The ultimate objective within this framework is to map the entanglement generated across a repeater chain onto two ion-based end nodes and certify entanglement between them. Since measurements along arbitrary directions are available in this platform, the main challenge is to acquire sufficient data to establish entanglement with statistically significant confidence. In this context, the development of an entanglement witness that explicitly accounts for finite-size effects could substantially reduce the number of experimental trials required for certification. Such a tool would greatly facilitate the demonstration of entanglement between remote metropolitan quantum networks, representing a major milestone for the field and a result of considerable interest to the quantum networking community.

Additional collaborations with experimental groups would help promote the adoption of the exact theoretical model of an elementary quantum network link developed in this thesis. Beyond the collaboration with the De Riedmatten group discussed previously, an ongoing collaboration with the Pan group at USTC aims to realize a record-length elementary network link. A first objective is to characterize the deployed setup by measuring the photon detection statistics obtained with local non-photon-number-resolving detectors over a range of pump powers. While significant efforts have been made to ensure that the photon-pair sources operate in a single spatial and spectral mode, extending our exact model to sources exhibiting multimode emission would be valuable for quantifying the impact of residual multimode effects on the network performance. The final objective is to characterize and certify entanglement within this architecture following the methodology proposed in this thesis. Here again, a key challenge is to account for finite-size effects

and, ideally, to develop methods that reduce the number of experimental trials required for reliable entanglement certification. Such advances would considerably facilitate the demonstration of entanglement over unprecedented distances and further strengthen the practical relevance of the theoretical framework developed in this work.

Quantum Circuit Synthesis

Moving away from quantum communication, the second part of this manuscript turns to circuit-level quantum compilation, a key ingredient of fault-tolerant quantum computing. We study exact unitary synthesis within the Clifford+ T gate set, focusing on decompositions of quantum algorithm subroutines. Unlike approximate synthesis, which relies on the continuous-parameter optimization of a parametrized oracle, exact synthesis is a discrete optimization problem. We frame this synthesis as a single-player game, in which an AlphaZero agent builds the circuit gate by gate, aiming to reach a target unitary U . Within this framework, we recover known optimal Clifford+ T constructions for small unitaries on all-to-all architectures. Targeting light non-Clifford resources decomposition, we train an agent on a dynamic-circuit architecture, leading to a four T -gate decomposition of the Toffoli gate on three qubits and a clean ancilla. This work highlight the relevance of model-based reinforcement learning agent for unitary synthesis, and their use for the exploration of promising architectures. The work has been released as a preprint [164], together with the code of both the synthesis agent [165] and the quantum-circuit-synthesis reinforcement learning environment [166].

Looking ahead, this work remains a proof of concept and several directions of improvement are still open. First, making the agent Clifford-aware by feeding it a pre-processed measure of Cliffordness would help for efficient MCTS exploration. Moreover, explicitly penalizing circuit depth and non-Clifford gates, could guide the search toward depth and T -gates optimized decompositions. Further, agents should scale toward larger subroutines, such as quantum adders of a few bits. This will likely require both more powerful learning architectures and more compact representations of quantum unitaries. In this respect, transformer-based neural networks with attention layers could help capture long-range correlations across larger qubit registers, while a compact encoding based on ZX-calculus, such as the one proposed in Tsim [161], could provide a scalable state representation. Finally, modern hardware (Tensor Cores, BF16/FP8, higher VRAM bandwidth) would also bring orders-of-magnitude inference speed-ups over the GPU used in this work, directly translating into broader tree exploration at equal compute budget. Together, these developments could enable RL-agent to help compile quantum algorithms, beyond simple subroutines.

Future research

From a broader perspective, quantum compilation is now a major research area in the quantum information community. On the hardware side, recent demonstrations have shown logical qubits operating below the surface-code threshold [167, 168]. On the algorithmic side, resource estimates get closer to realistic implementation: factoring a 2048-bit RSA integer is projected to require fewer than a million noisy qubits [169], where twenty million were needed few years earlier [5]. As discussed in this manuscript, beyond the combinatorial hardness of finding a gate sequence, the circuit must fit with the constraints of the underlying error-correcting code and physical platform. This shifted the community from circuit representation to spacetime representation, a layer closer to the physical hardware that captures the full set of constraints, optimizing the logical implementation at the physical level rather than at the circuit level [170]. For instance with lattice surgery compilation, where ZX-graphs and their translation into spacetime pipe-diagram is a combinatorial problem analogous to circuit synthesis. Translating an abstract ZX representation of a logical operation into a physical pipe-diagram realization aims to minimize the spacetime volume, that is, the qubit and time resources on the physical chip, and exact solvers quickly become impracticable for large circuits [171]. As a heuristic solver for large combinatorial problems with sparse solutions and unreliable local observations, AlphaZero agent naturally fit with this layer. A recent lattice-surgery compiler, TopoLS [172], addresses this optimization with an MCTS-based pipeline: a first ZX-level reduction, followed by a heuristic MCTS optimization of the spacetime volume during the ZX-to-pipe translation. Replacing the MCTS by an AlphaZero-style agent would turn this layer-by-layer local optimization into a long-horizon planning strategy, capable of anticipating optimizations several layers ahead, on top of the neural networks generalization capabilities. Moreover, since both the ZX-level reduction and the ZX-to-pipe translation operate on the same ZX-graph, a sequential two-stage pipeline may miss joint optimizations. This calls for a multi-task reinforcement learning approach, in which a shared latent space in the value and policy networks would let a single agent learn a co-optimization of the two stages [173]. A further direction would extend such a compiler to 3D codes [174], where the spacetime diagram becomes four-dimensional and the geometric intuition guiding current 2D pipe-diagram compilation no longer applies. On a broader note, we expect reinforcement learning agents to be part in the realization of fault-tolerant quantum computing, thanks to their performance in complex optimization problems but also with their ability to find solutions adapted to constraints imposed by the physical implementations.

Part III

Appendix - Quantum Networks

A -Helpers

A.1 .Non-photon-number-resolving detector

Let us consider a non-resolving photon detector with an efficiency η . The POVM element associated with a *no-click* event reads

$$\Pi_0^a(\eta) = (1 - \eta)^{a^\dagger a}.$$

The probability of missing the detection of a state $|N\rangle$ containing N photons in mode $a^\dagger$ is

$$\langle N | \Pi_0^a(\eta) | N \rangle = (1 - \eta)^N = R^N,$$

where R is defined as $R = 1 - \eta$.

Including a dark-count probability dc in the photon detector, the probability of missing the detection of a state $|N\rangle$ becomes $\langle N | \Pi_0^a(\eta, \text{dc}) | N \rangle = (1 - \text{dc}) \cdot (1 - \eta)^N$. The updated POVM element for a *no-click* event is therefore

$$\Pi_0^a(\eta, \text{dc}) = (1 - \text{dc}) \cdot (1 - \eta)^{a^\dagger a}.$$

The POVM element for a *click* event then follows from the completeness of the two outcomes:

$$\begin{aligned} \text{No-click event: } \quad \Pi_0^a(\eta, \text{dc}) &= (1 - \text{dc}) \cdot R^{a^\dagger a}, \\ \text{Click event: } \quad \Pi_1^a(\eta, \text{dc}) &= \mathbb{1} - (1 - \text{dc}) \cdot R^{a^\dagger a}, \end{aligned} \tag{A.1}$$

with $R = 1 - \eta$. For simplicity, we omit the parameters η and dc in the notation and refer to these POVM elements as Π_0^a and Π_1^a .

A.2 .Trace of thermal states

This appendix derives general results for the partial trace of a two-mode and full trace of a n -modes squeezed state.

Partial trace

First we derive a practical expression of the partial trace of a two-mode squeezed state

$$\rho_a = (1 - x^2) \cdot \text{Tr}_b \left[\exp \left(x \cdot a^\dagger b^\dagger \right) |0\rangle \langle 0| \exp (x \cdot ab) \right] \quad (\text{A.2})$$

with $x \in [0, 1]$, and $a^\dagger$ and $b^\dagger$ two modes. First the definition of $\exp(\cdot)$ gives

$$\exp \left(x \cdot a^\dagger b^\dagger \right) |00\rangle = \sum_n \frac{x^n}{n!} a^{\dagger n} b^{\dagger n} |00\rangle = \sum_n x^n |nn\rangle, \quad (\text{A.3})$$

we thus have

$$\begin{aligned} \frac{\rho_a}{1 - x^2} &= \text{Tr}_b \left[\left(\sum_i x^i |ii\rangle \right) \cdot \left(\sum_j x^j \langle jj| \right) \right] \\ &= \sum_n \mathbb{1}_a \otimes \langle n|_b \left(\sum_{i,j} x^{i+j} |ii\rangle \langle jj| \right) \mathbb{1}_a \otimes |n\rangle_b \\ &= \sum_{n,i,j} x^{i+j} \langle n|i\rangle \langle j|n\rangle |i\rangle \langle i|_a \\ &= \sum_n x^{2n} |n\rangle \langle n|_a \end{aligned} \quad (\text{A.4})$$

hence, using the expression of a thermal state ρ_{Therm} :

$$\rho_{\text{Therm}}^a(x) = \frac{1}{1 + \bar{n}(x)} \sum_n \frac{\bar{n}(x)}{1 + \bar{n}(x)} |n\rangle \langle n|_a = (1 - x^2) \sum_n (x^2)^n |n\rangle \langle n|_a \quad (\text{A.5})$$

with the mean number of excitation defined as $\bar{n}(x) = \frac{x^2}{1-x^2}$. We finally get

$$(1 - x^2) \cdot \text{Tr}_b \left[\exp \left(x \cdot a^\dagger b^\dagger \right) |0\rangle \langle 0| \exp (x \cdot ab) \right] = \rho_{\text{Therm}}^a(x) \quad (\text{A.6})$$

Full trace

Now we go further and focus on the derivation of the full trace of an unnormalized n -modes squeezed state. Let's consider $\{a_i^\dagger\}_i$ and $\{b_i^\dagger\}_i$ two sets of n modes, M a $n \times n$ matrix, x a

real in $[0, 1]$, and $|\psi\rangle$ a state defined

$$|\psi\rangle = \exp \left[x \cdot \begin{pmatrix} a_1^\dagger & \cdots & a_n^\dagger \end{pmatrix} M \begin{pmatrix} b_1^\dagger \\ \vdots \\ b_n^\dagger \end{pmatrix} \right] |\underline{0}\rangle, \quad (\text{A.7})$$

here we aim to compute the trace of such a state

$$T = \text{Tr}_{\text{tot}} (|\psi\rangle \langle\psi|). \quad (\text{A.8})$$

The first step of the calculation consists of performing a singular value decomposition on the matrix M

$$M = U^\dagger \Sigma V \quad (\text{A.9})$$

with U and V unitary matrices, and $\Sigma = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_n)$ the diagonal matrix composed of the singular values of M . Based on the unitary property of U and V , we define two new sets of modes $\{\bar{a}_i^\dagger\}_i$ and $\{\bar{b}_i^\dagger\}_i$ such that¹

$$\begin{pmatrix} \bar{a}_1^\dagger \\ \vdots \\ \bar{a}_n^\dagger \end{pmatrix} = U \cdot \begin{pmatrix} a_1^\dagger \\ \vdots \\ a_n^\dagger \end{pmatrix} \quad \begin{pmatrix} \bar{b}_1^\dagger \\ \vdots \\ \bar{b}_n^\dagger \end{pmatrix} = V \cdot \begin{pmatrix} b_1^\dagger \\ \vdots \\ b_n^\dagger \end{pmatrix} \quad (\text{A.10})$$

we thus get

$$\begin{pmatrix} a_1^\dagger & \cdots & a_n^\dagger \end{pmatrix} M \begin{pmatrix} b_1^\dagger \\ \vdots \\ b_n^\dagger \end{pmatrix} = \begin{pmatrix} \bar{a}_1^\dagger & \cdots & \bar{a}_n^\dagger \end{pmatrix} \begin{pmatrix} \sigma_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \sigma_n \end{pmatrix} \begin{pmatrix} \bar{b}_1^\dagger \\ \vdots \\ \bar{b}_n^\dagger \end{pmatrix} = \sum_i^n \sigma_i \cdot \bar{a}_i^\dagger \bar{b}_i^\dagger, \quad (\text{A.11})$$

this leads to a gentle expression of $|\psi\rangle$

$$|\psi\rangle = \sqrt{1-x^2} e^{x \cdot \sum_i \sigma_i \bar{a}_i^\dagger \bar{b}_i^\dagger} |\underline{0}\rangle = \sqrt{1-x^2} \cdot \bigotimes_i \left(e^{x \cdot \sigma_i \bar{a}_i^\dagger \bar{b}_i^\dagger} |\underline{0}\rangle \right). \quad (\text{A.12})$$

Using the basis-invariance of the trace, we switch the trace from the pre-transformed mode basis to the new one $\text{Tr}_{\text{tot}} = \text{Tr}_{\{a_i\}, \{b_i\}} = \text{Tr} [\{\bar{a}_i\}, \{\bar{b}_i\}]$ and thus get

$$\begin{aligned} T &= \text{Tr}_{\{\bar{a}_i\}, \{\bar{b}_i\}} \left[\bigotimes_i e^{x \cdot \sigma_i \bar{a}_i^\dagger \bar{b}_i^\dagger} |\underline{0}\rangle \langle \underline{0}| e^{x \cdot \sum_i \sigma_i \bar{a}_i \bar{b}_i} \right] \\ &= \prod_i \text{Tr}_{\bar{a}_i, \bar{b}_i} \left[e^{x \cdot \sigma_i \bar{a}_i^\dagger \bar{b}_i^\dagger} |\underline{0}\rangle \langle \underline{0}| e^{x \cdot \sigma_i \bar{a}_i \bar{b}_i} \right]. \end{aligned} \quad (\text{A.13})$$

¹A linear combination of modes gives proper new modes if the transformation is unitary.

Finally using

$$\begin{aligned}
\text{Tr}_{a,b} \left[e^{x \cdot \sigma_i \bar{a}_i^\dagger \bar{b}_i^\dagger} |0\rangle \langle 0| e^{x \cdot \sigma_i \bar{a}_i \bar{b}_i} \right] &= \text{Tr}_{a,b} \left[\left(\sum_i (x \sigma_i)^i |ii\rangle \right) \cdot \left((x \sigma_j)^j \langle jj| \right) \right] \\
&= \sum_n \langle nn| \left(\sum_{i,j} x^{i+j} \sigma_i^i \sigma_j^j |ii\rangle \langle jj| \right) |nn\rangle = \sum_n (x \sigma_n)^{2n} = \frac{1}{1-x^2},
\end{aligned} \tag{A.14}$$

we get the expression of the trace $T = \text{Tr}_{tot} [|\psi\rangle \langle \psi|]$

$$T = \prod_i \frac{1}{1 - (x \sigma_i)^2}, \tag{A.15}$$

which is coherent with $\mathcal{N} = \det(\mathbb{1} - x^2 M M^\dagger) = \prod_i \frac{1}{1 - (x \sigma_i)^2}$, with $\mathcal{N}$ the normalization constant of $|\psi\rangle$.

A.3 .Integral calculation

This appendix aims to compute the integral:

$$I = \int \int_{\mathbb{C}^2} \exp \left[- \left(a|x|^2 + b|y|^2 + c|\alpha + \delta x + \tau y|^2 + d|\beta + \tau x - \delta y|^2 \right) \right] d^2x d^2y \quad (\text{A.16})$$

with, $(a, b, c, d) \in \mathbb{R}_+$ and $(\alpha, \beta) \in \mathbb{C}$.

Decomposing real and imaginary parts:

$$\begin{cases} x = x_1 + ix_2 \\ y = y_1 + iy_2 \end{cases} \quad \begin{cases} \alpha = \alpha_1 + i\alpha_2 \\ \beta = \beta_1 + i\beta_2 \end{cases} \quad (\text{A.17})$$

Since δ and τ are real, the affine combinations $\alpha + \delta x + \tau y$ and $\beta + \tau x - \delta y$ split into independent real and imaginary parts. The integral therefore factorizes as $I = I_1 \cdot I_2$, where I_k ($k = 1, 2$) carries the components $\{x_k, y_k, \alpha_k, \beta_k\}$:

$$I_k = \int \int_{\mathbb{R}^2} \exp \left[-ax_k^2 - by_k^2 - c(\alpha_k + \delta x_k + \tau y_k)^2 - d(\beta_k + \tau x_k - \delta y_k)^2 \right] d^2x_k d^2y_k.$$

Expanding the bracket and regrouping it as a quadratic form Q_k such as

$$Q_k = Ax_k^2 + By_k^2 + 2Cx_ky_k + 2P_kx_k + 2R_ky_k + c\alpha_k^2 + d\beta_k^2, \quad (\text{A.18})$$

where we have introduced the coefficients

$$A = a + \delta^2c + \tau^2d, \quad B = b + \tau^2c + \delta^2d, \quad C = \delta\tau(c - d), \quad (\text{A.19})$$

and the k -dependents coefficients

$$P_k = \delta c \alpha_k + \tau d \beta_k, \quad R_k = \tau c \alpha_k - \delta d \beta_k. \quad (\text{A.20})$$

We integrate first over y_k . The y_k -dependent part of the exponent is $-By_k^2 - 2(Cx_k + R_k)y_k$, so the Gauss integral² yields

$$I_k = \sqrt{\frac{\pi}{B}} e^{-(c\alpha_k^2 + d\beta_k^2)} \int_{\mathbb{R}} \exp \left[-Ax_k^2 - 2P_kx_k + \frac{(Cx_k + R_k)^2}{B} \right] d^2x_k. \quad (\text{A.21})$$

The remaining exponent is quadratic in x_k with coefficients

$$-\frac{AB - C^2}{B} \text{ for } x_k^2 \quad \text{and} \quad -\frac{2(BP_k - CR_k)}{B} \text{ for } x_k. \quad (\text{A.22})$$

²Gauss integral: $\int_{\mathbb{R}} \exp [-ax^2 - bx - c] d^2x = \sqrt{\frac{\pi}{a}} \exp \left[\frac{b^2}{4a} - c \right].$

The second Gauss integration over x_k then gives the prefactor

$$\sqrt{\frac{\pi}{B}} \cdot \sqrt{\frac{\pi B}{AB - C^2}} = \frac{\pi}{\sqrt{AB - C^2}}, \quad (\text{A.23})$$

together with a residual exponent we call Q_k^{res} . After calculation, the residual exponent reduces, with $\Delta \equiv AB - C^2$, to the compact symmetric form

$$Q_k^{\text{res}} = \frac{1}{\Delta} [bcd(\tau\alpha_k - \delta\beta_k)^2 + acd(\delta\alpha_k + \tau\beta_k)^2 + ab(c\alpha_k^2 + d\beta_k^2)], \quad (\text{A.24})$$

where

$$\Delta = AB - C^2 = (a + \delta^2 c + \tau^2 d)(b + \delta^2 d + \tau^2 c) - \delta^2 \tau^2 (c - d)^2, \quad (\text{A.25})$$

with $\Delta > 0$ for convergence. Hence $I_k = \frac{\pi}{\sqrt{\Delta}} e^{-Q_k^{\text{res}}}$.

The full integral is the product $I = I_1 I_2$. The prefactors multiply to π^2/Δ , and the exponents add to reconstruct the squared moduli, for instance $(\tau\alpha_1 - \delta\beta_1)^2 + (\tau\alpha_2 - \delta\beta_2)^2 = |\tau\alpha - \delta\beta|^2$. We finally get:

$$I = \frac{\pi^2}{\Delta} \exp \left[-\frac{bcd|\tau\alpha - \delta\beta|^2 + acd|\delta\alpha + \tau\beta|^2 + ab(c|\alpha|^2 + d|\beta|^2)}{\Delta} \right] \quad (\text{A.26})$$

B -Perturbative study of an elementary link

B.1 .Block-diagonal structure of the heralded state

We define $\tilde{\mathcal{H}}(n)$ as the Hilbert space generated by the states carrying n excitations across two modes, formally

$$\tilde{\mathcal{H}}(n) = \text{Span} \{ |ij\rangle \text{ s.t. } i + j = n \}. \quad (\text{B.1})$$

This appendix establishes two complementary results regarding the action of the non-trace-preserving quantum channel of the elementary link

$$\mathcal{E}_{\text{Lnk}} : \rho_{\text{ini}} \mapsto \text{Tr}_{\text{CA,CB}} \left[\sqrt{\Pi_{\text{Lnk}}} U_{\text{BS}} \rho_{\text{ini}} U_{\text{BS}}^\dagger \sqrt{\Pi_{\text{Lnk}}} \right] \quad (\text{B.2})$$

on the initial state ρ_{ini} . First, the initial state admits a block decomposition along the subspaces $\tilde{\mathcal{H}}(n)$, with diagonal blocks scaling as p^n . Second, the channel $\mathcal{E}_{\text{Lnk}}$ preserves this block structure: each subspace $\tilde{\mathcal{H}}(n)$ is stable under $\mathcal{E}_{\text{Lnk}}$, and all off-block-diagonal cross-terms are erased. Physically, if k photons meet l photons at the central station, the resulting state shared between Alice's and Bob's memories is a delocalised state of exactly $k + l$ excitations. The argument relies on the photon-number conservation of the beam-splitter unitary, combined with the structure of the SPDC expansion in the symmetric regime $p_A = p_B = p$. Schematically,

$$\rho_{\text{ini}} = \frac{1}{\mathcal{N}} \begin{pmatrix} p^0 \rho_{\text{ini}}^{(0)} & \star & \star & \cdots \\ \star & p \rho_{\text{ini}}^{(1)} & \star & \cdots \\ \star & \star & p^2 \rho_{\text{ini}}^{(2)} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad (\text{B.3})$$

$$\xrightarrow{\mathcal{E}_{\text{Lnk}}} \frac{1}{\mathcal{N}} \begin{pmatrix} p^0 \mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(0)}) & 0 & 0 & \cdots \\ 0 & p \mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(1)}) & 0 & \cdots \\ 0 & 0 & p^2 \mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(2)}) & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

Equivalently, the heralded state lies in the direct sum of the photon-number sectors,

$$\rho_{\text{Lnk}} = \bigoplus_n \rho_{\text{Lnk}}^{(n)}. \quad (\text{B.4})$$

Block structure of the initial state

The two SPDC sources produce the initial state

$$|\psi_{\text{SPDC}}\rangle = \sqrt{1-p} \sum_k \sqrt{p}^k |kk\rangle, \quad (\text{B.5})$$

in the modes $(a^\dagger, b^\dagger)$ for Alice and $(c^\dagger, d^\dagger)$ for Bob. The full initial state is the tensor product of the two source contributions,

$$\rho_{\text{ini}} = \rho_{\text{SPDC}} \otimes \rho_{\text{SPDC}} \propto \sum_{k,k',l,l'} \sqrt{p}^{k+k'+l+l'} |kk\rangle \langle k'k'|_{ab} \otimes |ll\rangle \langle l'l'|_{cd}. \quad (\text{B.6})$$

The diagonal terms $k = k'$ and $l = l'$ scale as p^{k+l} and they live in $\tilde{\mathcal{H}}_{\text{idler}}(k+l) \otimes \tilde{\mathcal{H}}_{\text{signal}}(k+l)$ (with idler photons in the left and idler on the write), i.e. the subspace carrying exactly $2(k+l)$ excitations across the four modes. Collecting these diagonal contributions by total photon number $n = k + l$ yields the block decomposition

$$\rho_{\text{ini}} = \frac{1}{\mathcal{N}} \sum_{n \geq 0} p^n \rho_{\text{ini}}^{(n)} + \rho_{\text{ini}}^*(p), \quad (\text{B.7})$$

where $\rho_{\text{ini}}^{(n)}$ is the normalised density matrix of the n -photon sector, living in $\tilde{\mathcal{H}}_{\text{idler}}(n) \otimes \tilde{\mathcal{H}}_{\text{signal}}(n)$, and ρ_{ini}^* collects the off-block-diagonal cross-terms $k \neq k'$ or $l \neq l'$.

Stability of the channel operator

We now show that the block structure of ρ_{ini} is preserved by $\mathcal{E}_{\text{Lnk}}$: each n -photon sector is mapped onto itself, and the cross-terms ρ_{ini}^* are erased. From a certain perspective, the stability of the n -photon sectors reads

$$\mathcal{E}_{\text{Lnk}} \left(\tilde{\mathcal{H}}_{\text{idler}}(n) \otimes \tilde{\mathcal{H}}_{\text{signal}}(n) \right) \rightarrow \tilde{\mathcal{H}}_{\text{signal}}(n), \quad (\text{B.8})$$

meaning every element $\rho_{\text{ini}}^{(n)} \in \tilde{\mathcal{H}}_{\text{idler}}(n) \otimes \tilde{\mathcal{H}}_{\text{signal}}(n)$ of the $2n$ -photon sector (idler and signal n -photon sector) leads to a state in the signal n -photon sector $\rho_{\text{Lnk}}^{(n)} \in \tilde{\mathcal{H}}(n)$. By linearity, it is sufficient to study the action of $\mathcal{E}_{\text{Lnk}}$ on a single elementary term of the expansion,

$$\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}) \propto \sum_{kk'll'} \sqrt{p}^{k+k'+l+l'} \mathcal{E}_{\text{Lnk}} \left(\rho^{(kk'll')} \right), \quad (\text{B.9})$$

with

$$\rho^{(kk'll')} = |kkll\rangle \langle k'k'll'|_{abcd} \quad (\text{B.10})$$

$$= |kl\rangle \langle k'l'|_{bc} \otimes |kl\rangle \langle k'l'|_{ad}, \quad (\text{B.11})$$

where the last line groups the central station modes on the front, $|abcd\rangle \rightarrow |bc\rangle |ad\rangle$, in order to isolate the modes b and c that interfere at the beam splitter from the modes a and d that propagate to the memories. The beam splitter mixes the modes b and c together, but conserves the total photon number, so it acts within each subspace $\tilde{\mathcal{H}}(n)$ separately,

$$U_{\text{BS}} \rho^{(kk' ll')} U_{\text{BS}}^\dagger = U_{\text{BS}} |kl\rangle \langle k'l'|_{bc} U_{\text{BS}}^\dagger \otimes |kl\rangle \langle k'l'|_{ad} \quad (\text{B.12})$$

$$= |\psi^{(kl)}\rangle \langle \tilde{\psi}^{(k'l')}|_{bc} \otimes |kl\rangle \langle k'l'|_{ad}, \quad (\text{B.13})$$

with $|\psi^{(kl)}\rangle \in \tilde{\mathcal{H}}(k+l)$ and $|\tilde{\psi}^{(k'l')}\rangle \in \tilde{\mathcal{H}}(k'+l')$. The heralding measurement only registers click patterns, so Π_{Lnk} is diagonal in the Fock basis $\{|ij\rangle\}$ and cannot couple sectors of different total photon number; consequently, each $\tilde{\mathcal{H}}(n)$ is stable under Π_{Lnk} ,

$$\Pi_{\text{Lnk}} |\psi^{(kl)}\rangle = |\tilde{\psi}^{(kl)}\rangle \in \tilde{\mathcal{H}}(k+l), \quad (\text{B.14})$$

and similarly for $|\tilde{\psi}^{(k'l')}\rangle$. Finally, the partial trace over the central modes projects the ket and the bra onto the same photon-number subspace, so only terms with matching total excitation survive,

$$\mathcal{E}_{\text{Lnk}} \left(\rho^{(kk' ll')} \right) = \text{Tr}_{bc} |\tilde{\psi}^{(kl)}\rangle \langle \tilde{\psi}^{(k'l')}|_{bc} \cdot |kl\rangle \langle k'l'|_{ad}, \quad (\text{B.15})$$

with

$$\text{Tr}_{bc} |\tilde{\psi}^{(kl)}\rangle \langle \tilde{\psi}^{(k'l')}|_{bc} = 0 \quad \text{if } k+l \neq k'+l', \quad (\text{B.16})$$

since the bra and the ket belong to orthogonal subspaces of the Fock space. We have thus shown

$$\mathcal{E}_{\text{Lnk}} \left(\rho^{(kk' ll')} \right) = \begin{cases} 0 & \text{if } k+l \neq k'+l', \\ \rho_{\text{Lnk}}^{(k+l)} & \text{otherwise,} \end{cases} \quad (\text{B.17})$$

which establishes that $\mathcal{E}_{\text{Lnk}}$ maps each n -photon sector onto itself and annihilates all cross-terms between sectors of different total photon number.

B.2 .Elementary link state expansion

This appendix derives the perturbative expansion of the heralded state ρ_{Lnk} , the total heralding probability $\mathbb{P}_{\text{Lnk}}$, and the fidelity $\mathcal{F}_{\text{Lnk}}$ in powers of the pump parameter p , starting from the block-diagonal decomposition of the initial state established in appendix B.1. The argument relies on the linearity of the heralding map and a Bayesian reading of the conditional heralding probabilities attached to each i -photon sector.

Setup

Let $\mathcal{E}_{\text{Lnk}}$ denote the elementary link heralding map, namely the composition of the beam-splitter, the single-click pattern Π_{Lnk} , and the partial trace over the central station modes,

$$\mathcal{E}_{\text{Lnk}} : \rho_{\text{ini}} \longmapsto \text{Tr}_{\text{CA}, \text{CB}} \left[\sqrt{\Pi_{\text{Lnk}}} U_{\text{BS}} \rho_{\text{ini}} U_{\text{BS}}^\dagger \sqrt{\Pi_{\text{Lnk}}} \right] \quad (\text{B.18})$$

The map $\mathcal{E}_{\text{Lnk}}$ is completely positive and linear, but not trace-preserving: the trace of its output is precisely the heralding probability of the input state. Following the algebraic study of appendix B.1, the initial state admits the block-diagonal decomposition

$$\rho_{\text{ini}} = \frac{1}{\mathcal{N}} \sum_{i=0}^N p^i \rho_{\text{ini}}^{(i)} + \rho_{\text{ini}}^*(p) + o(p^N), \quad \mathcal{N} = \sum_{i=0}^N p^i, \quad (\text{B.19})$$

where each $\rho_{\text{ini}}^{(i)}$ is the normalised density matrix of the i -photon sector ($\text{Tr} [\rho_{\text{ini}}^{(i)}] = 1$), and ρ_{ini}^* collects the off-block-diagonal cross-terms. Since $\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^*) = 0$, the cross-terms drop out of the analysis and only the diagonal blocks $\rho_{\text{ini}}^{(i)}$ need to be tracked.

Action of the heralding map

By linearity of $\mathcal{E}_{\text{Lnk}}$, the unnormalised heralded state reads

$$\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}) = \frac{1}{\mathcal{N}} \sum_{i=0}^N p^i \mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(i)}) + o(p^N). \quad (\text{B.20})$$

Each term $\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(i)})$ encodes the contribution of the single i -photon sector to the heralded output, and can be split into a probability $\mathbb{P}_{\text{Lnk}}^{(i)}$ and a normalised state $\rho_{\text{Lnk}}^{(i)}$. We define the bare heralding probability of the i -photon sector as

$$\mathbb{P}_{\text{Lnk}}^{(i)} = \text{Tr} \left[\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(i)}) \right], \quad (\text{B.21})$$

together with the corresponding normalised heralded state

$$\rho_{\text{Lnk}}^{(i)} = \frac{\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(i)})}{\mathbb{P}_{\text{Lnk}}^{(i)}}, \quad \text{Tr}[\rho_{\text{Lnk}}^{(i)}] = 1. \quad (\text{B.22})$$

Each pair $(\mathbb{P}_{\text{Lnk}}^{(i)}, \rho_{\text{Lnk}}^{(i)})$ is an intrinsic object of the i -photon sector, independent of the other sectors, and satisfies

$$\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}}^{(i)}) = \mathbb{P}_{\text{Lnk}}^{(i)} \rho_{\text{Lnk}}^{(i)}. \quad (\text{B.23})$$

Heralding probability

The total heralding probability follows by linearity

$$\mathbb{P}_{\text{Lnk}} = \text{Tr}[\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}})] = \frac{1}{\mathcal{N}} \sum_{i=0}^N p^i \mathbb{P}_{\text{Lnk}}^{(i)} + o(p^N). \quad (\text{B.24})$$

This identity is the operator-level analogue of the law of total probability: $\mathbb{P}_{\text{Lnk}}$ averages the conditional heralding probabilities $\mathbb{P}_{\text{Lnk}}^{(i)}$ over the prior weights $p^i/\mathcal{N}$ of the SPDC decomposition, where each term is an independent contribution of the i -photon sector.

The total heralding probability $\mathbb{P}_{\text{Lnk}}$ admits a parallel perturbative expansion. Expanding $1/\mathcal{N} = 1 - p + o(p^N)$ at finite truncation order N , regrouping by powers of p yields

$$\mathbb{P}_{\text{Lnk}} = \sum_{i=0}^N p^i \tilde{\mathbb{P}}_{\text{Lnk}}^{(i)} + o(p^N), \quad \tilde{\mathbb{P}}_{\text{Lnk}}^{(i)} = \mathbb{P}_{\text{Lnk}}^{(i)} - \mathbb{P}_{\text{Lnk}}^{(i-1)}, \quad \mathbb{P}_{\text{Lnk}}^{(-1)} \equiv 0. \quad (\text{B.25})$$

The expansion is telescopic: each order extracts the increment in conditional heralding probability when one opens a new photon-number channel. The first orders read

$$\tilde{\mathbb{P}}_{\text{Lnk}}^{(0)} = \mathbb{P}_{\text{Lnk}}^{(0)}, \quad (\text{B.26})$$

$$\tilde{\mathbb{P}}_{\text{Lnk}}^{(1)} = \mathbb{P}_{\text{Lnk}}^{(1)} - \mathbb{P}_{\text{Lnk}}^{(0)}, \quad (\text{B.27})$$

$$\tilde{\mathbb{P}}_{\text{Lnk}}^{(2)} = \mathbb{P}_{\text{Lnk}}^{(2)} - \mathbb{P}_{\text{Lnk}}^{(1)}. \quad (\text{B.28})$$

This structure is markedly simpler than the perturbative expansion of ρ_{Lnk} : the bare probabilities $\mathbb{P}_{\text{Lnk}}^{(i)}$ enter linearly, with no Bayesian back-reaction. The reason is that $\mathbb{P}_{\text{Lnk}}$ is a scalar trace, which depends only on the marginal heralding probabilities and not on the relative weights of the sectors in the heralded state.

Heralded state as a statistical mixture

The normalised heralded state is, by definition,

$$\rho_{\text{Lnk}} = \frac{\mathcal{E}_{\text{Lnk}}(\rho_{\text{ini}})}{\mathbb{P}_{\text{Lnk}}}. \quad (\text{B.29})$$

Substituting equation (B.25) and the sector-wise expressions, the factor $1/\mathcal{N}$ cancels between numerator and denominator, leaving

$$\rho_{\text{Lnk}} = \sum_{i=0}^N w_i \rho_{\text{Lnk}}^{(i)} + \mathcal{O}(p^N), \quad w_i = \frac{p^i \mathbb{P}_{\text{Lnk}}^{(i)}}{\sum_{j=0}^N p^j \mathbb{P}_{\text{Lnk}}^{(j)}}. \quad (\text{B.30})$$

The weights w_i admit a Bayesian reading: $w_i = \mathbb{P}(i | \text{H})$ is the posterior probability that the heralded contribution originates from the i -photon sector, given the heralding event H. The decomposition of equation (B.30) is exact within the truncation $\mathcal{O}(p^N)$. Normalisation is manifest: $\sum_i w_i = 1$ by construction, so $\text{Tr}[\rho_{\text{Lnk}}] = 1$ identically and at every order of the perturbative expansion. Positivity follows from the positivity of each $\rho_{\text{Lnk}}^{(i)}$ and of each w_i .

Perturbative expansion of the heralded state

To exhibit the hierarchy of multi-photon corrections, the weights w_i are expanded in powers of p . Introducing the relative heralding probabilities $\hat{\mathbb{P}}_{\text{Lnk}}^{(i)} = \mathbb{P}_{\text{Lnk}}^{(i)} / \mathbb{P}_{\text{Lnk}}^{(0)}$ with $\hat{\mathbb{P}}_{\text{Lnk}}^{(0)} = 1$, the weights read

$$w_i = \frac{p^i \hat{\mathbb{P}}_{\text{Lnk}}^{(i)}}{1 + \sum_{j=1}^N p^j \hat{\mathbb{P}}_{\text{Lnk}}^{(j)}} + \mathcal{O}(p^N). \quad (\text{B.31})$$

The denominator is expanded as a formal inverse series,

$$\frac{1}{\sum_{j=0}^N p^j \hat{\mathbb{P}}_{\text{Lnk}}^{(j)}} = \sum_{k=0}^N p^k c_k + \mathcal{O}(p^N), \quad c_0 = 1, \quad c_k = - \sum_{j=1}^k \hat{\mathbb{P}}_{\text{Lnk}}^{(j)} c_{k-j}, \quad (\text{B.32})$$

which directly captures the Bayesian back-reaction induced by the renormalisation of w_i . Reorganising by powers of p , the heralded state takes the compact form

$$\rho_{\text{Lnk}} = \sum_{i=0}^N p^i \left(\hat{\mathbb{P}}_{\text{Lnk}}^{(i)} \rho_{\text{Lnk}}^{(i)} + \sum_{k=0}^{i-1} c_{i-k} \hat{\mathbb{P}}_{\text{Lnk}}^{(k)} \rho_{\text{Lnk}}^{(k)} \right) + \mathcal{O}(p^N). \quad (\text{B.33})$$

The first term $\hat{\mathbb{P}}_{\text{Lnk}}^{(i)} \rho_{\text{Lnk}}^{(i)}$ is the direct contribution of the i -photon sector, weighted by its relative heralding probability. The sum collects the *dilution corrections*: opening the i -photon channel reshuffles the weights of all lower sectors, with the coefficients c_{i-k} encoding the resulting redistribution. Defining the dilution-corrected contribution of the

i -photon sector as

$$\tilde{\rho}_{\text{Lnk}}^{(i)} = \hat{\mathbb{P}}_{\text{Lnk}}^{(i)} \rho_{\text{Lnk}}^{(i)} + \sum_{k=0}^{i-1} c_{i-k} \hat{\mathbb{P}}_{\text{Lnk}}^{(k)} \rho_{\text{Lnk}}^{(k)}, \quad (\text{B.34})$$

the expansion takes the compact form

$$\rho_{\text{Lnk}} = \sum_{i=0}^N p^i \tilde{\rho}_{\text{Lnk}}^{(i)} + \mathcal{O}(p^N). \quad (\text{B.35})$$

The first orders read

$$\tilde{\rho}_{\text{Lnk}}^{(0)} = \rho_{\text{Lnk}}^{(0)}, \quad (\text{B.36})$$

$$\tilde{\rho}_{\text{Lnk}}^{(1)} = \hat{\mathbb{P}}_{\text{Lnk}}^{(1)} \left(\rho_{\text{Lnk}}^{(1)} - \rho_{\text{Lnk}}^{(0)} \right), \quad (\text{B.37})$$

$$\tilde{\rho}_{\text{Lnk}}^{(2)} = \hat{\mathbb{P}}_{\text{Lnk}}^{(2)} \left(\rho_{\text{Lnk}}^{(2)} - \rho_{\text{Lnk}}^{(0)} \right) - \left(\hat{\mathbb{P}}_{\text{Lnk}}^{(1)} \right)^2 \left(\rho_{\text{Lnk}}^{(1)} - \rho_{\text{Lnk}}^{(0)} \right). \quad (\text{B.38})$$

The leading term is the heralded state of the vacuum sector. The first correction quantifies the contamination by the single-photon sector, weighted by the relative heralding probability $\hat{\mathbb{P}}_{\text{Lnk}}^{(1)}$. Higher-order corrections combine the direct contribution of each sector with the Bayesian back-reactions encoded in the c_k .

The recurrence $\sum_{k=0}^i \hat{\mathbb{P}}_{\text{Lnk}}^{(k)} c_{i-k} = \delta_{i,0}$ implies $\text{Tr} \left[\tilde{\rho}_{\text{Lnk}}^{(0)} \right] = 1$ and $\text{Tr} \left[\tilde{\rho}_{\text{Lnk}}^{(i)} \right] = 0$ for $i \geq 1$: the dilution corrections are traceless operators that redistribute weight across sectors without altering the global normalisation, in direct analogy with central moments in classical probability.

Perturbative expansion of the fidelity

The fidelity to the target Bell state $|\Psi^+\rangle$ inherits the same structure as the heralded state, since it is a linear functional of ρ_{Lnk} . Defining the sectorial fidelity

$$\mathcal{F}_{\text{Lnk}}^{(i)} = \langle \Psi^+ | \rho_{\text{Lnk}}^{(i)} | \Psi^+ \rangle, \quad (\text{B.39})$$

linearity propagates the mixture decomposition of equation (B.30) directly to the fidelity,

$$\mathcal{F}_{\text{Lnk}} = \sum_{i=0}^N w_i \mathcal{F}_{\text{Lnk}}^{(i)} + \mathcal{O}(p^N). \quad (\text{B.40})$$

Substituting the perturbative expansion of w_i derived above, the fidelity takes the compact form

$$\mathcal{F}_{\text{Lnk}} = \sum_{i=0}^N p^i \tilde{\mathcal{F}}_{\text{Lnk}}^{(i)} + \mathcal{O}(p^N), \quad \tilde{\mathcal{F}}_{\text{Lnk}}^{(i)} = \hat{\mathbb{P}}_{\text{Lnk}}^{(i)} \mathcal{F}_{\text{Lnk}}^{(i)} + \sum_{k=0}^{i-1} c_{i-k} \hat{\mathbb{P}}_{\text{Lnk}}^{(k)} \mathcal{F}_{\text{Lnk}}^{(k)}, \quad (\text{B.41})$$

with the same coefficients c_k as for the state expansion. The dilution-corrected sectorial fidelities $\tilde{\mathcal{F}}_{\text{Lnk}}^{(i)}$ play the role of the $\tilde{\rho}_{\text{Lnk}}^{(i)}$: the bare contribution $\hat{\mathbb{P}}_{\text{Lnk}}^{(i)} \mathcal{F}_{\text{Lnk}}^{(i)}$ is augmented by the Bayesian back-reaction from lower-order sectors.

Case of a vanishing vacuum sector

The expansion above assumes $\mathbb{P}_{\text{Lnk}}^{(0)} \neq 0$, which holds whenever dark counts are present. In their absence, the vacuum sector cannot trigger a heralding event and $\mathbb{P}_{\text{Lnk}}^{(0)} = 0$. The mixture decomposition of equation (B.30) remains valid, with the sum restricted to the active sectors $i \geq 1$. The perturbative expansion then proceeds with the lowest active sector as reference; redefining $\hat{\mathbb{P}}_{\text{Lnk}}^{(i)} = \mathbb{P}_{\text{Lnk}}^{(i)} / \mathbb{P}_{\text{Lnk}}^{(1)}$ yields

$$\rho_{\text{Lnk}} = \rho_{\text{Lnk}}^{(1)} + p \hat{\mathbb{P}}_{\text{Lnk}}^{(2)} \left(\rho_{\text{Lnk}}^{(2)} - \rho_{\text{Lnk}}^{(1)} \right) + \mathcal{O}(p^2). \quad (\text{B.42})$$

This regime corresponds to the standard DLCZ analysis, in which the leading heralded state is the single-photon Bell pair and the first correction encodes the multi-photon contamination.

B.3 .One- and two-photon contributions to an elementary link

This appendix derives the contribution of the 2-photon sector for the perturbative study of the elementary link state, fidelity, and heralding probability.

1-photon sector

Starting from the initial state

$$|\psi_{\text{ini}}^{(1)}\rangle = \frac{1}{\sqrt{2}} (|1100\rangle + |0011\rangle), \quad (\text{B.43})$$

the state $|\psi_{\text{ini}}^{(1)}\rangle$ after the beam-splitter reads

$$|0\rangle \frac{|01\rangle + |10\rangle}{\sqrt{2}} |1\rangle + |1\rangle \frac{|01\rangle - |10\rangle}{\sqrt{2}} |0\rangle. \quad (\text{B.44})$$

With a non-null dark-count dc_{Lnk} , the heralding measurement brings

$$\begin{aligned} |0\rangle \frac{|01\rangle + |10\rangle}{\sqrt{2}} |1\rangle &= \begin{cases} \frac{1}{\sqrt{2}} |0011\rangle \xrightarrow{\Pi_{\text{Lnk}}} \frac{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}}R_{\text{Lnk}})}{\sqrt{2}} |0011\rangle \\ \quad \rightarrow \frac{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}}R_{\text{Lnk}})}{\sqrt{2}} |01\rangle |01\rangle \\ \frac{1}{\sqrt{2}} |0101\rangle \xrightarrow{\Pi_{\text{Lnk}}} \frac{\kappa_{\text{Lnk}}R_{\text{Lnk}}(1-\kappa_{\text{Lnk}})}{\sqrt{2}} |0101\rangle \\ \quad \rightarrow \frac{\kappa_{\text{Lnk}}R_{\text{Lnk}}(1-\kappa_{\text{Lnk}})}{\sqrt{2}} |01\rangle |10\rangle \end{cases} \quad (\text{a}) \\ |1\rangle \frac{|01\rangle - |10\rangle}{\sqrt{2}} |0\rangle &= \begin{cases} \frac{1}{\sqrt{2}} |1010\rangle \xrightarrow{\Pi_{\text{Lnk}}} \frac{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}}R_{\text{Lnk}})}{\sqrt{2}} |1010\rangle \\ \quad \rightarrow \frac{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}}R_{\text{Lnk}})}{\sqrt{2}} |10\rangle |01\rangle \\ -\frac{1}{\sqrt{2}} |1100\rangle \xrightarrow{\Pi_{\text{Lnk}}} -\frac{\kappa_{\text{Lnk}}R_{\text{Lnk}}(1-\kappa_{\text{Lnk}})}{\sqrt{2}} |1100\rangle \\ \quad \rightarrow -\frac{\kappa_{\text{Lnk}}R_{\text{Lnk}}(1-\kappa_{\text{Lnk}})}{\sqrt{2}} |10\rangle |10\rangle \end{cases} \quad (\text{b}) \end{aligned} \quad (\text{B.45})$$

The last arrow re-orders the modes $|m_{\text{A}}, c_{\text{A}}, c_{\text{B}}, m_{\text{B}}\rangle \rightarrow |m_{\text{A}}, m_{\text{B}}\rangle |c_{\text{A}}, c_{\text{B}}\rangle$, helping for channels mode factorization for the trace computation. The calculation leads to a 1-photon sector contribution to the state such that

$$\rho_{\text{Lnk}}^{(1)} = \frac{1}{1 + \frac{\kappa_{\text{Lnk}}R_{\text{Lnk}}(1-\kappa_{\text{Lnk}})}{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}}R_{\text{Lnk}})}} |\Psi_1^+\rangle \langle \Psi_1^+| + \frac{1}{1 + \frac{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}}R_{\text{Lnk}})}{\kappa_{\text{Lnk}}R_{\text{Lnk}}(1-\kappa_{\text{Lnk}})}} |\Psi_1^-\rangle \langle \Psi_1^-|, \quad (\text{B.46})$$

with a heralding probability of

$$\mathbb{P}_{\text{Lnk}}^{(1)} = \kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}}R_{\text{Lnk}}) + \kappa_{\text{Lnk}}R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}}), \quad (\text{B.47})$$

and a fidelity

$$\mathcal{F}_{\text{Lnk}}^{(1)} = \frac{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}})}{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}) + \kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}})}. \quad (\text{B.48})$$

2-photon sector

Starting from the 2-photon sector

$$\left| \psi_{\text{ini}}^{(2)} \right\rangle = \frac{1}{\sqrt{3}} (|2200\rangle + |0022\rangle + |1111\rangle), \quad (\text{B.49})$$

the beam-splitter followed by single-click measurement leads to the kets

$$|2200\rangle \xrightarrow{\text{BS}} \left\{ \begin{array}{l} \frac{|2200\rangle}{2} \xrightarrow{\Pi_{\text{Lnk}}} \frac{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}^2)}{2} |2200\rangle \\ \quad \rightarrow \frac{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}^2)}{2} |20\rangle |20\rangle \\ - \frac{|2110\rangle}{\sqrt{2}} \xrightarrow{\Pi_{\text{Lnk}}} - \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}})}{\sqrt{2}} |2110\rangle \\ \quad \rightarrow - \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}})}{\sqrt{2}} |20\rangle |11\rangle \\ \frac{|2020\rangle}{2} \xrightarrow{\Pi_{\text{Lnk}}} \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}}^2 (1 - \kappa_{\text{Lnk}})}{2} |2020\rangle \\ \quad \rightarrow \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}}^2 (1 - \kappa_{\text{Lnk}})}{2} |20\rangle |02\rangle \end{array} \right. \quad (\text{a})$$

$$|0022\rangle \xrightarrow{\text{BS}} \left\{ \begin{array}{l} \frac{|0202\rangle}{2} \xrightarrow{\Pi_{\text{Lnk}}} \frac{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}^2)}{2} |0202\rangle \\ \quad \rightarrow \frac{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}^2)}{2} |02\rangle |20\rangle \\ \frac{|0112\rangle}{\sqrt{2}} \xrightarrow{\Pi_{\text{Lnk}}} \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}})}{\sqrt{2}} |0112\rangle \\ \quad \rightarrow \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}})}{\sqrt{2}} |02\rangle |11\rangle \\ \frac{|0022\rangle}{2} \xrightarrow{\Pi_{\text{Lnk}}} \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}}^2 (1 - \kappa_{\text{Lnk}})}{2} |0022\rangle \\ \quad \rightarrow \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}}^2 (1 - \kappa_{\text{Lnk}})}{2} |02\rangle |02\rangle \end{array} \right. \quad (\text{b})$$

$$|1111\rangle \xrightarrow{\text{BS}} \left\{ \begin{array}{l} \frac{|1201\rangle}{\sqrt{2}} \xrightarrow{\Pi_{\text{Lnk}}} \frac{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}^2)}{\sqrt{2}} |1201\rangle \\ \quad \rightarrow \frac{\kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}^2)}{\sqrt{2}} |11\rangle |20\rangle \\ - \frac{|1021\rangle}{\sqrt{2}} \xrightarrow{\Pi_{\text{Lnk}}} - \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}}^2 (1 - \kappa_{\text{Lnk}})}{\sqrt{2}} |1021\rangle \\ \quad \rightarrow - \frac{\kappa_{\text{Lnk}} R_{\text{Lnk}}^2 (1 - \kappa_{\text{Lnk}})}{\sqrt{2}} |11\rangle |02\rangle \end{array} \right. \quad (\text{c})$$

The last arrow re-orders the modes $|m_A, c_A, c_B, m_B\rangle \rightarrow |m_A, m_B\rangle |c_A, c_B\rangle$, helping for channels mode $|c_A, c_B\rangle$ factorization for the trace computation. The derivation of the 2-photon

contribution leads to the state

$$\begin{aligned}\rho_{\text{Lnk}} = & \frac{1}{1 + \frac{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}})R_{\text{Lnk}}^2 + \kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}})R_{\text{Lnk}}}{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}})R_{\text{Lnk}}^2}} |\Psi_{12}^-\rangle \langle \Psi_{12}^-| \\ & + \frac{1}{1 + \frac{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}})R_{\text{Lnk}}^2 + \kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}})R_{\text{Lnk}}}{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}})R_{\text{Lnk}}^2}} |\Psi_{12}^+\rangle \langle \Psi_{12}^+| \\ & + \frac{1}{1 + \frac{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}})R_{\text{Lnk}}^2 + \kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}})R_{\text{Lnk}}^2}{\kappa_{\text{Lnk}}(1-\kappa_{\text{Lnk}})R_{\text{Lnk}}}} |\Psi_2^-\rangle \langle \Psi_2^-| \end{aligned}$$

with a heralding probability

$$\begin{aligned}\mathbb{P}_{\text{Lnk}}^{(2)} = & \frac{2}{3} \left[\frac{1}{4} \kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}^2) + \frac{1}{2} \kappa_{\text{Lnk}} R_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}) + \frac{1}{4} \kappa_{\text{Lnk}} R_{\text{Lnk}}^2 (1 - \kappa_{\text{Lnk}}) \right] \\ & + \frac{1}{3} \left[\frac{1}{2} \kappa_{\text{Lnk}} (1 - \kappa_{\text{Lnk}} R_{\text{Lnk}}^2) + \frac{1}{2} \kappa_{\text{Lnk}} R_{\text{Lnk}}^2 (1 - \kappa_{\text{Lnk}}) \right]. \end{aligned} \quad (\text{B.50})$$

and the fidelity reads

$$\mathcal{F}_{\text{Lnk}}^{(2)} = 0. \quad (\text{B.51})$$

B.4 .Multi-link swapping

This appendix derives the state swapped in a multi-link repeater chain as presented in section 1.2. The result is established by induction on the binary-tree depth $d \in \{0, \dots, n\}$, with $N = 2^n$ the total number of elementary sub-links. In this appendix we do not consider dark counts ($\text{dc}_{\text{swp}} = 0$), and we keep, a priori, an asymmetric swap operation $\alpha^L \neq \alpha^R$ to retain generality in the structure. The symmetric situation is then derived from the general expression.

Induction setup

We aim to show that at every depth d , the heralded state propagated through 2^d sub-links reads

$$\rho_{\text{swp}}^{(d)} = (1 - \alpha_d) |00\rangle \langle 00| + \alpha_d |\Psi^+\rangle \langle \Psi^+|, \quad (\text{B.52})$$

with the Bell weight α_d following the recurrence

$$\alpha_d = \frac{1}{\frac{\alpha_{d-1}^L + \alpha_{d-1}^R}{2} + 1 - \frac{\eta_{\text{swp}}}{2}}, \quad (\text{B.53})$$

and the heralding probability of the depth- d swap reading

$$\mathbb{P}_{\text{swp}}^{(d)} = \frac{\eta_{\text{swp}}}{4} [(2 - \eta_{\text{swp}}) \alpha_{d-1}^L \alpha_{d-1}^R + \alpha_{d-1}^L + \alpha_{d-1}^R], \quad (\text{B.54})$$

where α_{d-1}^L and α_{d-1}^R denote the Bell weights of the left and right input states. Following the 1st-order expansion of section 1.1.1, the initial elementary-link state is a pure Bell state $\rho_{\text{swp}}^{(0)} = |\Psi^+\rangle \langle \Psi^+|$, so that $\alpha_0^L = \alpha_0^R = 1$.

Initialization ($d = 1$)

The depth-1 swap operates on two pure Bell pairs $\rho_{\text{Lnk}}^L = \rho_{\text{Lnk}}^R = |\Psi^+\rangle \langle \Psi^+|$, which is exactly the two-link configuration solved in equation (1.55) from section 1.2. Plugging $\alpha_0^L = \alpha_0^R = 1$ into the recurrence (B.53) gives

$$\alpha_1 = \frac{1}{1 + 1 - \frac{\eta_{\text{swp}}}{2}} = \frac{2}{4 - \eta_{\text{swp}}}, \quad 1 - \alpha_1 = \frac{2 - \eta_{\text{swp}}}{4 - \eta_{\text{swp}}}, \quad (\text{B.55})$$

which recovers the depth-1 heralded state

$$\rho_{\text{swp}}^{(1)} = \frac{2 - \eta_{\text{swp}}}{4 - \eta_{\text{swp}}} |00\rangle \langle 00| + \frac{2}{4 - \eta_{\text{swp}}} |\Psi^+\rangle \langle \Psi^+|. \quad (\text{B.56})$$

Similarly, the probability (B.54) reduces to

$$\mathbb{P}_{\text{Swp}}^{(1)} = \frac{\eta_{\text{Swp}}}{4} [(2 - \eta_{\text{Swp}}) + 1 + 1] = \frac{\eta_{\text{Swp}}(4 - \eta_{\text{Swp}})}{4}, \quad (\text{B.57})$$

matching the two-link result of equation (1.55) and completing the initialization.

Induction step

Let us assume that the depth- $(d - 1)$ heralded states from both sides of the swap follow the form of (B.52):

$$\begin{aligned} \rho_{\text{Swp}}^{\text{L}, (d-1)} &= (1 - \alpha_{d-1}^{\text{L}}) |00\rangle \langle 00| + \alpha_{d-1}^{\text{L}} |\Psi^+\rangle \langle \Psi^+|, \\ \rho_{\text{Swp}}^{\text{R}, (d-1)} &= (1 - \alpha_{d-1}^{\text{R}}) |00\rangle \langle 00| + \alpha_{d-1}^{\text{R}} |\Psi^+\rangle \langle \Psi^+|. \end{aligned} \quad (\text{B.58})$$

We show that the depth- d swap produces a state of the same form, with α_d given by (B.53) and heralding probability (B.54).

Expanding the tensor product $\rho_{\text{Swp}}^{\text{L}, (d-1)} \otimes \rho_{\text{Swp}}^{\text{R}, (d-1)}$ yields four contributions

$$\begin{aligned} \rho_{\text{Swp}}^{\text{L}, (d-1)} \otimes \rho_{\text{Swp}}^{\text{R}, (d-1)} &= (1 - \alpha_{d-1}^{\text{L}}) (1 - \alpha_{d-1}^{\text{R}}) |0000\rangle \langle 0000| \\ &\quad + \alpha_{d-1}^{\text{L}} (1 - \alpha_{d-1}^{\text{R}}) |\Psi^+\rangle \langle \Psi^+| \otimes |00\rangle \langle 00| \\ &\quad + (1 - \alpha_{d-1}^{\text{L}}) \alpha_{d-1}^{\text{R}} |00\rangle \langle 00| \otimes |\Psi^+\rangle \langle \Psi^+| \\ &\quad + \alpha_{d-1}^{\text{L}} \alpha_{d-1}^{\text{R}} |\Psi^+\rangle \langle \Psi^+| \otimes |\Psi^+\rangle \langle \Psi^+|. \end{aligned} \quad (\text{B.59})$$

The action of the swap channel $\mathcal{E}_{\text{Swp}}$ is computed term by term.

The vacuum-vacuum term $|0000\rangle \langle 0000|$ carries no excitation and cannot trigger the single-click heralding, hence

$$\mathcal{E}_{\text{Swp}} [|0000\rangle \langle 0000|] = 0. \quad (\text{B.60})$$

The fourth term is exactly the two-pure-Bell-pair input solved in section 1.2.2, leading to

$$\mathcal{E}_{\text{Swp}} [|\Psi^+\rangle \langle \Psi^+| \otimes |\Psi^+\rangle \langle \Psi^+|] = \frac{\eta_{\text{Swp}}(2 - \eta_{\text{Swp}})}{4} |00\rangle \langle 00| + \frac{\eta_{\text{Swp}}}{2} |\Psi^+\rangle \langle \Psi^+|. \quad (\text{B.61})$$

The second and third terms require explicit computation. Focusing on the second term, the beam-splitter U_{BS} acting on the central memory modes gives

$$\begin{aligned} |\Psi^+\rangle \otimes |00\rangle &= \frac{|0100\rangle + |1000\rangle}{\sqrt{2}} \\ &\xrightarrow{U_{\text{BS}}} \frac{|0\rangle \frac{|01\rangle + |10\rangle}{\sqrt{2}} |1\rangle + |1000\rangle}{2}, \end{aligned} \quad (\text{B.62})$$

followed by the single-click POVM $\Pi_{\text{Swp}} = R_{\text{Swp}}^{m_B^R \dagger m_B^R} \left(\mathbb{1} - R_{\text{Swp}}^{m_A^L \dagger m_A^L} \right)$ which selects the click/no-click event on the central pair:

$$\xrightarrow{\Pi_{\text{Swp}}} \frac{|0\rangle \langle \eta_{\text{Swp}} | 01 \rangle + 0 \cdot |10\rangle \rangle |1\rangle + 0 \cdot |1000\rangle}{2} = \frac{\eta_{\text{Swp}}}{2} \cdot |0010\rangle. \quad (\text{B.63})$$

Tracing over the central modes yields

$$\mathcal{E}_{\text{Swp}} [|\Psi^+\rangle \langle \Psi^+| \otimes |00\rangle \langle 00|] = \frac{\eta_{\text{Swp}}}{4} |00\rangle \langle 00|. \quad (\text{B.64})$$

By symmetry, the third term contributes the same amount

$$\mathcal{E}_{\text{Swp}} [|00\rangle \langle 00| \otimes |\Psi^+\rangle \langle \Psi^+|] = \frac{\eta_{\text{Swp}}}{4} |00\rangle \langle 00|. \quad (\text{B.65})$$

Summing the four contributions, the unnormalized output reads

$$\begin{aligned} \mathcal{E}_{\text{Swp}} \left[\rho_{\text{Swp}}^{L, (d-1)} \otimes \rho_{\text{Swp}}^{R, (d-1)} \right] &= \frac{\eta_{\text{Swp}}}{4} \left[(2 - \eta_{\text{Swp}}) \alpha_{d-1}^L \alpha_{d-1}^R + \alpha_{d-1}^L (1 - \alpha_{d-1}^R) \right. \\ &\quad \left. + (1 - \alpha_{d-1}^L) \alpha_{d-1}^R \right] |00\rangle \langle 00| \\ &\quad + \alpha_{d-1}^L \alpha_{d-1}^R \frac{\eta_{\text{Swp}}}{2} |\Psi^+\rangle \langle \Psi^+|. \end{aligned} \quad (\text{B.66})$$

The trace of this expression gives the depth- d heralding probability

$$\mathbb{P}_{\text{Swp}}^{(d)} = \frac{\eta_{\text{Swp}}}{4} \left[(2 - \eta_{\text{Swp}}) \alpha_{d-1}^L \alpha_{d-1}^R + \alpha_{d-1}^L + \alpha_{d-1}^R \right], \quad (\text{B.67})$$

which recovers (B.54). Normalizing by $\mathbb{P}_{\text{Swp}}^{(d)}$ yields the heralded state

$$\rho_{\text{Swp}}^{(d)} = (1 - \alpha_d) |00\rangle \langle 00| + \alpha_d |\Psi^+\rangle \langle \Psi^+|, \quad (\text{B.68})$$

with the Bell weight

$$\begin{aligned} \alpha_d &= \frac{\alpha_{d-1}^L \alpha_{d-1}^R \eta_{\text{Swp}}/2}{\mathbb{P}_{\text{Swp}}^{(d)}} \\ &= \frac{2 \alpha_{d-1}^L \alpha_{d-1}^R}{(2 - \eta_{\text{Swp}}) \alpha_{d-1}^L \alpha_{d-1}^R + \alpha_{d-1}^L + \alpha_{d-1}^R} \\ &= \frac{1}{\frac{\alpha_L + \alpha_R}{2\alpha_L \alpha_R} + 1 - \frac{\eta_{\text{Swp}}}{2}}, \end{aligned} \quad (\text{B.69})$$

which completes the induction and proves (B.52)-(B.54).

Symmetric situation

We now specialize to the fully symmetric case where both inputs of every swap share the same Bell weight, $\alpha_{d-1}^L = \alpha_{d-1}^R = \alpha_{d-1}$. The recurrence (B.53) simplifies to

$$\alpha_d = \frac{1}{\frac{1}{\alpha_{d-1}} + 1 - \frac{\eta_{\text{swp}}}{2}}. \quad (\text{B.70})$$

Introducing $\beta_d = 1/\alpha_d$, this turns into a plain arithmetic recurrence

$$\beta_d = \beta_{d-1} + 1 - \frac{\eta_{\text{swp}}}{2}. \quad (\text{B.71})$$

Using the initial condition $\beta_1 = 1/\alpha_1 = (4 - \eta_{\text{swp}})/2 = 2 - \eta_{\text{swp}}/2$ derived in the initialization step, the arithmetic sequence sums to

$$\beta_d = \beta_1 + (d - 1) \left(1 - \frac{\eta_{\text{swp}}}{2} \right) = (d + 1) - \frac{d \eta_{\text{swp}}}{2}. \quad (\text{B.72})$$

Inverting back yields the closed-form Bell weight at every depth d

$$\alpha_d = \frac{1}{(d + 1) - \frac{d \eta_{\text{swp}}}{2}} = \frac{2}{2(d + 1) - d \eta_{\text{swp}}}, \quad (\text{B.73})$$

which is the result quoted in equation (1.64). Note that the Bell weight decays *polynomially* as $\mathcal{O}(1/d)$ with the chain depth, reflecting the linear accumulation of vacuum admixture across successive swaps, rather than the exponential decay one might naively expect.

C -Study of the full repeater link

C.1 .Distribution time

This appendix computes the time T required to entangle Alice's and Bob's ions in the context of dual-rail repeater link with 2-links chain. Let L_{Lnk} denote the length of a single link, c the speed of light in the fiber, and M the number of modes. Assuming the heralding classical feed-back constitutes the bottleneck of the repeater click rate¹, the system attempts to distribute entanglement over a single repeater every² $\nu = M \cdot c/L_{\text{Lnk}}$. First, let's define the following events:

- A : the two elementary links of a chain have been loaded
- B : both chains have been loaded
- C : the four-click heralding has been detected.

The ion rate is then expressed as

$$\begin{aligned}\nu_{\text{ion}} &= \nu \cdot \mathbb{P}(A \cap B \cap C) \\ &= \nu \cdot \mathbb{P}(C | B) \cdot \mathbb{P}(B | A) \cdot \mathbb{P}(A).\end{aligned}\tag{C.1}$$

At the repeater level, event A is true when both elementary links of a single chain herald. Since the heralding of each elementary link occurs in parallel, the probability of loading both links of a chain reads

$$\mathbb{P}(A) = \frac{2}{3} \cdot \mathbb{P}_{\text{Lnk}},\tag{C.2}$$

where $\mathbb{P}_{\text{Lnk}}$ is the heralding probability of a single elementary link, and the $2/3$ factor stands as an effective loading time from the parallelization [13], called *synchronisation factor* in the manuscript. Similarly, since entanglement propagation in each chain occurs in parallel, the probability of loading both chains is

$$\mathbb{P}(B | A) = \frac{2}{3} \cdot \mathbb{P}_{\text{chn}}.\tag{C.3}$$

Finally, the probability of event C given B is simply $\mathbb{P}(C | B) = \mathbb{P}_{\text{ion}}$.

¹This assumes $c/L_{\text{Lnk}} > \nu_{\text{SPDC}}$, where ν_{SPDC} is the pair generation rate of an SPDC source.

²The total length of the repeater is L_{Lnk} , so the light travels a distance $L_{\text{Lnk}}/2$, an additional $L_{\text{Lnk}}/2$ is then needed for the heralding signal to return, hence a waiting time of L_{Lnk}/c .

The total rate is therefore

$$\nu_{\text{link}} = \nu \cdot \left(\frac{2}{3}\right)^2 \cdot \mathbb{P}_{\text{lon}} \cdot \mathbb{P}_{\text{chn}} \cdot \mathbb{P}_{\text{Lnk}}, \quad (\text{C.4})$$

leading to the distribution time

$$T = \left(\frac{3}{2}\right)^2 \cdot \frac{1}{M} \cdot \frac{L_{\text{Lnk}}}{c} \cdot \frac{1}{\mathbb{P}_{\text{lon}} \cdot \mathbb{P}_{\text{chn}} \cdot \mathbb{P}_{\text{Lnk}}}. \quad (\text{C.5})$$

C.2 .Strategy computation for the 4-clicks protocol

This appendix details the strategy used to compute the heralding probability $\mathbb{P}_{\text{lon}}$ and the fidelity $\mathcal{F}_{\text{lon}}$ for the 4-clicks protocol. We assume the two chains are loaded with the state ρ_{chn} , expressed as a statistical mixture of orthonormal pure states $\rho_i = |\psi_i\rangle \langle \psi_i|$. Using the linearity of the trace, we show that $\mathcal{F}_{\text{lon}}$ and $\mathbb{P}_{\text{lon}}$ reduce to the computation of the fidelities $\mathcal{F}_{\text{lon}}^{(i,j)}$ and probabilities $\mathbb{P}_{\text{lon}}^{(i,j)}$ attached to each pair of pure states $\rho_i \otimes \rho_j$. The fidelity and probability of the mixture are then derived as weighted averages of these elementary contributions. To emphasize the role of the two-photon corrections, we finally expand both quantities in powers of the squeezing parameter p .

Linearity

First, we recall that the final state shared by Alice's and Bob's ions reads

$$\rho_{\text{lon}} = \frac{1}{\mathbb{P}_{\text{lon}}} \cdot \text{Tr}_{\text{in}} \left[\sqrt{\Pi_{\text{lon}}} U_{\text{BS}} \rho_{\text{chn}} \otimes \rho_{\text{chn}} U_{\text{BS}}^\dagger \sqrt{\Pi_{\text{lon}}} \right], \quad (\text{C.6})$$

where U_{BS} is the beam-splitter unitary, Π_{lon} the POVM associated with the four single-click heralding event, $\text{Tr}_{\text{in}}[\cdot]$ the partial trace over the inside modes, and $\mathbb{P}_{\text{lon}}$ the normalisation constant given by the total heralding probability. Given the decomposition

$$\rho_{\text{chn}} = \sum_i a_i \rho_i, \quad (\text{C.7})$$

the fidelity of the final state ρ_{lon} expands as

$$\begin{aligned} \mathcal{F}_{\text{lon}} &= \langle \Psi^+ | \rho_{\text{lon}} | \Psi^+ \rangle \\ &= \frac{1}{\mathbb{P}_{\text{lon}}} \langle \Psi^+ | \text{Tr}_{\text{in}} \left[\Pi_{\text{lon}} U_{\text{BS}} \rho_{\text{chn}} \otimes \rho_{\text{chn}} U_{\text{BS}}^\dagger \right] | \Psi^+ \rangle \\ &= \frac{1}{\mathbb{P}_{\text{lon}}} \sum_{i,j} a_i a_j \cdot \underbrace{\langle \Psi^+ | \text{Tr}_{\text{in}} \left[\Pi_{\text{lon}} U_{\text{BS}} \rho_i \otimes \rho_j U_{\text{BS}}^\dagger \right] | \Psi^+ \rangle}_{=\mathbb{P}_{\text{lon}}^{(i,j)} \cdot \rho_{\text{lon}}^{(i,j)}}. \end{aligned} \quad (\text{C.8})$$

Denoting by $\mathcal{F}_{\text{lon}}^{(i,j)}$ the fidelity of the pairwise state $\rho_{\text{lon}}^{(i,j)}$

$$\mathcal{F}_{\text{lon}}^{(i,j)} = \langle \Psi^+ | \rho_{\text{lon}}^{(i,j)} | \Psi^+ \rangle = \frac{\langle \Psi^+ | \text{Tr}_{\text{in}} \left[\Pi_{\text{lon}} U_{\text{BS}} \rho_i \otimes \rho_j U_{\text{BS}}^\dagger \right] | \Psi^+ \rangle}{\mathbb{P}_{\text{lon}}^{(i,j)}}, \quad (\text{C.9})$$

with the pairwise heralding probability

$$\mathbb{P}_{\text{lon}}^{(i,j)} = \text{Tr}_{\text{tot}} \left[\Pi_{\text{lon}} U_{\text{BS}} \rho_i \otimes \rho_j U_{\text{BS}}^\dagger \right], \quad (\text{C.10})$$

the total fidelity reads

$$\mathcal{F}_{\text{lon}} = \sum_{i,j} a_i a_j \cdot \frac{\mathbb{P}_{\text{lon}}^{(i,j)} \mathcal{F}_{\text{lon}}^{(i,j)}}{\mathbb{P}_{\text{lon}}}. \quad (\text{C.11})$$

Similarly, the total heralding probability is the weighted sum

$$\mathbb{P}_{\text{lon}} = \sum_{i,j} a_i a_j \cdot \mathbb{P}_{\text{lon}}^{(i,j)}. \quad (\text{C.12})$$

Thus, $\mathcal{F}_{\text{lon}}$ and $\mathbb{P}_{\text{lon}}$ are entirely determined by the elementary contributions $\mathbb{P}_{\text{lon}}^{(i,j)}$ and $\mathcal{F}_{\text{lon}}^{(i,j)}$ attached to each pair $\rho_i \otimes \rho_j$.

Expansion

Expanding $\mathcal{F}_{\text{lon}}$ and $\mathbb{P}_{\text{lon}}$ in powers of the squeezing parameter p , the coefficients a_i read $a_i = \alpha_i + \beta_i \cdot p + \mathcal{O}(p)$, yielding

$$a_i a_j = \alpha_i \alpha_j + (\alpha_i \beta_j + \alpha_j \beta_i) \cdot p + \mathcal{O}(p). \quad (\text{C.13})$$

The numerator of $\mathcal{F}_{\text{lon}}$ then expands as

$$\begin{aligned} \mathbb{P}_{\text{lon}} \cdot \mathcal{F}_{\text{lon}} &= \sum_{i,j} \alpha_i \alpha_j \cdot \mathbb{P}_{\text{lon}}^{(i,j)} \mathcal{F}_{\text{lon}}^{(i,j)} + p \sum_{i,j} (\alpha_i \beta_j + \alpha_j \beta_i) \cdot \mathbb{P}_{\text{lon}}^{(i,j)} \mathcal{F}_{\text{lon}}^{(i,j)} + \mathcal{O}(p) \\ &= \mathcal{F}_{\text{lon}}^{(0)} + \mathcal{F}_{\text{lon}}^{(1)} \cdot p + \mathcal{O}(p), \end{aligned} \quad (\text{C.14})$$

where the barred quantities $\mathcal{F}_{\text{lon}}^{(k)}$ denote the bare expansion coefficients of the numerator $\mathbb{P}_{\text{lon}} \cdot \mathcal{F}_{\text{lon}}$. Note that these coefficients do not coincide with the expansion of $\mathcal{F}_{\text{lon}}$ itself, since the normalisation $\mathbb{P}_{\text{lon}}$ also depends on p . This effect has been previously discussed and corresponds to the dilution correction. Expanding $\mathbb{P}_{\text{lon}} = \mathbb{P}_{\text{lon}}^{(0)} + \mathbb{P}_{\text{lon}}^{(1)} \cdot p + \mathcal{O}(p)$ similarly, the fidelity becomes

$$\mathcal{F}_{\text{lon}} = \frac{\mathcal{F}_{\text{lon}}^{(0)}}{\mathbb{P}_{\text{lon}}^{(0)}} + \left(\frac{\mathcal{F}_{\text{lon}}^{(1)}}{\mathbb{P}_{\text{lon}}^{(0)}} - \frac{\mathcal{F}_{\text{lon}}^{(0)} \mathbb{P}_{\text{lon}}^{(1)}}{(\mathbb{P}_{\text{lon}}^{(0)})^2} \right) \cdot p + \mathcal{O}(p), \quad (\text{C.15})$$

and the probability

$$\mathbb{P}_{\text{lon}} = \mathbb{P}_{\text{lon}}^{(0)} + \mathbb{P}_{\text{lon}}^{(1)} \cdot p + \mathcal{O}(p). \quad (\text{C.16})$$

Final result Putting everything together, the heralding probability and fidelity of the 4-clicks protocol read, up to first order in the squeezing parameter,

$$\boxed{\mathbb{P}_{\text{lon}} = \mathbb{P}_{\text{lon}}^{(0)} + \mathbb{P}_{\text{lon}}^{(1)} \cdot p + \mathcal{O}(p)}, \quad (\text{C.17})$$

$$\mathcal{F}_{\text{lon}} = \frac{\mathcal{F}_{\text{lon}}^{(0)}}{\mathbb{P}_{\text{lon}}^{(0)}} + \left(\frac{\mathcal{F}_{\text{lon}}^{(1)}}{\mathbb{P}_{\text{lon}}^{(0)}} - \frac{\mathcal{F}_{\text{lon}}^{(0)} \mathbb{P}_{\text{lon}}^{(1)}}{(\mathbb{P}_{\text{lon}}^{(0)})^2} \right) \cdot p + \mathbf{o}(p), \quad (\text{C.18})$$

with the expansion coefficients

$$\begin{aligned} \mathcal{F}_{\text{lon}}^{(0)} &= \sum_{i,j} \alpha_i \alpha_j \cdot \mathbb{P}_{\text{lon}}^{(i,j)} \mathcal{F}_{\text{lon}}^{(i,j)}, & \mathbb{P}_{\text{lon}}^{(0)} &= \sum_{i,j} \alpha_i \alpha_j \cdot \mathbb{P}_{\text{lon}}^{(i,j)}, \\ \mathcal{F}_{\text{lon}}^{(1)} &= \sum_{i,j} (\alpha_i \beta_j + \alpha_j \beta_i) \cdot \mathbb{P}_{\text{lon}}^{(i,j)} \mathcal{F}_{\text{lon}}^{(i,j)}, & \mathbb{P}_{\text{lon}}^{(1)} &= \sum_{i,j} (\alpha_i \beta_j + \alpha_j \beta_i) \cdot \mathbb{P}_{\text{lon}}^{(i,j)}. \end{aligned} \quad (\text{C.19})$$

Note that we recover the general results of section 1.1.2 for the fidelity, with an explicit expression of the 2-photon contribution and its associated dilution correction:

$$\tilde{\mathcal{F}}_{\text{lon}}^{(1)} = \frac{\mathcal{F}_{\text{lon}}^{(1)}}{\mathbb{P}_{\text{lon}}^{(1)}}, \quad \tilde{\mathcal{F}}_{\text{lon}}^{(2)} = \frac{\mathcal{F}_{\text{lon}}^{(2)}}{\mathbb{P}_{\text{lon}}^{(1)}} - \frac{\mathcal{F}_{\text{lon}}^{(1)}}{\mathbb{P}_{\text{lon}}^{(1)}} \cdot \frac{\mathbb{P}_{\text{lon}}^{(2)}}{\mathbb{P}_{\text{lon}}^{(1)}}. \quad (\text{C.20})$$

The probability, on the contrary, does not carry any dilution factor, this is a direct consequence of the assumption that the initial state from the repeater chains is already normalized

$$\tilde{\mathbb{P}}_{\text{lon}}^{(1)} = \mathbb{P}_{\text{lon}}^{(1)}, \quad \tilde{\mathbb{P}}_{\text{lon}}^{(2)} = \mathbb{P}_{\text{lon}}^{(2)}. \quad (\text{C.21})$$

Also note that the zeroth-order coefficient of the fidelity is unity by construction: $\tilde{\mathcal{F}}_{\text{lon}}^{(1)} = \mathcal{F}_{\text{lon}}^{(1)} / \mathbb{P}_{\text{lon}}^{(1)} = 1$. This approach reduces the computation of $\mathcal{F}_{\text{lon}}$ and $\mathbb{P}_{\text{lon}}$ to the evaluation of the elementary contributions $\mathbb{P}_{\text{lon}}^{(i,j)}$ and $\mathcal{F}_{\text{lon}}^{(i,j)}$ attached to each pure-state pair $\rho_i \otimes \rho_j$, for which a closed-form expression can be derived directly from the beam-splitter and heralding-measurement actions.

C.3 .Modelling temporal-mode mismatch

This appendix derives general results to model partial photon-shape overlap. Concretely the beam-splitter relation between two photons of mismatched shapes, and the shape-independent detection POVM.

Beam-Splitter relation

The standard beam-splitter relation describing the interference of two photons assumes that both photons share the same temporal shape. Denoting the spatio-temporal modes of the two photons $a^\dagger$ and $b^\dagger$, this assumption means their modes coincide at the interference plane, i.e. $a^\dagger = b^\dagger$ at that time, so the photons are temporally identical. In this appendix we study the case where two photons with mismatched shapes interfere at a beam-splitter.

Since photons with orthogonal shapes cannot interfere, we decompose $b^\dagger$ into two components³

- an *interfering component* $b_0^\dagger$, aligned with the shape of $a^\dagger$, the *reference mode*,
- a *non-interfering component* $b_\perp^\dagger$, orthogonal to $a^\dagger$, the *orthogonal mode*.

This decomposition splits the shape of $b^\dagger$ into a part matching the shape of $a^\dagger$, namely $b_0^\dagger$, and a residual part orthogonal to $a^\dagger$, namely $b_\perp^\dagger$. Algebraically, it amounts to projecting $b^\dagger$ orthogonally onto $a^\dagger$ to get $b_0^\dagger$, then reconstructing $b^\dagger$ from $b_0^\dagger$ and $b_\perp^\dagger$ gives

$$\begin{aligned} a^\dagger &= a_0^\dagger, \\ b^\dagger &= \sqrt{x} \cdot b_0^\dagger + \sqrt{1-x} \cdot b_\perp^\dagger, \end{aligned} \tag{C.22}$$

where $x = |\langle 0 | a b^\dagger | 0 \rangle|^2$ is the overlap intensity between the photons, $x = 1$ when the photons have identical shapes and fully interfere, and $x = 0$ when they are orthogonal and do not interfere. Using the standard beam-splitter relations, the transformation becomes

$$\begin{aligned} a^\dagger &\xrightarrow{\text{BS}} \frac{a_0^\dagger + b_0^\dagger}{\sqrt{2}}, \\ b^\dagger &\xrightarrow{\text{BS}} \sqrt{x} \cdot \frac{a_0^\dagger - b_0^\dagger}{\sqrt{2}} + \sqrt{1-x} \cdot \frac{a_\perp^\dagger - b_\perp^\dagger}{\sqrt{2}}. \end{aligned} \tag{C.23}$$

³Although $a^\dagger$ and $b^\dagger$ differ in shape, they are not necessarily orthogonal, which motivates this decomposition.

POVM element

Let us consider a non-resolving photon detector with efficiency η and no dark counts. The POVMs respectively corresponding to a *click* and *non-click* event in a mode $a^\dagger$ are given by:

$$\begin{aligned}\Pi_{nc}^a(\eta) &= R^{a^\dagger a}, \\ \Pi_c^a(\eta) &= \mathbb{1} - R^{a^\dagger a},\end{aligned}\tag{C.24}$$

where $R = 1 - \eta$ is the losses associated with detection efficiency η . Now we consider a scenario with two input modes: the reference mode $a_0^\dagger$ and the orthogonal mode $a_n^\dagger$. Since photons in $a_0^\dagger$ and $a_n^\dagger$ have orthogonal shapes, we assume distinct detection efficiencies η_0 and η_n for these modes. Thus, the POVM for the overall mode $a^\dagger$ can be written as $\Pi_0^a = R_0^{a_0^\dagger a_0} \cdot R_n^{a_n^\dagger a_n}$, resulting in a probability $R_0^{N_0} \cdot R_n^{N_n}$ for a state $|N_0, N_n\rangle$, composed of N_0 photons in $a_0^\dagger$ and N_n photons in $a_n^\dagger$. Resulting expressions for the POVMs

$$\begin{aligned}\text{Non-click event:} \quad \Pi_{nc}^a &= R_0^{a_0^\dagger a_0} \cdot R_n^{a_n^\dagger a_n}, \\ \text{Click event:} \quad \Pi_c^a &= \mathbb{1} - R_0^{a_0^\dagger a_0} \cdot R_n^{a_n^\dagger a_n}.\end{aligned}\tag{C.25}$$

Shaping and Beam-Splitter

Calling x and x the photon shaping mismatch respectively at Alice and Bob interface, the beam-splitter relation leads to

$$\begin{aligned}
 1^{\text{st}} \text{ B.S.: } & \begin{cases} \tilde{b}^\dagger \longrightarrow \sqrt{x} \cdot \frac{b_0^\dagger + d_0^\dagger}{\sqrt{2}} + \sqrt{1-x} \cdot \frac{b_\perp^\dagger + d_\perp^\dagger}{\sqrt{2}} \\ \tilde{d}^\dagger \longrightarrow \frac{b_0^\dagger - d_0^\dagger}{\sqrt{2}} \end{cases} \\
 2^{\text{nd}} \text{ B.S.: } & \begin{cases} \tilde{c}^\dagger \longrightarrow \sqrt{x} \cdot \frac{c_0^\dagger + f_0^\dagger}{\sqrt{2}} + \sqrt{1-x} \cdot \frac{c_\perp^\dagger + f_\perp^\dagger}{\sqrt{2}} \\ \tilde{f}^\dagger \longrightarrow \frac{c_0^\dagger - f_0^\dagger}{\sqrt{2}} \end{cases} \\
 3^{\text{rd}} \text{ B.S.: } & \begin{cases} \tilde{h}^\dagger \longrightarrow \sqrt{x} \cdot \frac{h_0^\dagger + e_0^\dagger}{\sqrt{2}} + \sqrt{1-x} \cdot \frac{h_\perp^\dagger + e_\perp^\dagger}{\sqrt{2}} \\ \tilde{e}^\dagger \longrightarrow \frac{h_0^\dagger - e_0^\dagger}{\sqrt{2}} \end{cases} \\
 4^{\text{th}} \text{ B.S.: } & \begin{cases} \tilde{i}^\dagger \longrightarrow \sqrt{x} \cdot \frac{i_0^\dagger + g_0^\dagger}{\sqrt{2}} + \sqrt{1-x} \cdot \frac{i_\perp^\dagger + g_\perp^\dagger}{\sqrt{2}} \\ \tilde{g}^\dagger \longrightarrow \frac{i_0^\dagger - g_0^\dagger}{\sqrt{2}} \end{cases}
 \end{aligned} \tag{C.27}$$

Note we use the atomic ensemble shape as reference, for clarity the reference modes $\{a_0, b_0, \dots, i_0, j_0\}$ will simply be written $\{a, b, \dots, i, j\}$. The BS transformation leads to the state

$$\begin{aligned}
 |\psi\rangle = & \frac{1}{4} \left[\sqrt{x} \left(\frac{b+d}{\sqrt{2}} + a \cdot \frac{c+f}{\sqrt{2}} \right) + \sqrt{1-x} \left(\frac{b_\perp + d_\perp}{\sqrt{2}} + a \cdot \frac{c_\perp + f_\perp}{\sqrt{2}} \right) \right] \dots \\
 & \times \left[\frac{b-d}{\sqrt{2}} + \frac{h-e}{\sqrt{2}} \right] \dots \\
 & \times \left[\frac{c-f}{\sqrt{2}} + \frac{i-g}{\sqrt{2}} \right] \dots \\
 & \times \left[\sqrt{x} \left(\frac{h+e}{\sqrt{2}} + j \cdot \frac{i+g}{\sqrt{2}} \right) + \sqrt{1-x} \left(\frac{h_\perp + e_\perp}{\sqrt{2}} + j \cdot \frac{i_\perp + g_\perp}{\sqrt{2}} \right) \right] |0\rangle
 \end{aligned} \tag{C.28}$$

Detection

We recall the expression of the POVM element associated to the ion swap heralding is

$$\Pi_{\text{Ion}} = \prod_{(\alpha, \beta) \in \mathcal{C}} \Pi_1^\alpha(\eta_{\text{Ion}}) \Pi_0^\beta(\eta_{\text{Ion}}), \tag{C.29}$$

with η_{lon} the heralding efficiency and $\mathcal{C} = \{(b, d), (c, f), (h, e), (g, i)\}$ the set of click/no-click mode pairs. At first glance, the action of Π_{lon} on $|\psi\rangle$ appears non-trivial: Π_{lon} acts on eight distinct modes, while $|\psi\rangle$ consists of $8 \times 4 \times 4 \times 8 = 1024$ terms. However, $\Pi_{\text{lon}} |\psi\rangle$ contains mainly null terms, which significantly simplifying the calculation. Specifically, for $i = 0, 1$, the action of Π_0 and Π_1 on $|i\rangle$ yields

$$\begin{aligned} \text{Reference mode:} \quad & \Pi_1 |i\rangle = 1 - (1 - \eta_{\text{lon}})^i, \quad \Pi_0 |i\rangle = (1 - \eta_{\text{lon}})^i, \\ \text{Orthogonal mode:} \quad & \Pi_1 |i\rangle = 1 - (1 - \eta_{\perp})^i, \quad \Pi_0 |i\rangle = (1 - \eta_{\perp})^i. \end{aligned} \quad (\text{C.30})$$

This a direct consequence of neglecting dark counts. Thus, since the POVM element Π_{lon} clicks for the modes $b^\dagger$, $c^\dagger$, $h^\dagger$, and $i^\dagger$, the non-zero terms in $\Pi_{\text{lon}} |\psi\rangle$ are the states storing one excitation in each of these modes, or in their respective orthogonal modes. For clarity, here are few examples of terms that yield zero and non-zero outcomes when acted upon by Π_{lon} :

$$\begin{aligned} \Pi_{\text{lon}} \cdot b^\dagger c^\dagger h^\dagger i^\dagger & \longrightarrow \eta_{\text{lon}}^4 \cdot b^\dagger c^\dagger h^\dagger i^\dagger & \text{All modes } b, c, h, i \text{ are present} \\ & & \text{(or resp. } b_{\perp}, c_{\perp}, h_{\perp}, i_{\perp}) \\ \Pi_{\text{lon}} \cdot b^\dagger c_{\perp}^\dagger h^\dagger i_{\perp}^\dagger & \longrightarrow \eta_{\text{lon}}^2 \eta_{\perp}^2 \cdot b^\dagger c_{\perp}^\dagger h^\dagger i_{\perp}^\dagger & \text{All modes } b, c, h, i \text{ are present} \\ & & \text{(or resp. } b_{\perp}, c_{\perp}, h_{\perp}, i_{\perp}) \\ \Pi_{\text{lon}} \cdot b_{\perp}^\dagger h^\dagger i^\dagger & \longrightarrow 0 & \text{Mode } c \text{ (or } c_{\perp}) \text{ is missing} \\ \Pi_{\text{lon}} \cdot a^\dagger b^\dagger c_{\perp}^\dagger d^\dagger f^\dagger i_{\perp}^\dagger & \longrightarrow 0 & \text{Mode } h \text{ (or } h_{\perp}) \text{ is missing} \end{aligned} \quad (\text{C.31})$$

Identifying these zero terms reduces the number of non-zero vectors in $\Pi_{\text{lon}} |\psi\rangle$ from 1024 down to 8. Yet, all 1024 vectors must in principle be evaluated to identify the surviving ones. The method outlined below provides a systematic approach to this enumeration.

Non-Zero factors

Since equation (C.28) is composed of four factors, a convenient strategy is to track the modes b , c , h , and i from each factor and reconstruct the four-mode product $b^\dagger c^\dagger h^\dagger i^\dagger$, or its orthogonal counterparts. Listing all admissible modes in each factor: $\{b, d, b_{\perp}, d_{\perp}\}$ in

the first, $\{b, h\}$ in the second, $\{c, i\}$ in the third, and $\{h, i, h_\perp, i_\perp\}$ in the fourth,

$$\begin{aligned}
& \left[\sqrt{x} \left(\frac{\boxed{b} + d}{\sqrt{2}} + a \cdot \frac{\boxed{c} + f}{\sqrt{2}} \right) + \sqrt{1-x} \left(\frac{\boxed{b_\perp} + d_\perp}{\sqrt{2}} + a \cdot \frac{\boxed{c_\perp} + f_\perp}{\sqrt{2}} \right) \right] \longleftarrow b, d, b_\perp, d_\perp, \\
& \left[\frac{\boxed{b} - d}{\sqrt{2}} + \frac{\boxed{h} - e}{\sqrt{2}} \right] \longleftarrow b, h, \\
& \left[\frac{\boxed{c} - f}{\sqrt{2}} + \frac{\boxed{i} - g}{\sqrt{2}} \right] \longleftarrow c, i, \\
& \left[\sqrt{x} \left(\frac{\boxed{h} + e}{\sqrt{2}} + j \cdot \frac{\boxed{i} + g}{\sqrt{2}} \right) + \sqrt{1-x} \left(\frac{\boxed{h_\perp} + e_\perp}{\sqrt{2}} + j \cdot \frac{\boxed{i_\perp} + g_\perp}{\sqrt{2}} \right) \right] \longleftarrow h, i, h_\perp, i_\perp.
\end{aligned} \tag{C.32}$$

For instance, one builds a non-zero term by picking b from the first factor, h from the second, c from the third, and $i \cdot j$ from the fourth, yielding the mode combination $b^\dagger c^\dagger h^\dagger i^\dagger j^\dagger$, with coefficient $\frac{x}{16}$ and POVM term η_{ion}^4 . Iterating this procedure enumerates the eight combinations of non-zero terms, along with their coefficients and POVM contributions. Table C.1 summarises the eight resulting terms. The column *factor* lists the mode combinations built from the four factors as described above, the column *coefficient* gives the weight of the associated vector in $|\psi\rangle$, the column *POVM* reports the eigenvalue of Π_{ion} on that vector, and the column *ion* indicates the ion mode involved in the combination. For instance, the previous example $\Pi_{\text{ion}} \cdot \frac{x}{16} \cdot b^\dagger c^\dagger h^\dagger i^\dagger j^\dagger |0\rangle = \eta_{\text{ion}}^4 \cdot \frac{x}{16} \cdot b^\dagger c^\dagger h^\dagger i^\dagger j^\dagger |0\rangle$ corresponds to η_{ion}^4 for the POVM term, $\frac{x}{16}$ for the coefficient, $b^\dagger c^\dagger h^\dagger i^\dagger$ for the combination, and $j^\dagger$ for the ion.

Trace and state

At this point we have identified all the non-zero vectors of $\Pi_{\text{ion}} |\psi\rangle$ with their associated coefficient, POVM terms, and ion's mode involve. The last step of the calculation aims compute the partial trace $\text{Tr}_{\text{in}} [\Pi_{\text{ion}} |\psi\rangle \langle \psi|]$. As the trace Tr_{in} is over the *inside modes* $\{b, c \cdots h, i\}$ and $\{b_\perp, c_\perp \cdots h_\perp, i_\perp\}$, we need to factorize the sum regarding those modes, otherwise we might miss cross terms. The *ion* column of table C.1 shows that only the 1st and 5th non-zero terms involve the same inside modes (the modes c, b, i, h), where the others terms involve distinct inside modes. Finally, squaring the coefficients and spotting

	1 st factor	2 nd factor	3 rd factor	4 th factor	Coefficient	POVM term	Ion involved
	$b, c, b_{\perp}, c_{\perp}$	b, h	c, i	$h, i, h_b, i_{\perp}$			
1 st	b	h	c	i	$\frac{1}{16} \cdot x$	η_{ion}^4	Mode j
2 nd	b	h	c	$i_{\perp}$	$\frac{1}{16} \cdot \sqrt{x(1-x)}$	$\eta_{\text{ion}}^3 \cdot \eta_{\text{ion}}^{\perp}$	Mode j
3 rd	$b_{\perp}$	h	c	i	$\frac{1}{16} \cdot \sqrt{x(1-x)}$	$\eta_{\text{ion}}^3 \cdot \eta_{\text{ion}}^{\perp}$	Mode j
4 th	$b_{\perp}$	h	c	$i_{\perp}$	$\frac{1}{16} \cdot (1-x)$	$\eta_{\text{ion}}^2 \cdot (\eta_{\text{ion}}^{\perp})^2$	Mode j
5 th	c	b	i	h	$\frac{1}{16} \cdot x$	η_{ion}^4	Mode a
6 th	c	b	i	$h_{\perp}$	$\frac{1}{16} \cdot \sqrt{x(1-x)}$	$\eta_{\text{ion}}^3 \cdot \eta_{\text{ion}}^{\perp}$	Mode a
7 th	$c_{\perp}$	b	i	h	$\frac{1}{16} \cdot \sqrt{x(1-x)}$	$\eta_{\text{ion}}^3 \cdot \eta_{\text{ion}}^{\perp}$	Mode a
8 th	$c_{\perp}$	b	i	$h_{\perp}$	$\frac{1}{16} \cdot (1-x)$	$\eta_{\text{ion}}^2 \cdot (\eta_{\text{ion}}^{\perp})^2$	Mode a

Table C.1: Summary of all the 8 combinations built from the modes b, c, h, i or their respective orthogonal modes, representing the non-zero terms of $\Pi|\Psi\rangle$. The *coefficient column* corresponds to the associated coefficient of $|\Psi\rangle$, the *POVM term column* shows the eigenvalue of Π from the associated vector, and the *ion involve column* shows the ion's mode involve in the combination.

the ion's mode involve for each terms, we get the matrix terms

$$\begin{aligned}
\eta_{\text{lon}}^4 \cdot \frac{x^2}{156} \cdot (a+j) \cdot b h c i & \longleftarrow (|e, g\rangle + |g, e\rangle)(\langle e, g| + \langle g, e|) \\
\eta_{\text{lon}}^3 \eta_{\text{lon}}^\perp \cdot \frac{x(1-x)}{256} \cdot j \cdot b h c i_\perp & \longleftarrow |g, e\rangle \langle g, e| \\
\eta_{\text{lon}}^3 \eta_{\text{lon}}^\perp \cdot \frac{x(1-x)}{256} \cdot j \cdot b_\perp h c i & \longleftarrow |g, e\rangle \langle g, e| \\
\eta_{\text{lon}}^2 (\eta_{\text{lon}}^\perp)^2 \cdot \frac{(1-x)^2}{256} \cdot j \cdot b_\perp h c i_\perp & \longleftarrow |g, e\rangle \langle g, e| \\
\eta_{\text{lon}}^3 \eta_{\text{lon}}^\perp \cdot \frac{x(1-x)}{256} \cdot a \cdot c b i h_\perp & \longleftarrow |e, g\rangle \langle e, g| \\
\eta_{\text{lon}}^3 \eta_{\text{lon}}^\perp \cdot \frac{x(1-x)}{256} \cdot a \cdot c_\perp b i h & \longleftarrow |e, g\rangle \langle e, g| \\
\eta_{\text{lon}}^2 (\eta_{\text{lon}}^\perp)^2 \cdot \frac{(1-x)^2}{256} \cdot a \cdot c_\perp b i h_\perp & \longleftarrow |e, g\rangle \langle e, g|
\end{aligned}$$

Leading to the partial trace

$$\begin{aligned}
\text{Tr}_{\text{in}} [\Pi_{\text{lon}} \cdot |\psi\rangle \langle \psi|] &= \frac{\eta_{\text{lon}}^4 x^2 + 2\eta_{\text{lon}}^3 \eta_{\text{lon}}^\perp x(1-x) + \eta_{\text{lon}}^2 (\eta_{\text{lon}}^\perp)^2 (1-x)^2}{256} \cdot (|01\rangle \langle 01| + |10\rangle \langle 10|) \dots \\
&+ \frac{\eta_{\text{lon}}^4 x^2}{256} \cdot (|01\rangle \langle 10| + |10\rangle \langle 01|)
\end{aligned} \tag{C.33}$$

calling $\delta = \eta_{\text{lon}} - \eta_{\text{lon}}^\perp$ the mismatch between the reference mode and orthogonal mode, we notice the equality $\eta_{\text{lon}}^4 x + 2\eta_{\text{lon}}^3 \eta_{\text{lon}}^\perp x(1-x) + \eta_{\text{lon}}^2 (\eta_{\text{lon}}^\perp)^2 (1-x)^2 = \eta_{\text{lon}}^2 \cdot (\eta_{\text{lon}}^\perp + x \cdot \delta)^2$, leading to a the lon-lon density matrix

$$\rho_{\text{lon}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1/2 & \frac{\eta_{\text{lon}}^2 x^2}{2(\eta_{\text{lon}}^\perp + x \cdot \delta)^2} & 0 \\ 0 & \frac{\eta_{\text{lon}}^2 x^2}{2(\eta_{\text{lon}}^\perp + x \cdot \delta)^2} & 1/2 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \tag{C.34}$$

Probability and fidelity

We recall the POVM element Π_{lon} actually represents on 1 of the 16 possible event. We also recall the initial state from the ions and links chain are actually loaded with a vacuum component, which bring a dilution factor α_{chn}^2 and α_{lon}^2 to the final results. Finally the ion swap probability event reads

$$\mathbb{P}_{\text{lon}} = \alpha_{\text{chn}}^2 \alpha_{\text{lon}}^2 \frac{\eta_{\text{lon}}^2 \cdot (\eta_{\text{lon}}^\perp + x \cdot \delta)^2}{8}, \tag{C.35}$$

which heralds the state

$$\begin{aligned} \rho_{\text{ion}} = & \left(\frac{1}{2} + \frac{\eta_{\text{ion}}^2 \cdot x^2}{2 \cdot (\eta_{\text{ion}}^\perp + x \cdot \delta)^2} \right) \cdot |\Psi_+\rangle \langle \Psi_+| \\ & + \left(\frac{1}{2} - \frac{\eta_{\text{ion}}^2 \cdot x^2}{2 \cdot (\eta_{\text{ion}}^\perp + x \cdot \delta)^2} \right) \cdot |\Psi_-\rangle \langle \Psi_-|, \end{aligned} \quad (\text{C.36})$$

leading to the fidelity

$$\mathcal{F}_{\text{ion}} = \frac{1}{2} + \frac{\eta_{\text{ion}}^2 \cdot x^2}{2 \cdot (\eta_{\text{ion}}^\perp + x \cdot \delta)^2}. \quad (\text{C.37})$$

C.5 .Shape-independent construction of the photodetection POVM

This appendix constructs a POVM describing the click and no-click outcomes of a non-resolving photodetector. The construction holds for any continuous family of non-orthogonal spatio-temporal modes $\{a_\gamma^\dagger\}$, regardless of the shape γ of the incident excitation. Let's introduce an orthonormal basis $\{\alpha_\nu\}_\nu$ satisfying

$$[\alpha_\nu, \alpha_{\tilde{\nu}}^\dagger] = \delta(\nu - \tilde{\nu}) \mathbb{1}. \quad (\text{C.38})$$

Kets and operators written in the basis $\{\alpha_\nu\}_\nu$ are indexed by ν , while those written in the basis $\{a_\gamma\}_\gamma$ are indexed by γ . Thus, for any shape γ , the mode $a_\gamma^\dagger$ decomposed as

$$a_\gamma^\dagger = \int P_\gamma(\nu) \alpha_\nu^\dagger d\nu, \quad \text{with} \quad \int |P_\gamma(\nu)|^2 d\nu = 1. \quad (\text{C.39})$$

As we work in the single-excitation subspace, we look for a POVM element acting non-trivially on $\text{Span} \left\{ \alpha_\nu^\dagger |0\rangle \right\}_\nu$. Concretely, for any shape γ and any single-mode Fock state $|i\rangle_\gamma$ with $i \in \{0, 1\}$, we want

$$\Pi_0 |i\rangle_\gamma = (1 - \eta_{\text{lon}})^i |i\rangle_\gamma. \quad (\text{C.40})$$

We propose the candidate

$$\Pi_0 = R_{\text{lon}}^{\int \alpha_\nu^\dagger \alpha_\nu d\nu}, \quad (\text{C.41})$$

where $R_{\text{lon}} = 1 - \eta_{\text{lon}}$. Let's show such an operator satisfies the required action on the single-excitation subspace. Consider a state $|1\rangle_\gamma = a_\gamma^\dagger |0\rangle$ carrying one excitation in an arbitrary spatio-temporal mode $a_\gamma^\dagger = \int P_\gamma(\nu) \alpha_\nu^\dagger d\nu$. The action of Π_0 on this state reads

$$\begin{aligned} \Pi_0 |1\rangle_\gamma &= R_{\text{lon}}^{\int \alpha_\nu^\dagger \alpha_\nu d\nu} a_\gamma^\dagger |0\rangle \\ &= R_{\text{lon}}^{\int \alpha_\nu^\dagger \alpha_\nu d\nu} \int P_\gamma(\nu) \alpha_\nu^\dagger |0\rangle d\nu \\ &= \int P_\gamma(\nu) R_{\text{lon}}^{\int \alpha_{\nu'}^\dagger \alpha_{\nu'} d\nu'} |1\rangle_\nu d\nu \\ &= \int P_\gamma(\nu) R_{\text{lon}} |1\rangle_\nu d\nu \\ &= R_{\text{lon}} \int P_\gamma(\nu) \alpha_\nu^\dagger |0\rangle d\nu \\ &= R_{\text{lon}} |1\rangle_\gamma. \end{aligned} \quad (\text{C.42})$$

The third line uses the fact that $|1\rangle_\nu = \alpha_\nu^\dagger |0\rangle$ is an eigenstate of the photon-number operator $\int \alpha_{\nu'}^\dagger \alpha_{\nu'} d\nu'$ with eigenvalue 1, and therefore an eigenstate of Π_0 with eigenvalue R_{lon} . The vacuum $|0\rangle$ is similarly an eigenstate with eigenvalue 1. This shows that Π_0 acts diagonally on $\text{Span} \left\{ \alpha_\nu^\dagger |0\rangle \right\}_\nu$ with eigenvalue R_{lon} , regardless of the temporal shape γ of

the incident excitation. Finally, the pair

$$\begin{aligned} \text{No-click: } \Pi_0 &= R_{\text{lon}}^{\int \alpha_{\nu}^{\dagger} \alpha_{\nu} d\nu}, \\ \text{Click: } \Pi_1 &= \mathbb{1} - R_{\text{lon}}^{\int \alpha_{\nu}^{\dagger} \alpha_{\nu} d\nu}, \end{aligned} \tag{C.43}$$

forms a valid POVM describing the click and no-click outcomes of a non-resolving photodetector, valid for any spatio-temporal shape γ in the family $\{a_{\gamma}\}$.

D -Exact study of an elementary link

D.1 .Derivation of the state

The derivation of ρ_{Lnk} proceeds in four steps: the joint state of the two SPDC sources $\rho_{\text{SPDC}}^A \otimes \rho_{\text{SPDC}}^B$, the mode mixing on the beam-splitter U_{BS} , the action of the heralding POVM Π_{Lnk} at the central station, and the partial trace over the central modes $c_A^\dagger$ and $c_B^\dagger$:

$$\rho_{\text{Lnk}} = \frac{1}{\mathcal{N}} \text{Tr}_{c_A, c_B} \left[\sqrt{\Pi_{\text{Lnk}}} \cdot U_{\text{BS}} \cdot \rho_{\text{SPDC}}^A \otimes \rho_{\text{SPDC}}^B \cdot U_{\text{BS}}^\dagger \sqrt{\Pi_{\text{Lnk}}} \right], \quad (\text{D.1})$$

with $\mathcal{N}$ the normalized constant.

State from the SPDC sources:

The state generated by a single SPDC source on modes a, b reads

$$|\psi_{\text{SPDC}}^{a,b}\rangle = \sqrt{1-p} \cdot \exp\left(\sqrt{p} \cdot a^\dagger b^\dagger\right) |0_{a,b}\rangle, \quad (\text{D.2})$$

so that the joint state of Alice's and Bob's sources is

$$\begin{aligned} \rho_{\text{SPDC}} &= |\psi_{\text{SPDC}}^{m_A, c_A}\rangle \langle \psi_{\text{SPDC}}^{m_A, c_A}| \otimes |\psi_{\text{SPDC}}^{c_B, m_B}\rangle \langle \psi_{\text{SPDC}}^{c_B, m_B}| \\ &= (1-p_A)(1-p_B) \cdot \exp\left(\sqrt{p_A} m_A^\dagger c_A^\dagger + \sqrt{p_B} c_B^\dagger m_B^\dagger\right) |0\rangle \langle 0| \exp\left(\sqrt{p_A} m_A c_A + \sqrt{p_B} c_B m_B\right). \end{aligned} \quad (\text{D.3})$$

State after the beam-splitter:

The 50/50 beam-splitter mixes the central modes c_A and c_B through the unitary map

$$\begin{pmatrix} c_A \\ c_B \end{pmatrix} \longrightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} c_A \\ c_B \end{pmatrix}. \quad (\text{D.4})$$

Applied to the SPDC exponent, this gives

$$\begin{aligned} \sqrt{p_A} m_A^\dagger c_A^\dagger + \sqrt{p_B} c_B^\dagger m_B^\dagger &\xrightarrow{\text{BS}} \sqrt{p_A} m_A^\dagger \frac{c_A + c_B}{\sqrt{2}} + \sqrt{p_B} \frac{c_A - c_B}{\sqrt{2}} m_B^\dagger \\ &= \left(m_A^\dagger \quad m_B^\dagger \right) \underbrace{\frac{1}{\sqrt{2}} \begin{pmatrix} \sqrt{p_A} & \sqrt{p_A} \\ \sqrt{p_B} & -\sqrt{p_B} \end{pmatrix}}_M \begin{pmatrix} c_A^\dagger \\ c_B^\dagger \end{pmatrix}, \end{aligned}$$

leading to the post-beam-splitter state

$$\rho_{BS} = (1 - p_A)(1 - p_B) \cdot \exp \left[\begin{pmatrix} m_A^\dagger & m_B^\dagger \end{pmatrix} M \begin{pmatrix} c_A^\dagger \\ c_B^\dagger \end{pmatrix} \right] |0\rangle \langle 0| \exp \left[\begin{pmatrix} m_A & m_B \end{pmatrix} M \begin{pmatrix} c_A \\ c_B \end{pmatrix} \right]. \quad (\text{D.5})$$

POVMs and measurement:

From section A.1, the POVM elements associated with the *no-click* and *click* events on mode a , for a photon-detector of efficiency η and dark-count rate dc , are

$$\begin{aligned} \Pi_0^a(\eta, dc) &= \kappa \cdot R^{a^\dagger a}, \\ \Pi_1^a(\eta, dc) &= \mathbb{1} - \kappa \cdot R^{a^\dagger a}, \end{aligned} \quad (\text{D.6})$$

where $R = 1 - \eta$ and $\kappa = 1 - dc$. The heralding event at the central station corresponds to a *click* on mode c_A and a *no-click* on mode c_B (or vice versa), so

$$\rho_{Lnk} = \frac{1}{\mathcal{N}} \text{Tr}_{c_A, c_B} [\Pi_1^{c_A}(\eta_{Lnk}, dc_{Lnk}) \otimes \Pi_0^{c_B}(\eta_{Lnk}, dc_{Lnk}) \cdot \rho_{BS}]. \quad (\text{D.7})$$

Expanding the *click* POVM yields

$$\begin{aligned} \rho_{Lnk} &= \frac{1}{\mathcal{N}} \text{Tr}_{c_A, c_B} \left[\left(\mathbb{1} - \kappa_{Lnk} \cdot R_{Lnk}^{c_A^\dagger c_A} \right) \otimes \kappa_{Lnk} \cdot R_{Lnk}^{c_B^\dagger c_B} \cdot \rho_{BS} \right] \\ &= \frac{\kappa_{Lnk}}{\mathcal{N}} \cdot \text{Tr}_{c_A, c_B} \left[R_{Lnk}^{c_B^\dagger c_B} \cdot \rho_{BS} \right] - \frac{\kappa_{Lnk}^2}{\mathcal{N}} \cdot \text{Tr}_{c_A, c_B} \left[R_{Lnk}^{c_A^\dagger c_A} R_{Lnk}^{c_B^\dagger c_B} \cdot \rho_{BS} \right]. \end{aligned} \quad (\text{D.8})$$

Trace computation:

We first focus on $\text{Tr}_{c_A, c_B} \left[R_{Lnk}^{c_B^\dagger c_B} \cdot \rho_{BS} \right]$. Using the cyclic property of the trace,

$$\text{Tr}_{c_A, c_B} \left[R_{Lnk}^{c_B^\dagger c_B} \cdot \rho_{BS} \right] = \text{Tr}_{c_A, c_B} \left[R_{Lnk}^{c_B^\dagger c_B/2} \cdot |\psi_{BS}\rangle \langle \psi_{BS}| \cdot R_{Lnk}^{c_B^\dagger c_B/2} \right]. \quad (\text{D.9})$$

The identity $x^{a^\dagger a} f(a^\dagger) = f(x a^\dagger) x^{a^\dagger a}$ scales $c_B^\dagger$ by $\sqrt{R_{Lnk}}$ and annihilates the remaining $R_{Lnk}^{c_B^\dagger c_B/2}$ on the vacuum:

$$R_{Lnk}^{c_B^\dagger c_B/2} \exp \left[\begin{pmatrix} m_A^\dagger & m_B^\dagger \end{pmatrix} M \begin{pmatrix} c_A^\dagger \\ c_B^\dagger \end{pmatrix} \right] |0\rangle = \exp \left[\begin{pmatrix} m_A^\dagger & m_B^\dagger \end{pmatrix} M \begin{pmatrix} c_A^\dagger \\ \sqrt{R_{Lnk}} c_B^\dagger \end{pmatrix} \right] \underbrace{R_{Lnk}^{c_B^\dagger c_B/2} |0\rangle}_{=|0\rangle}, \quad (\text{D.10})$$

which defines a new matrix M_σ :

$$\begin{pmatrix} m_A^\dagger & m_B^\dagger \end{pmatrix} M \begin{pmatrix} c_A^\dagger \\ \sqrt{R_{Lnk}} c_B^\dagger \end{pmatrix} = \begin{pmatrix} m_A^\dagger & m_B^\dagger \end{pmatrix} \underbrace{\frac{1}{\sqrt{2}} \begin{pmatrix} \sqrt{p_A} & \sqrt{p_A R_{Lnk}} \\ \sqrt{p_B} & -\sqrt{p_B R_{Lnk}} \end{pmatrix}}_{M_\sigma} \begin{pmatrix} c_A^\dagger \\ c_B^\dagger \end{pmatrix}. \quad (\text{D.11})$$

We now exploit the singular value decomposition (SVD) of M_σ to diagonalize this bilinear form, then use the basis invariance of the trace to simplify the partial trace. For a 2×2 matrix, the SVD $M_\sigma = U_\sigma^\dagger \Sigma V_\sigma$ provides

$$U_\sigma = \begin{pmatrix} \cos \theta & \sin \theta \\ \sin \theta & -\cos \theta \end{pmatrix} \quad \Sigma = \begin{pmatrix} \sigma_+ & 0 \\ 0 & \sigma_- \end{pmatrix} \quad V_\sigma = \begin{pmatrix} \cos \phi & \sin \phi \\ \sin \phi & -\cos \phi \end{pmatrix}, \quad (\text{D.12})$$

with U_σ, V_σ unitary and Σ diagonal. Applied to M_σ , this gives

$$\begin{pmatrix} m_A^\dagger & m_B^\dagger \end{pmatrix} M_\sigma \begin{pmatrix} c_A^\dagger \\ c_B^\dagger \end{pmatrix} = \begin{pmatrix} m_A^\dagger & m_B^\dagger \end{pmatrix} U_\sigma^\dagger \Sigma V_\sigma \begin{pmatrix} c_A^\dagger \\ c_B^\dagger \end{pmatrix}. \quad (\text{D.13})$$

The unitarity of U_σ and V_σ defines the rotated modes

$$\begin{pmatrix} m_\sigma^+ \\ m_\sigma^- \end{pmatrix} = U_\sigma \cdot \begin{pmatrix} m_A \\ m_B \end{pmatrix} \quad \begin{pmatrix} c_\sigma^+ \\ c_\sigma^- \end{pmatrix} = V_\sigma^\dagger \cdot \begin{pmatrix} c_A \\ c_B \end{pmatrix}, \quad (\text{D.14})$$

so that

$$R_{\text{Lnk}}^{c_B^\dagger c_B/2} \cdot |\psi_{\text{BS}}\rangle = \sqrt{1-p_A} \sqrt{1-p_B} \cdot \exp \left[\sigma_+ m_\sigma^{+\dagger} c_\sigma^{+\dagger} + \sigma_- c_\sigma^{-\dagger} m_\sigma^{-\dagger} \right] |0\rangle. \quad (\text{D.15})$$

Since the rotated modes mutually commute, the corresponding density operator factorizes:

$$R_{\text{Lnk}}^{c_B^\dagger c_B/2} \cdot \rho_{\text{BS}} \cdot R_{\text{Lnk}}^{c_B^\dagger c_B/2} = \underbrace{\exp \left[\sigma_+ m_\sigma^{+\dagger} c_\sigma^{+\dagger} \right] |0\rangle \langle 0| \exp \left[\sigma_+ m_\sigma^+ c_\sigma^+ \right]}_{\rho_{\sigma_+}} \otimes \underbrace{\exp \left[\sigma_- c_\sigma^{-\dagger} m_\sigma^{-\dagger} \right] |0\rangle \langle 0| \exp \left[\sigma_- c_\sigma^- m_\sigma^- \right]}_{\rho_{\sigma_-}}, \quad (\text{D.16})$$

up to the factor $(1-p_A)(1-p_B)$. The Fock space generated by $c_A^\dagger, c_B^\dagger$ is stable under the unitary mapping to $c_\sigma^{+\dagger}, c_\sigma^{-\dagger}$, so the partial trace can be taken in the rotated basis:

$$\text{Tr}_{c_A, c_B} \left[R_{\text{Lnk}}^{c_B^\dagger c_B} \cdot \rho_{\text{BS}} \right] = (1-p_A)(1-p_B) \cdot \text{Tr}_{c_\sigma^+} [\rho_{\sigma_+}] \otimes \text{Tr}_{c_\sigma^-} [\rho_{\sigma_-}]. \quad (\text{D.17})$$

From section A.2, the partial trace of a two-mode squeezed state yields a thermal state:

$$(1-x^2) \cdot \text{Tr}_b \left[\exp \left(x \cdot a^\dagger b^\dagger \right) |0\rangle \langle 0| \exp \left(x \cdot a b \right) \right] = \rho_{\text{Therm}}^a(x), \quad (\text{D.18})$$

where $\rho_{\text{Therm}}^a(x)$ is the thermal state on mode a with average photon number $\bar{n} = x^2/(1-x^2)$. Applied to both factors,

$$\text{Tr}_{c_\sigma^+} [\rho_{\sigma_+}] = \frac{1}{1-\sigma_+^2} \rho_{\text{Therm}}^{m_\sigma^+}(\sigma_+) \quad \text{Tr}_{c_\sigma^-} [\rho_{\sigma_-}] = \frac{1}{1-\sigma_-^2} \rho_{\text{Therm}}^{m_\sigma^-}(\sigma_-), \quad (\text{D.19})$$

giving the first term of Equation D.8:

$$\text{Tr}_{c_A, c_B} \left[R_{\text{Lnk}}^{c_B^\dagger c_B} \cdot \rho_{\text{BS}} \right] = \frac{(1-p_A)(1-p_B)}{(1-\sigma_+^2)(1-\sigma_-^2)} \cdot \rho_{\text{Therm}}^{m_\sigma^+}(\sigma_+) \otimes \rho_{\text{Therm}}^{m_\sigma^-}(\sigma_-). \quad (\text{D.20})$$

The second term follows the same logic. Since $R_{\text{Lnk}}^{c_A^\dagger c_A}$ and $R_{\text{Lnk}}^{c_B^\dagger c_B}$ commute,

$$R_{\text{Lnk}}^{c_A^\dagger c_A/2} R_{\text{Lnk}}^{c_B^\dagger c_B/2} \cdot |\psi_{\text{BS}}\rangle = \sqrt{1-p_A} \sqrt{1-p_B} \cdot \exp \left[\omega_+ m_\omega^{+\dagger} c_\omega^{+\dagger} + \omega_- c_\omega^{-\dagger} m_\omega^{-\dagger} \right] |0\rangle, \quad (\text{D.21})$$

where $\omega_\pm$ and $m_\omega^+, m_\omega^-, c_\omega^+, c_\omega^-$ are the singular elements of the matrix

$$\begin{pmatrix} m_A^\dagger & m_B^\dagger \end{pmatrix} \underbrace{\frac{\sqrt{R_{\text{Lnk}}}}{\sqrt{2}} \begin{pmatrix} \sqrt{p_A} & \sqrt{p_A} \\ \sqrt{p_B} & -\sqrt{p_B} \end{pmatrix}}_{M_\omega} \begin{pmatrix} c_A^\dagger \\ c_B^\dagger \end{pmatrix} = \begin{pmatrix} m_\omega^{+\dagger} & m_\omega^{-\dagger} \end{pmatrix} \begin{pmatrix} \omega_+ & 0 \\ 0 & \omega_- \end{pmatrix} \begin{pmatrix} c_\omega^{+\dagger} \\ c_\omega^{-\dagger} \end{pmatrix}, \quad (\text{D.22})$$

leading to

$$\text{Tr}_{c_A, c_B} \left[R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{Lnk}}^{c_B^\dagger c_B} \cdot \rho_{\text{BS}} \right] = \frac{(1-p_A)(1-p_B)}{(1-\omega_+^2)(1-\omega_-^2)} \cdot \rho_{\text{Therm}}^{m_\omega^+}(\omega_+) \otimes \rho_{\text{Therm}}^{m_\omega^-}(\omega_-). \quad (\text{D.23})$$

Combining both terms gives the shared state, up to normalization:

$$\begin{aligned} \rho_{\text{Lnk}} = \frac{\kappa_{\text{Lnk}}}{\mathcal{N}} & \left[\frac{(1-p_A)(1-p_B)}{(1-\sigma_+^2)(1-\sigma_-^2)} \cdot \rho_{\text{Therm}}^{m_\sigma^+}(\sigma_+) \otimes \rho_{\text{Therm}}^{m_\sigma^-}(\sigma_-) \right. \\ & \left. - \kappa_{\text{Lnk}} \cdot \frac{(1-p_A)(1-p_B)}{(1-\omega_+^2)(1-\omega_-^2)} \cdot \rho_{\text{Therm}}^{m_\omega^+}(\omega_+) \otimes \rho_{\text{Therm}}^{m_\omega^-}(\omega_-) \right], \end{aligned} \quad (\text{D.24})$$

where $\mathcal{N} = \text{Tr} [\rho_{\text{Lnk}}]$ is the normalization constant.

Normalization:

The trace of the unnormalized expression yields

$$\mathcal{N} = \kappa_{\text{Lnk}} \cdot (1-p_A)(1-p_B) \left[\frac{1}{(1-\sigma_+^2)(1-\sigma_-^2)} - \frac{\kappa_{\text{Lnk}}}{(1-\omega_+^2)(1-\omega_-^2)} \right]. \quad (\text{D.25})$$

Introducing $\xi = \kappa_{\text{Lnk}} \cdot \frac{(1-\sigma_+^2)(1-\sigma_-^2)}{(1-\omega_+^2)(1-\omega_-^2)}$, the normalized state finally reads

$$\rho_{\text{Lnk}} = \frac{1}{1-\xi} \cdot \rho_{\text{Therm}}^{m_\sigma^+}(\sigma_+) \otimes \rho_{\text{Therm}}^{m_\sigma^-}(\sigma_-) - \frac{\xi}{1-\xi} \cdot \rho_{\text{Therm}}^{m_\omega^+}(\omega_+) \otimes \rho_{\text{Therm}}^{m_\omega^-}(\omega_-). \quad (\text{D.26})$$

D.2 .Fidelity of the elementary link to a Bell pair

This appendix derives the fidelity $\mathcal{F}_{\text{Lnk}}$ of the heralded elementary link ρ_{Lnk} from the exact modelization. The starting point is the exact state derivation of the state from section D.1, concretely

$$\rho_{\text{Lnk}} = \frac{1}{1-\xi} \cdot \rho_{\text{Therm}}^{m_{\sigma}^+}(\sigma_+) \otimes \rho_{\text{Therm}}^{m_{\sigma}^-}(\sigma_-) - \frac{\xi}{1-\xi} \cdot \rho_{\text{Therm}}^{m_{\omega}^+}(\omega_+) \otimes \rho_{\text{Therm}}^{m_{\omega}^-}(\omega_-), \quad (\text{D.27})$$

with $\rho_{\text{Therm}}^a(x)$ the thermal state on mode a

$$\rho_{\text{Therm}}^a(x) = (1-x^2) \sum_n (x^2)^n |n\rangle \langle n|_a, \quad (\text{D.28})$$

and $\xi = (1 - \text{dc}_{\text{Lnk}}) \cdot \frac{(1-\sigma_+^2)(1-\sigma_-^2)}{(1-\omega_+^2)(1-\omega_-^2)}$. The fidelity is defined as $\mathcal{F}_{\text{Lnk}} = \langle \Psi^+ | \rho_{\text{Lnk}} | \Psi^+ \rangle$, with the single-excitation Bell pair

$$|\Psi^+\rangle = \frac{m_A^\dagger + m_B^\dagger}{\sqrt{2}} |0\rangle. \quad (\text{D.29})$$

The heralded state ρ_{Lnk} is diagonal in the collective modes $m_{\sigma}^{\pm}$ and $m_{\omega}^{\pm}$, defined from the SVD

$$\begin{aligned} \begin{pmatrix} m_{\sigma}^+ \\ m_{\sigma}^- \end{pmatrix} &= U_{\sigma} \begin{pmatrix} m_A \\ m_B \end{pmatrix} = \begin{pmatrix} \Sigma_A^+ & \Sigma_B^+ \\ \Sigma_A^- & \Sigma_B^- \end{pmatrix} \begin{pmatrix} m_A \\ m_B \end{pmatrix}, \\ \begin{pmatrix} m_{\omega}^+ \\ m_{\omega}^- \end{pmatrix} &= U_{\omega} \begin{pmatrix} m_A \\ m_B \end{pmatrix} = \begin{pmatrix} \Omega_A^+ & \Omega_B^+ \\ \Omega_A^- & \Omega_B^- \end{pmatrix} \begin{pmatrix} m_A \\ m_B \end{pmatrix}. \end{aligned} \quad (\text{D.30})$$

The fidelity therefore requires expressing $|\Psi^+\rangle$ in the rotated bases. From the orthogonality of U_{σ} we have $U_{\sigma}^\dagger = U_{\sigma}^T$, thus

$$m_A^\dagger = \Sigma_A^+ m_{\sigma}^{+\dagger} + \Sigma_A^- m_{\sigma}^{-\dagger}, \quad m_B^\dagger = \Sigma_B^+ m_{\sigma}^{+\dagger} + \Sigma_B^- m_{\sigma}^{-\dagger}, \quad (\text{D.31})$$

thus the Bell pair becomes is a delocalized single excitation over the two collective modes

$$|\Psi^+\rangle = \frac{1}{\sqrt{2}} \left[(\Sigma_A^+ + \Sigma_B^+) |1\rangle_{m_{\sigma}^+} |0\rangle_{m_{\sigma}^-} + (\Sigma_A^- + \Sigma_B^-) |0\rangle_{m_{\sigma}^+} |1\rangle_{m_{\sigma}^-} \right]. \quad (\text{D.32})$$

From the decomposition of Equation D.32, the contirbution of the σ term of the state ρ_{Lnk} to the fidelits is

$$\begin{aligned} \mathcal{F}_{\sigma} &= \langle \Psi^+ | \rho_{\text{Therm}}^{m_{\sigma}^+}(\sigma_+) \otimes \rho_{\text{Therm}}^{m_{\sigma}^-}(\sigma_-) | \Psi^+ \rangle \\ &= \frac{(\Sigma_A^+ + \Sigma_B^+)^2}{2} P_{10}^{\sigma} + \frac{(\Sigma_A^- + \Sigma_B^-)^2}{2} P_{01}^{\sigma}, \end{aligned} \quad (\text{D.33})$$

where

$$P_{ij}^{\sigma} = \langle i | \rho_{\text{Therm}}^{m_{\sigma}^+}(\sigma_+) | i \rangle_{m_{\sigma}^+} \langle j | \rho_{\text{Therm}}^{m_{\sigma}^-}(\sigma_-) | j \rangle_{m_{\sigma}^-}. \quad (\text{D.34})$$

From the expansion of the thermal population from Equation D.28, we have $\langle 1 | \rho_{\text{Therm}}(x) | 1 \rangle = (1 - x^2)x^2$ and $\langle 0 | \rho_{\text{Therm}}(x) | 0 \rangle = (1 - x^2)$, hence

$$P_{10}^\sigma = \sigma_+^2(1 - \sigma_+^2)(1 - \sigma_-^2), \quad P_{01}^\sigma = \sigma_-^2(1 - \sigma_+^2)(1 - \sigma_-^2), \quad (\text{D.35})$$

giving

$$\begin{aligned} \mathcal{F}_\sigma &= \frac{(1 - \sigma_+^2)(1 - \sigma_-^2)}{2} [(\Sigma_A^+ + \Sigma_B^+)^2 \sigma_+^2 + (\Sigma_A^- + \Sigma_B^-)^2 \sigma_-^2] \\ \mathcal{F}_\omega &= \frac{(1 - \omega_+^2)(1 - \omega_-^2)}{2} [(\Omega_A^+ + \Omega_B^+)^2 \omega_+^2 + (\Omega_A^- + \Omega_B^-)^2 \omega_-^2] \end{aligned} \quad (\text{D.36})$$

The previous calculation finally gives

$$\begin{aligned} \mathcal{F}_{\text{Lnk}} &= \frac{1}{2(1 - \xi)} \left[(1 - \sigma_+^2)(1 - \sigma_-^2) ((\Sigma_A^+ + \Sigma_B^+)^2 \sigma_+^2 + (\Sigma_A^- + \Sigma_B^-)^2 \sigma_-^2) \right. \\ &\quad \left. - \xi (1 - \omega_+^2)(1 - \omega_-^2) ((\Omega_A^+ + \Omega_B^+)^2 \omega_+^2 + (\Omega_A^- + \Omega_B^-)^2 \omega_-^2) \right]. \end{aligned} \quad (\text{D.37})$$

For readability, we introduce the thermal envelopes

$$C_\sigma = (1 - \sigma_+^2)(1 - \sigma_-^2), \quad C_\omega = (1 - \omega_+^2)(1 - \omega_-^2), \quad (\text{D.38})$$

so that $\xi = (1 - \text{dc}_{\text{Lnk}}) C_\sigma / C_\omega$, together with the Bell-pair overlaps onto the collective modes

$$S_\pm = \frac{(\Sigma_A^\pm + \Sigma_B^\pm)^2}{2}, \quad W_\pm = \frac{(\Omega_A^\pm + \Omega_B^\pm)^2}{2}. \quad (\text{D.39})$$

The orthogonality of U_σ and U_ω enforces $S_+ + S_- = 1$ and $W_+ + W_- = 1$. The fidelity then compacts to

$$\boxed{\mathcal{F}_{\text{Lnk}} = \frac{C_\sigma C_\omega}{C_\omega - (1 - \text{dc}_{\text{Lnk}}) C_\sigma} \left[S_+ \sigma_+^2 + S_- \sigma_-^2 - (1 - \text{dc}_{\text{Lnk}}) (W_+ \omega_+^2 + W_- \omega_-^2) \right].} \quad (\text{D.40})$$

D.3 .Derivation of the displaced local probabilities

The entanglement witness of section 3.3 is built from the displaced local probabilities $\mathbb{P}_{\text{Loc}}^{ij}(\alpha, \beta)$, the joint statistics of a *click/no-click* measurement on Alice's mode m_A displaced by α and Bob's mode m_B displaced by β . This appendix derives their closed form on the heralded state ρ_{Lnk} of equation (3.4). The derivation proceeds in three steps: the displaced POVM elements are expanded by linearity of the trace into a few elementary traces, these traces are evaluated on a generic thermal *toy state*, and the result is mapped back onto the two thermal components of ρ_{Lnk} .

Displaced POVM and trace expansion

The displaced local probabilities read

$$\mathbb{P}_{\text{Loc}}^{ij}(\alpha, \beta) = \text{Tr}_{\text{tot}} \left[\Pi_{\text{Loc}}^{ij} \cdot \mathcal{D}(\alpha) \mathcal{D}(\beta) \cdot \rho_{\text{Lnk}} \cdot \mathcal{D}^\dagger(\alpha) \mathcal{D}^\dagger(\beta) \right], \quad (\text{D.41})$$

with $i, j \in \{0, 1\}$ the *no-click/click* indices, ρ_{Lnk} the normalized heralded state from equation (3.4)

$$\rho_{\text{Lnk}} = \frac{1}{1 - \xi} \cdot \rho_{\text{Therm}}^{m_\sigma^+}(\sigma_+) \otimes \rho_{\text{Therm}}^{m_\sigma^-}(\sigma_-) - \frac{\xi}{1 - \xi} \cdot \rho_{\text{Therm}}^{m_\omega^+}(\omega_+) \otimes \rho_{\text{Therm}}^{m_\omega^-}(\omega_-), \quad (\text{D.42})$$

and Π_{Loc}^{ij} the POVM elements of the two non-photon-number-resolving detectors. Denoting η_A (resp. η_B) the detector efficiency on Alice's (resp. Bob's) side and $R_A = 1 - \eta_A$, $R_B = 1 - \eta_B$ the associated losses, these elements are

$$\begin{aligned} \Pi_{\text{Loc}}^{00} &= R_A^{m_A^\dagger m_A} \cdot R_B^{m_B^\dagger m_B}, & \Pi_{\text{Loc}}^{01} &= \left(\mathbb{1} - R_A^{m_A^\dagger m_A} \right) \cdot R_B^{m_B^\dagger m_B}, \\ \Pi_{\text{Loc}}^{10} &= R_A^{m_A^\dagger m_A} \cdot \left(\mathbb{1} - R_B^{m_B^\dagger m_B} \right), & \Pi_{\text{Loc}}^{11} &= \left(\mathbb{1} - R_A^{m_A^\dagger m_A} \right) \cdot \left(\mathbb{1} - R_B^{m_B^\dagger m_B} \right). \end{aligned} \quad (\text{D.43})$$

By linearity of the trace, expanding the four elements reduces the computation to three elementary traces: the action of $R_A^{m_A^\dagger m_A} R_B^{m_B^\dagger m_B}$, $R_A^{m_A^\dagger m_A} \cdot \mathbb{1}$, and $\mathbb{1} \cdot R_B^{m_B^\dagger m_B}$ on the displaced state. The state ρ_{Lnk} is itself a mixture of two thermal terms, $\rho_{\text{Therm}}^{m_\sigma^+}(\sigma_+) \otimes \rho_{\text{Therm}}^{m_\sigma^-}(\sigma_-)$ and $\rho_{\text{Therm}}^{m_\omega^+}(\omega_+) \otimes \rho_{\text{Therm}}^{m_\omega^-}(\omega_-)$, so we evaluate each trace on a single generic *toy state*

$$\rho_0 = \rho_{\text{Therm}}^{\bar{a}}(x_a) \otimes \rho_{\text{Therm}}^{\bar{b}}(x_b), \quad (\text{D.44})$$

with mean photon numbers $n_a = \bar{n}(x_a)$, $n_b = \bar{n}(x_b)$, and rotated modes $\bar{a}^\dagger = \delta a^\dagger + \tau b^\dagger$, $\bar{b}^\dagger = \tau a^\dagger - \delta b^\dagger$ with $|\delta|^2 + |\tau|^2 = 1$. The two thermal components of ρ_{Lnk} are recovered at the end by specializing (x_a, x_b, δ, τ) .

Action of $R_A^{a^\dagger a} R_B^{b^\dagger b}$

We start with the two-detector term on the toy state,

$$\text{Tr} \left[R_A^{a^\dagger a} R_B^{b^\dagger b} \cdot \mathcal{D}(\alpha) \mathcal{D}(\beta) \cdot \rho_0 \cdot \mathcal{D}^\dagger(\alpha) \mathcal{D}^\dagger(\beta) \right]. \quad (\text{D.45})$$

Using the cyclic property of the trace, we symmetrize the loss operators

$$\text{Tr} \left[R_A^{a^\dagger a/2} R_B^{b^\dagger b/2} \cdot \mathcal{D}(\alpha) \mathcal{D}(\beta) \cdot \rho_0 \cdot \mathcal{D}^\dagger(\alpha) \mathcal{D}^\dagger(\beta) \cdot R_A^{a^\dagger a/2} R_B^{b^\dagger b/2} \right]. \quad (\text{D.46})$$

This form exhibits four steps: express ρ_0 in the coherent-state basis, apply the displacements, apply the loss operators, then trace over the modes a, b . We define the intermediate states

$$\begin{cases} \rho_0 = \rho_{\text{Therm}}^{\bar{a}}(x_a) \otimes \rho_{\text{Therm}}^{\bar{b}}(x_b), \\ \rho_1 = \mathcal{D}(\alpha) \mathcal{D}(\beta) \cdot \rho_0 \cdot \mathcal{D}^\dagger(\alpha) \mathcal{D}^\dagger(\beta), \\ \rho_2 = R_A^{a^\dagger a/2} R_B^{b^\dagger b/2} \cdot \rho_1 \cdot R_A^{a^\dagger a/2} R_B^{b^\dagger b/2}. \end{cases} \quad (\text{D.47})$$

Expression of ρ_0 : Using the P -representation of the thermal state,

$$\rho_0 = \frac{1}{\pi^2 n_a n_b} \int \int_{\mathbb{C}^2} e^{-|\gamma_a|^2/n_a - |\gamma_b|^2/n_b} |\gamma_a, \gamma_b\rangle \langle \gamma_a, \gamma_b| d^2\gamma_a d^2\gamma_b. \quad (\text{D.48})$$

The coherent states $|\gamma_a, \gamma_b\rangle$ live in the rotated modes $\bar{a}, \bar{b}$, while detection and displacement act on a, b . Applying the mode-switching identity of appendix D.4,

$$|\bar{a} : \gamma_a, \bar{b} : \gamma_b\rangle = e^{i\phi} \cdot |a : \delta\gamma_a + \tau\gamma_b, b : \tau\gamma_a - \delta\gamma_b\rangle, \quad (\text{D.49})$$

so, writing $\omega = \delta\gamma_a + \tau\gamma_b$ and $\nu = \tau\gamma_a - \delta\gamma_b$, the toy state in the a, b modes reads

$$\rho_0 = \frac{1}{\pi^2 n_a n_b} \int \int_{\mathbb{C}^2} e^{-|\gamma_a|^2/n_a - |\gamma_b|^2/n_b} |\omega, \nu\rangle \langle \omega, \nu| d^2\gamma_a d^2\gamma_b. \quad (\text{D.50})$$

Expression of ρ_1 : The displacement operators commute with the integrals, so their action reduces to a shift of the coherent-state labels,

$$\mathcal{D}(\alpha) \mathcal{D}(\beta) |\omega, \nu\rangle = |\alpha + \omega, \beta + \nu\rangle, \quad (\text{D.51})$$

leading to

$$\rho_1 = \frac{1}{\pi^2 n_a n_b} \int \int_{\mathbb{C}^2} e^{-|\gamma_a|^2/n_a - |\gamma_b|^2/n_b} |\alpha + \omega, \beta + \nu\rangle \langle \alpha + \omega, \beta + \nu| d^2\gamma_a d^2\gamma_b. \quad (\text{D.52})$$

Expression of ρ_2 : The action of a loss operator on a coherent state is given in appendix D.4

$$R_A^{a^\dagger a/2} R_B^{b^\dagger b/2} |u, v\rangle = e^{-\left(\frac{\eta_A}{2}|u|^2 + \frac{\eta_B}{2}|v|^2\right)} \left| \sqrt{R_A} u, \sqrt{R_B} v \right\rangle. \quad (\text{D.53})$$

Setting $\omega' = \sqrt{R_A}(\alpha + \omega)$ and $\nu' = \sqrt{R_B}(\beta + \nu)$,

$$\rho_2 = \frac{1}{\pi^2 n_a n_b} \int \int_{\mathbb{C}^2} e^{-|\gamma_a|^2/n_a - |\gamma_b|^2/n_b} \cdot e^{-\left(\frac{\eta_A}{2}|\alpha + \omega|^2 + \frac{\eta_B}{2}|\beta + \nu|^2\right)} |\omega', \nu'\rangle \langle \omega', \nu'| d^2\gamma_a d^2\gamma_b. \quad (\text{D.54})$$

Trace of ρ_2 : The final step traces ρ_2 over the modes a, b . Assuming the sum and the integral commute, the Fock-basis trace of a normalized coherent state gives

$$\sum_{i,j} |\langle i, j | \omega', \nu' \rangle|^2 = 1, \quad (\text{D.55})$$

so that, recalling $\omega = \delta\gamma_a + \tau\gamma_b$ and $\nu = \tau\gamma_a - \delta\gamma_b$,

$$\text{Tr}[\rho_2] = \frac{1}{\pi^2 n_a n_b} \int \int_{\mathbb{C}^2} e^{-|\gamma_a|^2/n_a - |\gamma_b|^2/n_b} \cdot e^{-\left(\frac{\eta_A}{2}|\alpha + \delta\gamma_a + \tau\gamma_b|^2 + \frac{\eta_B}{2}|\beta + \tau\gamma_a - \delta\gamma_b|^2\right)} d^2\gamma_a d^2\gamma_b. \quad (\text{D.56})$$

Evaluating the Gaussian integral of appendix A.3 with $a = 1/n_a$, $b = 1/n_b$, $c = \eta_A/2$, $d = \eta_B/2$ yields

$$\text{Tr}[\rho_2] = \frac{1}{\Gamma} \exp\left[-\frac{n_a \eta_A \eta_B}{\Gamma} |\tau\alpha - \delta\beta|^2\right] \exp\left[-\frac{n_b \eta_A \eta_B}{\Gamma} |\delta\alpha + \tau\beta|^2\right] \exp\left[-\frac{\eta_A |\alpha|^2 + \eta_B |\beta|^2}{\Gamma}\right], \quad (\text{D.57})$$

with

$$\Gamma = (1 + n_a(\eta_A \delta^2 + \eta_B \tau^2)) (1 + n_b(\eta_A \tau^2 + \eta_B \delta^2)) - n_a n_b \delta^2 \tau^2 (\eta_A - \eta_B)^2. \quad (\text{D.58})$$

We denote this quantity f , an eight-variable function with four arguments characterizing the link (δ, τ, x_a, x_b) and four characterizing the witness $(\alpha, \beta, \eta_A, \eta_B)$,

$$f(\alpha, \beta, \eta_A, \eta_B, \delta, \tau, x_a, x_b) = \text{Tr}[\rho_2]. \quad (\text{D.59})$$

Action of $R_A^{a^\dagger a} \cdot \mathbb{1}$ and $\mathbb{1} \cdot R_B^{b^\dagger b}$

The single-detector terms follow from f by a limit. The loss operator is diagonal in the Fock basis,

$$R^{a^\dagger a} = \sum_n R^n |n\rangle \langle n|, \quad \mathbb{1} = \sum_n |n\rangle \langle n| = \lim_{R \rightarrow 1} R^{a^\dagger a}, \quad (\text{D.60})$$

so $R_A^{a^\dagger a} \cdot \mathbb{1} = \lim_{R_B \rightarrow 1} (R_A^{a^\dagger a} R_B^{b^\dagger b})$. Assuming the trace and the limit commute, the single-detector terms are obtained by setting $\eta_B \rightarrow 0$ (resp. $\eta_A \rightarrow 0$) in f ,

$$\begin{aligned} \text{Tr} \left[R_A^{a^\dagger a} \cdot \mathbb{1} \cdot \mathcal{D}(\alpha) \mathcal{D}(\beta) \cdot \rho_0 \cdot \mathcal{D}^\dagger(\alpha) \mathcal{D}^\dagger(\beta) \right] &= f(\alpha, \beta, \eta_A, 0, \delta, \tau, x_a, x_b), \\ \text{Tr} \left[\mathbb{1} \cdot R_B^{b^\dagger b} \cdot \mathcal{D}(\alpha) \mathcal{D}(\beta) \cdot \rho_0 \cdot \mathcal{D}^\dagger(\alpha) \mathcal{D}^\dagger(\beta) \right] &= f(\alpha, \beta, 0, \eta_B, \delta, \tau, x_a, x_b). \end{aligned} \quad (\text{D.61})$$

Assembling the probabilities

We now piece the elementary traces together. Expanding the POVM elements of equation (D.41) in terms of the three computed traces,

$$\begin{aligned} \Pi_{\text{Loc}}^{00} &= R_A^{m_A^\dagger m_A} R_B^{m_B^\dagger m_B}, \\ \Pi_{\text{Loc}}^{01} &= \mathbb{1} \cdot R_B^{m_B^\dagger m_B} - R_A^{m_A^\dagger m_A} R_B^{m_B^\dagger m_B}, \\ \Pi_{\text{Loc}}^{10} &= R_A^{m_A^\dagger m_A} \cdot \mathbb{1} - R_A^{m_A^\dagger m_A} R_B^{m_B^\dagger m_B}, \\ \Pi_{\text{Loc}}^{11} &= \mathbb{1} - \mathbb{1} \cdot R_B^{m_B^\dagger m_B} - R_A^{m_A^\dagger m_A} \cdot \mathbb{1} + R_A^{m_A^\dagger m_A} R_B^{m_B^\dagger m_B}. \end{aligned} \quad (\text{D.62})$$

The toy state ρ_0 maps onto the two thermal components of ρ_{Lnk} (cf. equation (3.4)) through

$$\begin{aligned} \rho_{\text{Therm}}^{m_\sigma^+}(\sigma_+) \otimes \rho_{\text{Therm}}^{m_\sigma^-}(\sigma_-) &\longrightarrow \rho_0 \text{ with } x_a = \sigma_+, x_b = \sigma_-, \\ \rho_{\text{Therm}}^{m_\omega^+}(\omega_+) \otimes \rho_{\text{Therm}}^{m_\omega^-}(\omega_-) &\longrightarrow \rho_0 \text{ with } x_a = \omega_+, x_b = \omega_-, \end{aligned} \quad (\text{D.63})$$

with (δ, τ) fixed by the rotation matrices U_σ, U_ω of equation (3.4). Exploiting the Alice/Bob symmetry, we collapse f into reduced functions indexed by the singular component, σ or ω ,

$$\begin{aligned} f(\alpha, \beta, \eta_A, \eta_B, \delta, \tau, \sigma_+, \sigma_-) &\longrightarrow g^\sigma(\alpha, \beta, \eta_A, \eta_B), \\ f(\alpha, \beta, \eta_A, 0, \delta, \tau, \sigma_+, \sigma_-) &\longrightarrow h^\sigma(\alpha, \eta_A), \\ f(\alpha, \beta, 0, \eta_B, \delta, \tau, \sigma_+, \sigma_-) &\longrightarrow h^\sigma(\beta, \eta_B), \end{aligned} \quad (\text{D.64})$$

and likewise g^ω, h^ω for the ω component. With the normalization $\mathbb{P}_{\text{Lnk}}$ written through $\xi = (1 - \text{dc}_{\text{Lnk}}) C_\sigma / C_\omega$, the displaced local probabilities read

$$\begin{aligned} \mathbb{P}_{\text{Loc}}^{00}(\alpha, \beta) &= \frac{1}{1-\xi} g^\sigma(\alpha, \beta, \eta_A, \eta_B) - \frac{\xi}{1-\xi} g^\omega(\alpha, \beta, \eta_A, \eta_B), \\ \mathbb{P}_{\text{Loc}}^{01}(\alpha, \beta) &= \frac{1}{1-\xi} \left[h^\sigma(\beta, \eta_B) - g^\sigma(\alpha, \beta, \eta_A, \eta_B) \right] - \frac{\xi}{1-\xi} \left[h^\omega(\beta, \eta_B) - g^\omega(\alpha, \beta, \eta_A, \eta_B) \right], \\ \mathbb{P}_{\text{Loc}}^{10}(\alpha, \beta) &= \frac{1}{1-\xi} \left[h^\sigma(\alpha, \eta_A) - g^\sigma(\alpha, \beta, \eta_A, \eta_B) \right] - \frac{\xi}{1-\xi} \left[h^\omega(\alpha, \eta_A) - g^\omega(\alpha, \beta, \eta_A, \eta_B) \right], \\ \mathbb{P}_{\text{Loc}}^{11}(\alpha, \beta) &= \frac{1}{1-\xi} \left[1 - h^\sigma(\alpha, \eta_A) - h^\sigma(\beta, \eta_B) + g^\sigma(\alpha, \beta, \eta_A, \eta_B) \right] \\ &\quad - \frac{\xi}{1-\xi} \left[1 - h^\omega(\alpha, \eta_A) - h^\omega(\beta, \eta_B) + g^\omega(\alpha, \beta, \eta_A, \eta_B) \right]. \end{aligned}$$

(D.65)

The reduced functions inherit the structure of f ,

$$g^\sigma(\alpha, \beta, \eta_A, \eta_B) = \frac{1}{\Gamma_g} \exp \left[-\frac{\bar{n}(\sigma_+) \eta_A \eta_B}{2\Gamma_g} |\alpha - \beta|^2 \right] \exp \left[-\frac{\bar{n}(\sigma_-) \eta_A \eta_B}{2\Gamma_g} |\alpha + \beta|^2 \right] \exp \left[-\frac{\eta_A |\alpha|^2 + \eta_B |\beta|^2}{\Gamma_g} \right],$$

$$h^\sigma(\alpha, \eta) = \frac{1}{\Gamma_h} \exp \left[-\frac{\eta |\alpha|^2}{\Gamma_h} \right],$$

(D.66)

with $\Gamma_g = \left(1 + \bar{n}(\sigma_+) \frac{\eta_A + \eta_B}{2}\right) \left(1 + \bar{n}(\sigma_-) \frac{\eta_A + \eta_B}{2}\right) - \left(\frac{\eta_A - \eta_B}{2}\right)^2$ and $\Gamma_h = \left(1 + \bar{n}(\sigma_+) \frac{\eta}{2}\right) \left(1 + \bar{n}(\sigma_-) \frac{\eta}{2}\right) - \left(\frac{\eta}{2}\right)^2$, and g^ω, h^ω obtained by $\sigma_\pm \rightarrow \omega_\pm$.

D.4 .Useful coherent states indentities

This appendix gives general results on the representation of coherent states in the Fock basis, used for the derivation of the displaced local probabilities of appendix D.3.

Mode switching

Here we aim to express a two-mode coherent state in a rotated basis, i.e. to find $\bar{\alpha}, \bar{\beta}$ such that

$$|a : \alpha, b : \beta\rangle = |\bar{a} : \bar{\alpha}, \bar{b} : \bar{\beta}\rangle, \quad (\text{D.67})$$

with the unitary mode rotation

$$\bar{a}^\dagger = \delta a^\dagger + \nu b^\dagger, \quad \bar{b}^\dagger = -\nu^* a^\dagger + \delta^* b^\dagger, \quad |\delta|^2 + |\nu|^2 = 1, \quad (\text{D.68})$$

whose inverse is $a^\dagger = \delta^* \bar{a}^\dagger - \nu \bar{b}^\dagger, b^\dagger = \nu^* \bar{a}^\dagger + \delta \bar{b}^\dagger$.

Since the modes commute, the displacement is a single exponential of the generator, $\mathcal{D}_a(\alpha) \mathcal{D}_b(\beta) = \exp[\alpha a^\dagger + \beta b^\dagger - \text{h.c.}]$. Injecting the inverse relations and collecting terms,

$$\alpha a^\dagger + \beta b^\dagger = (\delta^* \alpha + \nu^* \beta) \bar{a}^\dagger + (-\nu \alpha + \delta \beta) \bar{b}^\dagger, \quad (\text{D.69})$$

so the generator is that of $\mathcal{D}_{\bar{a}}(\bar{\alpha}) \mathcal{D}_{\bar{b}}(\bar{\beta})$, with no phase ($[\bar{a}^\dagger, \bar{b}] = 0$). The vacuum being common to both bases, we get

$$|a : \alpha, b : \beta\rangle = |\bar{a} : \bar{\alpha}, \bar{b} : \bar{\beta}\rangle, \quad \begin{cases} \bar{\alpha} = \delta^* \alpha + \nu^* \beta, \\ \bar{\beta} = -\nu \alpha + \delta \beta. \end{cases} \quad (\text{D.70})$$

Action of $R^{a^\dagger a}$ on a coherent state

First we compute the action of $x^{a^\dagger a}$ on a coherent state $|\alpha\rangle$,

$$\begin{aligned} x^{a^\dagger a} |\alpha\rangle &= x^{a^\dagger a} e^{-\frac{|\alpha|^2}{2}} \sum_n \frac{\alpha^n}{\sqrt{n!}} |n\rangle \\ &= e^{-\frac{|\alpha|^2}{2}} \sum_n \frac{\alpha^n}{\sqrt{n!}} x^n |n\rangle \\ &= e^{-\frac{|\alpha|^2}{2} + \frac{|x\alpha|^2}{2}} e^{-\frac{|x\alpha|^2}{2}} \sum_n \frac{(x\alpha)^n}{\sqrt{n!}} |n\rangle \\ &= e^{(|x|^2 - 1) \frac{|\alpha|^2}{2}} |x\alpha\rangle. \end{aligned} \quad (\text{D.71})$$

Hence, using the equality $x^{a^\dagger a} |\alpha\rangle = e^{(|x|^2-1)\frac{|\alpha|^2}{2}} |x\alpha\rangle$, we get

$$\begin{aligned} R^{\frac{a^\dagger a}{2}} R^{\frac{d^\dagger d}{2}} |\alpha, \beta\rangle &= R^{\frac{a^\dagger a}{2}} |\alpha\rangle R^{\frac{d^\dagger d}{2}} |\beta\rangle \\ &= e^{(|R|-1)\frac{|\alpha|^2}{2}} \cdot e^{(|R|-1)\frac{|\beta|^2}{2}} \left| \sqrt{R}\alpha, \sqrt{R}\beta \right\rangle. \end{aligned} \quad (\text{D.72})$$

Finally, using $R - 1 = -\eta$, we get

$$\begin{aligned} R^{\frac{a^\dagger a}{2}} \cdot R^{\frac{d^\dagger d}{2}} |\alpha, \beta\rangle &= e^{-\frac{\eta}{2}(|\alpha|^2+|\beta|^2)} \left| \sqrt{R}\alpha, \sqrt{R}\beta \right\rangle, \\ R^{\frac{a^\dagger a}{2}} \cdot \mathbb{1} |\alpha, \beta\rangle &= e^{-\frac{\eta}{2}|\alpha|^2} \left| \sqrt{R}\alpha, \beta \right\rangle, \\ \mathbb{1} \cdot R^{\frac{d^\dagger d}{2}} |\alpha, \beta\rangle &= e^{-\frac{\eta}{2}|\beta|^2} \left| \alpha, \sqrt{R}\beta \right\rangle. \end{aligned} \quad (\text{D.73})$$

D.5 .Perturbative coincidence probability

In this appendix we derive the local coincidence probability $\mathbb{P}_{\text{Coinc}}^A$ and $\mathbb{P}_{\text{Coinc}}^B$ required for the evaluation of the Cauchy Schwartz innequality. This probability corresponds to the joint *click* on the two output modes m_T (Tom) and m_J (Jerry), built from the output of a 50/50 beam-splitter placed after the memory's mode m_A or m_B . In the calculation below, we derive such a coincidence in a low-pump perturbative regime, up to the 2-photon sector.

Setup and context

We first notice the setup is symmetric between the two hubs, thus Alice's coincidence probability $\mathbb{P}_{\text{Coinc}}^A$ follows the same derivation, up to the substitutions side efficiencies $dc_A, \eta_A \rightarrow dc_B, \eta_B$. We thereofre focus on Alice's side m_A with the derivation of $\mathbb{P}_{\text{Coinc}}^A$. As usually done in this manuscript, the modelization of the losses and dark-count have been pushed to the POVM element, consequently the efficiencies η_T, η_J and dark-counts dc_T, dc_J carry together the detection setup measuring the coincidence and the loss and dark-counts of Alice's memory itself η_A and dc_A . For clarity, we here assume symmetric situation $\eta_T = \eta_J = \eta_{TJ}$ with no dark-count $dc_T = dc_J = 0$, as including these quantities does not require much more work while weighing the derivation. The coincidence is heralded by a *click* on both beam-splitter outputs m_T and m_J

$$\mathbb{P}_{\text{Coinc}}^A = \langle \Pi_1^{m_T}(\eta_{TJ}) \Pi_1^{m_J}(\eta_{TJ}) \rangle_{\rho_{\text{Lnk}}}, \quad (\text{D.74})$$

with $\Pi_1^{m_T}(\eta_{TJ}) = \mathbb{1} - R_{TJ}^{m_T^\dagger m_T}$ the click POVM on mode m_T (cf. section A.1), and similarly on m_J . For clarity we call $\Pi_{TJ} = \Pi_1^{m_T}(\eta_{TJ}) \Pi_1^{m_J}(\eta_{TJ})$.

Second-order expansion

Since we do not consider dark-count, the joint click requires at least one photon in each output, hence at least two photons in m_A is needed. The 1-photon sector therefore does not contribute: in the low-pump regime $\rho_{\text{Lnk}} = |\Psi^+\rangle \langle \Psi^+| + o(p)$ carries a single delocalized excitation, giving $\mathbb{P}_{\text{Coinc}}^A = o(p)$. The witness must thus be evaluated at the 2-photon sector, which requires the second-order expansion of the SPDC state. We recall the 2-photon sector carries the state

$$|\psi_{\text{ini}}^{(2)}\rangle = \frac{1}{\sqrt{3}} (|2200\rangle + |0022\rangle + |1111\rangle). \quad (\text{D.75})$$

We propagate the three kets through the central beam-splitter, then weight them by the heralding POVM Π_{Lnk} . Recalling the two-photon beam-splitter relations

$$|20\rangle, |02\rangle \xrightarrow{\text{BS}} \frac{\sqrt{2}|20\rangle \pm 2|11\rangle + \sqrt{2}|02\rangle}{2\sqrt{2}}, \quad |11\rangle \xrightarrow{\text{BS}} \frac{|20\rangle - |02\rangle}{\sqrt{2}}, \quad (\text{D.76})$$

the heralded weights of the three components are

$$\begin{aligned}
|2200\rangle &\xrightarrow{\text{BS}} \begin{cases} \frac{1}{4} |2200\rangle \langle 2200| \xrightarrow{\Pi_{\text{Lnk}}} \frac{1}{4}(1 - R_{\text{Lnk}}^2) = \frac{\eta_{\text{Lnk}}(2-\eta_{\text{Lnk}})}{4} \\ \frac{1}{2} |2110\rangle \langle 2110| \xrightarrow{\Pi_{\text{Lnk}}} \frac{1}{2}R_{\text{Lnk}}(1 - R_{\text{Lnk}}) = \frac{\eta_{\text{Lnk}}(1-\eta_{\text{Lnk}})}{2} \\ \frac{1}{4} |2020\rangle \langle 2020| \xrightarrow{\Pi_{\text{Lnk}}} 0 \end{cases} \quad (\text{a}) \\
|0022\rangle &\xrightarrow{\text{BS}} \begin{cases} \frac{1}{4} |0202\rangle \langle 0202| \xrightarrow{\Pi_{\text{Lnk}}} \frac{\eta_{\text{Lnk}}(2-\eta_{\text{Lnk}})}{4} \\ \frac{1}{2} |0112\rangle \langle 0112| \xrightarrow{\Pi_{\text{Lnk}}} \frac{\eta_{\text{Lnk}}(1-\eta_{\text{Lnk}})}{2} \\ \frac{1}{4} |0022\rangle \langle 0022| \xrightarrow{\Pi_{\text{Lnk}}} 0 \end{cases} \quad (\text{b}) \quad (\text{D.77}) \\
|1111\rangle &\xrightarrow{\text{BS}} \begin{cases} \frac{1}{2} |1201\rangle \langle 1201| \xrightarrow{\Pi_{\text{Lnk}}} \frac{1}{2}(1 - R_{\text{Lnk}}^2) = \frac{\eta_{\text{Lnk}}(2-\eta_{\text{Lnk}})}{2} \\ \frac{1}{2} |1021\rangle \langle 1021| \xrightarrow{\Pi_{\text{Lnk}}} 0 \end{cases} \quad (\text{c})
\end{aligned}$$

Only the component (c) leaves a photon in each memory mode m_A and m_B , the configuration relevant to a coincidence in m_A .

Coincidence probability

A joint *click* on m_T, m_j requires at least two photons in m_A before the beam-splitter. Among the heralded second-order components, only (c) populates m_A with a single photon and the others with two, so the only term feeding two photons into m_A are the two kets from (a). The beam-splitter followed with the coincidence POVM leads to

$$|2\rangle \xrightarrow{\text{BS}} \begin{cases} \frac{1}{4} |20\rangle \langle 20| \xrightarrow{\Pi_{\text{Tj}}} 0 \\ \frac{1}{2} |11\rangle \langle 11| \xrightarrow{\Pi_{\text{Tj}}} \frac{\eta_{\text{Tj}}^2}{2} \\ \frac{1}{4} |02\rangle \langle 02| \xrightarrow{\Pi_{\text{Tj}}} 0 \end{cases} \quad (\text{D.78})$$

Collecting the heralded weight of (a) and the local coincidence weight

$$\mathbb{P}_{\text{Coinc}}^A = \frac{1}{\mathbb{P}_{\text{Lnk}}} \cdot \frac{\eta_{\text{Tj}}^2}{2} \cdot \frac{\eta_{\text{Lnk}}(4 - 3\eta_{\text{Lnk}})}{4}, \quad (\text{D.79})$$

where the factor $1/\mathbb{P}_{\text{Lnk}}$ is the elementary-link normalization as the coincidence is measured once the state has been heralded.

D.6 .Exact coincidence probability

In this appendix we give a non-perturbative derivation of the local coincidence probability $\mathbb{P}_{\text{Coinc}}^A$ and $\mathbb{P}_{\text{Coinc}}^B$ required for the evaluation of the Cauchy Schwartz inequality. As in appendix D.6, this probability corresponds to the joint *click* on the two output modes m_T (Tom) and m_J (Jerry), built from the output of a 50/50 beam-splitter placed after the memory's mode m_A or m_B .

Setup and context

As before, the setup is symmetric between the two hubs, thus Bob's coincidence probability $\mathbb{P}_{\text{Coinc}}^B$ follows the same derivation, up to the substitutions side efficiencies $dc_A, \eta_A \rightarrow dc_B, \eta_B$. We therefore focus on Alice's side m_A with the derivation of $\mathbb{P}_{\text{Coinc}}^A$. The modelization of the losses and dark-count have been pushed to the POVM element, consequently the efficiencies η_T, η_J and dark-counts dc_T, dc_J carry together the detection setup measuring the coincidence and the loss and dark-counts of Alice's memory itself, η_A and dc_A . For clarity, we here assume symmetric situation $\eta_T = \eta_J = \eta_{TJ}$ with no dark-count $dc_T = dc_J = 0$, as including these quantities does not require much more work while weighting the derivation. We call $\Pi_{TJ} = \Pi_1^{m_T}(\eta_{TJ}) \Pi_1^{m_J}(\eta_{TJ})$ the coincidence POVM, with $\Pi_1^{m_T}(\eta_{TJ}) = \mathbb{1} - R_{TJ}^{m_T^\dagger m_T}$ the click POVM on mode m_T (cf. appendix A.1), and similarly on m_J .

Strategy

This derivation fold the heralding and the coincidence measurement into a single POVM acting on the post-beam-splitter state ρ_{BS} of appendix D.1

$$\mathbb{P}_{\text{Coinc}}^A = \frac{1}{\mathbb{P}_{\text{Lnk}}} \text{Tr} \left[\Pi_{TJ} U_{BS}^{m_T, m_J} \Pi_{\text{Lnk}} U_{BS}^{m_A, m_B} \rho_{\text{Lnk}} (U_{BS}^{m_A, m_B})^\dagger U_{BS}^{m_T, m_J \dagger} \right] = \frac{1}{\mathbb{P}_{\text{Lnk}}} \text{Tr} [\Pi_{TJ} \Pi_{\text{Lnk}} \cdot \rho_{BS}], \quad (\text{D.80})$$

as $U_{BS}^{m_A, m_B}$ and $U_{BS}^{m_T, m_J \dagger}$ commute, with Π_{Lnk} the central click/no-click heralding POVM and $1/\mathbb{P}_{\text{Lnk}}$ the elementary-link normalization. The derivation mirrors appendix D.1: we build a non-unitary matrix M_i for each term of the expanded POVM, then evaluate its trace through the singular value decomposition of M_i .

Mode mixing

The local beam-splitter maps Alice's memory mode onto the two output modes

$$m_A^\dagger \longrightarrow \frac{m_T^\dagger + m_J^\dagger}{\sqrt{2}}, \quad (\text{D.81})$$

so the SPDC exponent, already mixed by the central beam-splitter, becomes

$$\begin{aligned}
m_A^\dagger c_A^\dagger + c_B^\dagger m_B^\dagger &\xrightarrow{\text{BS}_{\text{CS}}} m_A^\dagger \frac{c_A + c_B}{\sqrt{2}} + m_B^\dagger \frac{c_A - c_B}{\sqrt{2}} \\
&\xrightarrow{\text{BS}_A} \frac{m_T^\dagger + m_J^\dagger}{\sqrt{2}} \cdot \frac{c_A + c_B}{\sqrt{2}} + m_B^\dagger \frac{c_A - c_B}{\sqrt{2}} \\
&= \begin{pmatrix} m_T^\dagger & m_J^\dagger & m_B^\dagger & m_B^\dagger \end{pmatrix} \underbrace{\frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & -1 \\ 1 & -1 \end{pmatrix}}_M \begin{pmatrix} c_A^\dagger \\ c_B^\dagger \end{pmatrix}.
\end{aligned}$$

POVM expansion

Following appendix D.1, we expand $\Pi_{\text{TJ}} \Pi_{\text{Lnk}}$ into a sum of loss operators $R^{\alpha^\dagger \alpha}$

$$\begin{aligned}
\Pi_{\text{TJ}} \Pi_{\text{Lnk}} &= R_{\text{Lnk}}^{c_B^\dagger c_B} \left(\mathbb{1} - R_{\text{Lnk}}^{c_A^\dagger c_A} \right) \left(\mathbb{1} - R_{\text{TJ}}^{m_T^\dagger m_T} \right) \left(\mathbb{1} - R_{\text{TJ}}^{m_J^\dagger m_J} \right) \\
&= R_{\text{Lnk}}^{c_B^\dagger c_B} - R_{\text{Lnk}}^{c_B^\dagger c_B} R_{\text{Lnk}}^{c_A^\dagger c_A} - R_{\text{Lnk}}^{c_B^\dagger c_B} R_{\text{TJ}}^{m_T^\dagger m_T} - R_{\text{Lnk}}^{c_B^\dagger c_B} R_{\text{TJ}}^{m_J^\dagger m_J} \\
&\quad + R_{\text{Lnk}}^{c_B^\dagger c_B} R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{TJ}}^{m_T^\dagger m_T} + R_{\text{Lnk}}^{c_B^\dagger c_B} R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{TJ}}^{m_J^\dagger m_J} + R_{\text{Lnk}}^{c_B^\dagger c_B} R_{\text{TJ}}^{m_T^\dagger m_T} R_{\text{TJ}}^{m_J^\dagger m_J} - R_{\text{Lnk}}^{c_B^\dagger c_B} R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{TJ}}^{m_T^\dagger m_T} R_{\text{TJ}}^{m_J^\dagger m_J}.
\end{aligned} \tag{D.82}$$

Using $x^{\alpha^\dagger \alpha} f(\alpha^\dagger) = f(x \alpha^\dagger) x^{\alpha^\dagger \alpha}$, each term scales the corresponding column of M and defines a matrix M_i , whose trace follows from its singular value decomposition (cf. appendix A.2). For instance, the first term $R_{\text{Lnk}}^{c_A^\dagger c_A}$ gives

$$M_1 = \frac{1}{2} \begin{pmatrix} 1 & \sqrt{R_{\text{Lnk}}} \\ 1 & \sqrt{R_{\text{Lnk}}} \\ 1 & -\sqrt{R_{\text{Lnk}}} \\ 1 & -\sqrt{R_{\text{Lnk}}} \end{pmatrix} = U \begin{pmatrix} \sigma_+ & 0 \\ 0 & \sigma_- \\ 0 & 0 \\ 0 & 0 \end{pmatrix} V, \tag{D.83}$$

with singular values $\sigma_+^2 = R_{\text{Lnk}}$ and $\sigma_-^2 = 1$, leading to

$$\text{Tr} \left[R_{\text{Lnk}}^{c_A^\dagger c_A} \rho_{\text{BS}} \right] = \frac{(1-p)^2}{(1-pR_{\text{Lnk}})(1-p)} = \frac{1-p}{1-pR_{\text{Lnk}}}. \tag{D.84}$$

The matrices and singular values of the eight terms are collected in table D.1.

Coincidence probability

Summing the eight contributions yields the coincidence probability as follow

$$\begin{aligned}
\frac{\mathbb{P}^A}{(1-p)^2} \cdot \mathbb{P}_{\text{Lnk}} &= \frac{1}{(1-pR_{\text{Lnk}})(1-p)} - \frac{1}{(1-pR_{\text{Lnk}})^2} \\
&\quad - \frac{1 - \frac{p}{4} [3 + R_{\eta} + 3R_{\text{Lnk}} + R_{\text{Lnk}}R_{\eta}] + \frac{p^2}{64} [(3 + R_{\eta} + 3R_{\text{Lnk}} + R_{\text{Lnk}}R_{\eta})^2 + 32(1 + R_{\eta})R_{\text{Lnk}} - (3 + R_{\eta})^2(1 + R_{\text{Lnk}})^2]}{2} \\
&\quad + \frac{1}{2} \frac{1 - pR_{\text{Lnk}} \left(1 - p \frac{R_{\text{Lnk}}(1+R_{\eta})}{2}\right)}{(1-pR_{\text{Lnk}})^2} \\
&\quad + \frac{1}{1} \frac{1 - \frac{p}{2} [1 + R_{\text{Lnk}} + R_{\eta}(1 + R_{\text{Lnk}})] + \frac{p^2}{16} [(1 + R_{\text{Lnk}} + R_{\eta}(1 + R_{\text{Lnk}}))^2 + 16R_{\eta}R_{\text{Lnk}} - (1 + R_{\eta})^2(1 + R_{\text{Lnk}})^2]}{(1-pR_{\text{Lnk}}R_{\eta})(1-pR_{\text{Lnk}})} \\
&\quad - \frac{1}{(1-pR_{\text{Lnk}}R_{\eta})(1-pR_{\text{Lnk}})}
\end{aligned}$$

(D.85)

Term	Matrix M_i	Singular values
$R_{\text{Lnk}}^{c_A^\dagger c_A}$	$\frac{1}{2} \begin{pmatrix} 1 & \sqrt{R_{\text{Lnk}}} \\ 1 & \sqrt{R_{\text{Lnk}}} \\ 1 & -\sqrt{R_{\text{Lnk}}} \\ 1 & -\sqrt{R_{\text{Lnk}}} \end{pmatrix}$	$\sigma_+^2 = R_{\text{Lnk}}$ $\sigma_-^2 = 1$
$R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{Lnk}}^{c_B^\dagger c_B}$	$\frac{1}{2} \begin{pmatrix} \sqrt{R_{\text{Lnk}}} & \sqrt{R_{\text{Lnk}}} \\ \sqrt{R_{\text{Lnk}} R_{\text{Tj}}} & \sqrt{R_{\text{Lnk}} R_{\text{Tj}}} \\ \sqrt{R_{\text{Lnk}}} & -\sqrt{R_{\text{Lnk}}} \\ \sqrt{R_{\text{Lnk}}} & -\sqrt{R_{\text{Lnk}}} \end{pmatrix}$	$\sigma_+^2 = R_{\text{Lnk}}$ $\sigma_-^2 = R_{\text{Lnk}}$
$R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{Tj}}^{m_\text{T}^\dagger m_\text{T}}$	$\frac{1}{2} \begin{pmatrix} \sqrt{R_{\text{Tj}}} & \sqrt{R_{\text{Lnk}} R_{\text{Tj}}} \\ 1 & \sqrt{R_{\text{Lnk}}} \\ 1 & -\sqrt{R_{\text{Lnk}}} \\ 1 & -\sqrt{R_{\text{Lnk}}} \end{pmatrix}$	$\sigma_\pm^2 = \frac{1}{8} (3 + R_{\text{Tj}} + 3R_{\text{Lnk}} + R_{\text{Tj}} R_{\text{Lnk}} \pm \sqrt{-32(1 + R_{\text{Tj}})R_{\text{Lnk}} + (3 + R_{\text{Tj}})^2(1 + R_{\text{Lnk}})^2})$
$R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{Tj}}^{m_\text{j}^\dagger m_\text{j}}$	$\frac{1}{2} \begin{pmatrix} 1 & \sqrt{R_{\text{Lnk}}} \\ \sqrt{R_{\text{Tj}}} & \sqrt{R_{\text{Lnk}} R_{\text{Tj}}} \\ 1 & -\sqrt{R_{\text{Lnk}}} \\ 1 & -\sqrt{R_{\text{Lnk}}} \end{pmatrix}$	$\sigma_\pm^2 = \frac{1}{8} (3 + R_{\text{Tj}} + 3R_{\text{Lnk}} + R_{\text{Tj}} R_{\text{Lnk}} \pm \sqrt{-32(1 + R_{\text{Tj}})R_{\text{Lnk}} + (3 + R_{\text{Tj}})^2(1 + R_{\text{Lnk}})^2})$
$R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{Lnk}}^{c_B^\dagger c_B} R_{\text{Tj}}^{m_\text{T}^\dagger m_\text{T}}$	$\frac{1}{2} \begin{pmatrix} \sqrt{R_{\text{Lnk}}} & \sqrt{R_{\text{Lnk}}} \\ \sqrt{R_{\text{Tj}} R_{\text{Lnk}}} & \sqrt{R_{\text{Tj}} R_{\text{Lnk}}} \\ \sqrt{R_{\text{Lnk}}} & -\sqrt{R_{\text{Lnk}}} \\ \sqrt{R_{\text{Lnk}}} & -\sqrt{R_{\text{Lnk}}} \end{pmatrix}$	$\sigma_+^2 = \frac{R_{\text{Lnk}}(1 + R_{\text{Tj}})}{2}$ $\sigma_-^2 = R_{\text{Lnk}}$
$R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{Lnk}}^{c_B^\dagger c_B} R_{\text{Tj}}^{m_\text{j}^\dagger m_\text{j}}$	$\frac{1}{2} \begin{pmatrix} \sqrt{R_{\text{Tj}} R_{\text{Lnk}}} & \sqrt{R_{\text{Tj}} R_{\text{Lnk}}} \\ \sqrt{R_{\text{Lnk}}} & \sqrt{R_{\text{Lnk}}} \\ \sqrt{R_{\text{Lnk}}} & -\sqrt{R_{\text{Lnk}}} \\ \sqrt{R_{\text{Lnk}}} & -\sqrt{R_{\text{Lnk}}} \end{pmatrix}$	$\sigma_+^2 = \frac{R_{\text{Lnk}}(1 + R_{\text{Tj}})}{2}$ $\sigma_-^2 = R_{\text{Lnk}}$
$R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{Tj}}^{m_\text{T}^\dagger m_\text{T}} R_{\text{Tj}}^{m_\text{j}^\dagger m_\text{j}}$	$\frac{1}{2} \begin{pmatrix} \sqrt{R_{\text{Tj}}} & \sqrt{R_{\text{Tj}} R_{\text{Lnk}}} \\ \sqrt{R_{\text{Tj}}} & \sqrt{R_{\text{Tj}} R_{\text{Lnk}}} \\ 1 & -\sqrt{R_{\text{Lnk}}} \\ 1 & -\sqrt{R_{\text{Lnk}}} \end{pmatrix}$	$\sigma_\pm^2 = \frac{1}{4} (1 + R_{\text{Lnk}} + R_{\text{Tj}}(1 + R_{\text{Lnk}}) \pm \sqrt{-16R_{\text{Tj}}R_{\text{Lnk}} + (1 + R_{\text{Tj}})^2(1 + R_{\text{Lnk}})^2})$
$R_{\text{Lnk}}^{c_A^\dagger c_A} R_{\text{Lnk}}^{c_B^\dagger c_B} R_{\text{Tj}}^{m_\text{T}^\dagger m_\text{T}} R_{\text{Tj}}^{m_\text{j}^\dagger m_\text{j}}$	$\frac{1}{2} \begin{pmatrix} \sqrt{R_{\text{Tj}} R_{\text{Lnk}}} & \sqrt{R_{\text{Tj}} R_{\text{Lnk}}} \\ \sqrt{R_{\text{Tj}} R_{\text{Lnk}}} & \sqrt{R_{\text{Tj}} R_{\text{Lnk}}} \\ \sqrt{R_{\text{Lnk}}} & -\sqrt{R_{\text{Lnk}}} \\ \sqrt{R_{\text{Lnk}}} & -\sqrt{R_{\text{Lnk}}} \end{pmatrix}$	$\sigma_+^2 = R_{\text{Tj}} R_{\text{Lnk}}$ $\sigma_-^2 = R_{\text{Lnk}}$

Table D.1: Non-unitary matrices M_i and their squared singular values for each term of the expanded POVM $\Pi_{\text{Tj}} \Pi_{\text{Lnk}}$. Each M_i is the 4×2 matrix coupling the output modes ($m_\text{T}, m_\text{j}, m_\text{B}, m_\text{A}$) to the central modes (c_A, c_B), scaled by the loss operators of the corresponding term.

D.7 .Matrix elements of the displacement-based observable

This appendix derives the matrix elements of the single-mode displacement-based observable $\hat{o}_\alpha$ in the photon-number basis, restricted to the qubit subspace $\{|0\rangle, |1\rangle\}$. These elements are the building blocks of the entanglement witness w developed in appendix D.8 and appendix D.9.

Observable and setup

From appendix A.1, the POVM elements of a non-resolving-photodetector on mode a for a detector of efficiency η and dark-count rate dc read

$$\Pi_0^a(\eta, \text{dc}) = \kappa \cdot R^{a\dagger a}, \quad \Pi_1^a(\eta, \text{dc}) = \mathbb{1} - \kappa \cdot R^{a\dagger a}, \quad (\text{D.86})$$

with $R = 1 - \eta$ and $\kappa = 1 - \text{dc}$. The POVM elements in the displaced mode follows

$$\Pi_i^{a,\alpha}(\eta, \text{dc}) = \mathcal{D}^\dagger(\alpha) \Pi_i^a(\eta, \text{dc}) \mathcal{D}(\alpha), \quad i \in \{0, 1\}. \quad (\text{D.87})$$

Assigning the outcomes $+1$ to a no-click event and -1 to a click event, the single-mode observable reads

$$\hat{o}_\alpha = \Pi_0^{a,\alpha}(\eta, \text{dc}) - \Pi_1^{a,\alpha}(\eta, \text{dc}) \quad (\text{D.88})$$

$$= \mathcal{D}^\dagger(\alpha) \left(2\kappa \cdot R^{a\dagger a} - \mathbb{1} \right) \mathcal{D}(\alpha). \quad (\text{D.89})$$

The three matrix elements $\langle 0 | \hat{o}_\alpha | 0 \rangle$, $\langle 0 | \hat{o}_\alpha | 1 \rangle$, and $\langle 1 | \hat{o}_\alpha | 1 \rangle$ fully characterise $\hat{o}_\alpha$ in the qubit subspace, and the remaining element $\langle 1 | \hat{o}_\alpha | 0 \rangle$ is recovered by conjugation.

Action of $R^{a\dagger a}$ on a displaced state

The derivation relies on a single recurring identity. The operator $R^{a\dagger a}$ is diagonal in the Fock basis with eigenvalues R^n , so its action on the coherent state $\mathcal{D}(\alpha) | 0 \rangle = |\alpha\rangle$ reads

$$R^{a\dagger a} |\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n \geq 0} \frac{\alpha^n}{\sqrt{n!}} R^n |n\rangle \quad (\text{D.90})$$

$$= e^{-\eta|\alpha|^2/2} \left| \sqrt{R} \alpha \right\rangle. \quad (\text{D.91})$$

The loss R^n is absorbed in the ket by rewriting the series as the coherent state $\left| \sqrt{R} \alpha \right\rangle$ up to its normalisation factor $e^{-R|\alpha|^2/2}$. The scaling $\alpha \rightarrow \sqrt{R} \alpha$ inside the ket combined with the exponential prefactor $e^{-\eta|\alpha|^2/2}$ is the key relation used throughout the calculation.

Vacuum-vacuum element

The vacuum-vacuum element follows directly. Applying (D.89) to $|0\rangle$,

$$\langle 0 | \mathcal{D}^\dagger(\alpha) R^{a^\dagger a} \mathcal{D}(\alpha) | 0 \rangle = e^{-\eta|\alpha|^2/2} \langle 0 | \mathcal{D}^\dagger(\alpha) \left| \sqrt{R} \alpha \right\rangle \quad (\text{D.92})$$

$$= e^{-\eta|\alpha|^2/2} \left\langle -\alpha \left| \sqrt{R} \alpha \right\rangle e^{i \text{Im}[\alpha^* \sqrt{R} \alpha]} \quad (\text{D.93})$$

$$= e^{-\eta|\alpha|^2/2} e^{-\eta|\alpha|^2/2} \quad (\text{D.94})$$

$$= e^{-\eta|\alpha|^2}. \quad (\text{D.95})$$

Inserting this into (D.89),

$$\boxed{\langle 0 | \hat{o}_\alpha | 0 \rangle = 2\kappa \cdot e^{-\eta|\alpha|^2} - 1.} \quad (\text{D.96})$$

Vacuum-single element

For the off-diagonal element we use the commutation relation

$$\mathcal{D}(\alpha) |1\rangle = (a^\dagger - \alpha^*) \mathcal{D}(\alpha) |0\rangle. \quad (\text{D.97})$$

Expanding $\mathcal{D}(\alpha) |0\rangle = |\alpha\rangle$ in the Fock basis and rearranging the two contributions onto a common index n gives

$$\mathcal{D}(\alpha) |1\rangle = e^{-|\alpha|^2/2} \left[\sum_{n \geq 0} \frac{\alpha^n}{\sqrt{n!}} \sqrt{n+1} |n+1\rangle - \sum_{n \geq 0} |\alpha|^2 \frac{\alpha^{n-1}}{\sqrt{n!}} |n\rangle \right] \quad (\text{D.98})$$

$$= e^{-|\alpha|^2/2} \left[\sum_{n \geq 0} n \frac{\alpha^{n-1}}{\sqrt{n!}} |n\rangle - |\alpha|^2 \sum_{n \geq 0} \frac{\alpha^{n-1}}{\sqrt{n!}} |n\rangle \right] \quad (\text{D.99})$$

$$= e^{-|\alpha|^2/2} \sum_{n \geq 0} (n - |\alpha|^2) \frac{\alpha^{n-1}}{\sqrt{n!}} |n\rangle. \quad (\text{D.100})$$

The second line uses the re-indexing $\sum_{n \geq 0} \sqrt{n+1} u_{n+1} = \sum_{n \geq 0} \sqrt{n} u_n$, which aligns both sums on the same Fock index. Acting on $R^{a^\dagger a}$ leads to

$$R^{a^\dagger a} \mathcal{D}(\alpha) |1\rangle = e^{-|\alpha|^2/2} \sum_{n \geq 0} (n - |\alpha|^2) \frac{\alpha^{n-1}}{\sqrt{n!}} R^n |n\rangle. \quad (\text{D.101})$$

Projecting on $\langle 0 | \mathcal{D}^\dagger(\alpha) = \langle \alpha |$ collapses the sum onto the term $\langle \alpha | n \rangle = e^{-|\alpha|^2/2} (\alpha^*)^n / \sqrt{n!}$,

$$\langle 0 | \mathcal{D}^\dagger(\alpha) R^{a^\dagger a} \mathcal{D}(\alpha) | 1 \rangle = e^{-|\alpha|^2} \alpha^* \sum_{n \geq 0} (n - |\alpha|^2) \frac{(|\alpha|^2)^{n-1}}{n!} R^n \quad (\text{D.102})$$

$$= e^{-|\alpha|^2} \alpha^* \left[\sum_{n \geq 0} (n+1) \frac{(|\alpha|^2)^n}{(n+1)!} R^{n+1} - \sum_{n \geq 0} \frac{(|\alpha|^2)^n}{n!} R^n \right] \quad (\text{D.103})$$

$$= e^{-|\alpha|^2} \alpha^* \sum_{n \geq 0} (R-1) \frac{(R|\alpha|^2)^n}{n!} \quad (\text{D.104})$$

$$= -\eta \alpha^* e^{-\eta |\alpha|^2}. \quad (\text{D.105})$$

The second line splits the n -prefactor as $n = (n+1) - 1$ and uses the re-indexing $\sum_{n \geq 0} n u_n = \sum_{n \geq 0} (n+1) u_{n+1}$ to factor out a common power R^{n+1} , allowing both series to be summed as a single exponential $e^{R|\alpha|^2}$. We thus get

$$\boxed{\langle 0 | \hat{o}_\alpha | 1 \rangle = -2\kappa \cdot \eta \alpha^* e^{-\eta |\alpha|^2}.} \quad (\text{D.106})$$

Single-single element

For the diagonal single-photon element, projecting on $\langle 1 | \mathcal{D}^\dagger(\alpha)$ instead of $\langle 0 | \mathcal{D}^\dagger(\alpha)$ gives

$$\langle 1 | \mathcal{D}^\dagger(\alpha) R^{a^\dagger a} \mathcal{D}(\alpha) | 1 \rangle = e^{-|\alpha|^2} \sum_{n \geq 0} R^n (n - |\alpha|^2)^2 \frac{(|\alpha|^2)^{n-1}}{n!}. \quad (\text{D.107})$$

Expanding the square produces three separate series in n^2 , n , and a constant. Each of them admits a closed form, obtained by the same kind of re-indexing as in the previous element. Setting $x = |\alpha|^2$,

$$\sum_{n \geq 0} R^n \frac{x^n}{n!} = e^{Rx}, \quad (\text{D.108})$$

$$\sum_{n \geq 0} R^n n \frac{x^{n-1}}{n!} = \sum_{n \geq 0} R^{n+1} \frac{x^n}{n!} = R e^{Rx}, \quad (\text{D.109})$$

$$\sum_{n \geq 0} R^n n^2 \frac{x^{n-1}}{n!} = \sum_{n \geq 0} R^{n+1} (n+1) \frac{x^n}{n!} = R \sum_{n \geq 0} R^n n \frac{x^n}{n!} + R e^{Rx} = (R^2 x + R) e^{Rx}. \quad (\text{D.110})$$

Substituting these three results,

$$\langle 1 | \mathcal{D}^\dagger(\alpha) R^{a^\dagger a} \mathcal{D}(\alpha) | 1 \rangle = \left[R^2 |\alpha|^2 + R - 2|\alpha|^2 R + |\alpha|^2 \right] e^{-\eta |\alpha|^2} \quad (\text{D.111})$$

$$= \left[1 + \eta \left(\eta |\alpha|^2 - 1 \right) \right] e^{-\eta |\alpha|^2}, \quad (\text{D.112})$$

leading to

$$\boxed{\langle 1 | \hat{\sigma}_\alpha | 1 \rangle = 2\kappa \cdot \left[1 + \eta \left(\eta |\alpha|^2 - 1 \right) \right] e^{-\eta |\alpha|^2} - 1.} \quad (\text{D.113})$$

Summary

The three matrix elements of $\hat{\sigma}_\alpha$ in the qubit subspace read

$$\begin{aligned} \langle 0 | \hat{\sigma}_\alpha | 0 \rangle &= 2\kappa \cdot e^{-\eta |\alpha|^2} - 1, \\ \langle 0 | \hat{\sigma}_\alpha | 1 \rangle &= -2\kappa \cdot \eta \alpha^* e^{-\eta |\alpha|^2}, \\ \langle 1 | \hat{\sigma}_\alpha | 1 \rangle &= 2\kappa \cdot \left[1 + \eta \left(\eta |\alpha|^2 - 1 \right) \right] e^{-\eta |\alpha|^2} - 1, \end{aligned} \quad (\text{D.114})$$

with $\langle 1 | \hat{\sigma}_\alpha | 0 \rangle = \langle 0 | \hat{\sigma}_\alpha | 1 \rangle^*$.

D.8 .Entanglement witness in the qubit subspace

This appendix derives the closed-form expression of the entanglement witness w within the qubit subspace $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$. The full witness extending beyond the qubit subspace and accounting for higher photon-number contributions is treated in appendix D.9.

Setup and notations

Let's consider the ideal low-pump scenario where the elementary link carries the Bell state such that

$$\rho_{\text{Lnk}} = \frac{1}{2} \left[|01\rangle \langle 01| + |10\rangle \langle 10| + |01\rangle \langle 10| + |10\rangle \langle 01| \right]. \quad (\text{D.115})$$

Let's also consider the displacement-based observable on the bipartite memory modes m_A and m_B , with displacement amplitudes α and β , efficiencies η_A and η_B , and dark-count rates dc_A and dc_B . The joint observable reads

$$\mathcal{O}_{\alpha,\beta} = \hat{\sigma}_{\alpha}^{m_A} \otimes \hat{\sigma}_{\beta}^{m_B}, \quad (\text{D.116})$$

where $\hat{\sigma}_{\alpha}^{m_A}$ acts on mode m_A with parameters an efficiency η_A and a dark-count dc_A , and similarly for $\hat{\sigma}_{\beta}^{m_B}$. The phase-averaged observable reads

$$\overline{\mathcal{O}}_{\alpha,\beta} = \frac{1}{2\pi} \int_0^{2\pi} e^{i\phi(m_A^\dagger m_A + m_B^\dagger m_B)} \mathcal{O}_{\alpha,\beta} e^{-i\phi(m_A^\dagger m_A + m_B^\dagger m_B)} d\phi. \quad (\text{D.117})$$

From this observable, we build a witness w associated to the setting (α, β)

$$w_{\rho_{\text{Lnk}}}(\alpha, \beta) = \langle \overline{\mathcal{O}}_{\alpha,\beta} \rangle_{\rho_{\text{Lnk}}} - w_{\text{PPT}}^{\rho_{\text{Lnk}}}(\alpha, \beta), \quad (\text{D.118})$$

where w_{PPT} defines the separability boundary

$$\begin{aligned} w_{\text{PPT}}^{\rho_{\text{Lnk}}}(\alpha, \beta) &= \max_{\sigma \succeq 0} \langle \overline{\mathcal{O}}_{\alpha,\beta} \rangle_{\sigma} \\ \text{s.t.} \quad & \text{(a) } \sigma^{T_A} \succeq 0, \\ & \text{(b) } \sigma_{ii} = \mathbb{P}_{\text{Loc}}^{ii}, \end{aligned} \quad (\text{D.119})$$

with (a) the PPT criterion, and (b) an additional assumption on the state where the diagonal terms match the actual local probabilities. For clarity, we simply write w and w_{PPT} keeping in mind that these quantities are defined for ρ_{Lnk} with the setting (α, β) .

Expectation value

First, let's compute the expectation value of $\overline{\mathcal{O}}$ for the state ρ_{Lnk} . The phase averaging keeps only the n -photon sectors non-zero, concretely $\langle i, j | \overline{\mathcal{O}} | k, l \rangle = 0$ for $i + j \neq k + l$.

Thus, the observable takes the block-diagonal form

$$\overline{\mathcal{O}}_{\alpha,\beta} = \begin{pmatrix} o_{00} & 0 & 0 & 0 \\ 0 & o_{01} & o'_{01} & 0 \\ 0 & o'_{10} & o_{10} & 0 \\ 0 & 0 & 0 & o_{11} \end{pmatrix}, \quad (\text{D.120})$$

with the convention $o_{ij} = \langle ii | \overline{\mathcal{O}}_{\alpha,\beta} | jj \rangle$ on the diagonal and $o'_{ij} = \langle ij | \overline{\mathcal{O}}_{\alpha,\beta} | ji \rangle$ off the diagonal. From this decomposition, the expectation value of $\overline{\mathcal{O}}_{\alpha,\beta}$ on ρ_{Lnk} reads

$$\langle \overline{\mathcal{O}}_{\alpha,\beta} \rangle_{\rho_{\text{Lnk}}} = \sum_{ij} o_{ij} \rho_{\text{Lnk}}^{ij} + o'_{01} \langle 10 | \rho_{\text{Lnk}} | 01 \rangle + o'_{10} \langle 01 | \rho_{\text{Lnk}} | 10 \rangle, \quad (\text{D.121})$$

with the standard convention $\rho_{\text{Lnk}}^{ij} = \langle ij | \rho_{\text{Lnk}} | ij \rangle$.

PPT separability bound

Let's derive an expression of w_{PPT} by maximising $\langle \overline{\mathcal{O}}_{\alpha,\beta} \rangle_{\sigma}$ over all states σ satisfying the two constraints (a) and (b) from equation (D.119). Based the previous computation, the expectation value reads

$$\langle \overline{\mathcal{O}}_{\alpha,\beta} \rangle_{\sigma} = \sum_{ij} o_{ij} \sigma_{ij} + o'_{01} \langle 10 | \sigma | 01 \rangle + o'_{10} \langle 01 | \sigma | 10 \rangle. \quad (\text{D.122})$$

The assumption (b) fixes the diagonal terms to the local probabilities of ρ_{Lnk} such that $\sigma_{ij} = \mathbb{P}_{\text{Loc}}^{ij}$, therefore the first sum is constant under and the optimisation reduces to the off-diagonal terms. Standard complex inequalities give

$$o'_{01} \langle 10 | \sigma | 01 \rangle + o'_{10} \langle 01 | \sigma | 10 \rangle = 2 \text{Re}[o'_{01} \langle 10 | \sigma | 01 \rangle] \leq 2 |o'_{01}| |\langle 10 | \sigma | 01 \rangle|, \quad (\text{D.123})$$

where the upper bound is reached when o'_{01} and $\langle 10 | \sigma | 01 \rangle$ are collinear to the real axis. The value of $|\langle 01 | \sigma | 10 \rangle|$ is bounded using positivity and the Cauchy-Schwarz inequality

$$\begin{aligned} \text{Positivity of } \sigma : |\langle 01 | \sigma | 10 \rangle| &\leq \sqrt{\sigma_{01} \sigma_{10}} = \sqrt{\mathbb{P}_{\text{Loc}}^{01} \mathbb{P}_{\text{Loc}}^{10}}, \\ \text{Positivity of } \sigma^{TA} : |\langle 01 | \sigma | 10 \rangle| &\leq \sqrt{\sigma_{00} \sigma_{11}} = \sqrt{\mathbb{P}_{\text{Loc}}^{00} \mathbb{P}_{\text{Loc}}^{11}}, \end{aligned} \quad (\text{D.124})$$

where the second bound follows from the partial transposition $\langle 01 | \sigma^{TA} | 10 \rangle = \langle 11 | \sigma | 00 \rangle$, mapping the off-diagonal block to the diagonal block of σ . The effective bound is the minimum of the two, leading to

$$|\langle 01 | \sigma | 10 \rangle| \leq \min \left(\sqrt{\mathbb{P}_{\text{Loc}}^{00} \mathbb{P}_{\text{Loc}}^{11}}, \sqrt{\mathbb{P}_{\text{Loc}}^{01} \mathbb{P}_{\text{Loc}}^{10}} \right), \quad (\text{D.125})$$

which finally gives

$$w_{\text{PPT}} = \sum_{ij} o_{ij} \mathbb{P}_{\text{Loc}}^{ij} + 2|o'_{01}| \cdot \min \left(\sqrt{\mathbb{P}_{\text{Loc}}^{00} \mathbb{P}_{\text{Loc}}^{11}}, \sqrt{\mathbb{P}_{\text{Loc}}^{01} \mathbb{P}_{\text{Loc}}^{10}} \right). \quad (\text{D.126})$$

Witness

Combining equations (D.119), (D.121) and (D.126), the witness reads

$$w = o'_{01} \langle 10 | \rho_{\text{Lnk}} | 01 \rangle + o'_{10} \langle 01 | \rho_{\text{Lnk}} | 10 \rangle - 2|o'_{01}| \cdot \min \left(\sqrt{\mathbb{P}_{\text{Loc}}^{00} \mathbb{P}_{\text{Loc}}^{11}}, \sqrt{\mathbb{P}_{\text{Loc}}^{01} \mathbb{P}_{\text{Loc}}^{10}} \right). \quad (\text{D.127})$$

From the Bell-pair structure of ρ_{Lnk} , the local probabilities and coherences read

$$\begin{aligned} \mathbb{P}_{\text{Loc}}^{00} &= \frac{1}{2} \kappa_A \kappa_B (R_A + R_B), \\ \mathbb{P}_{\text{Loc}}^{01} &= \frac{1}{2} \kappa_A (R_A + 1) - \frac{1}{2} \kappa_A \kappa_B (R_A + R_B), \\ \mathbb{P}_{\text{Loc}}^{10} &= \frac{1}{2} \kappa_B (R_B + 1) - \frac{1}{2} \kappa_A \kappa_B (R_A + R_B), \\ \mathbb{P}_{\text{Loc}}^{11} &= 1 - \mathbb{P}_{\text{Loc}}^{00} - \mathbb{P}_{\text{Loc}}^{01} - \mathbb{P}_{\text{Loc}}^{10}, \\ \langle 01 | \rho_{\text{Lnk}} | 10 \rangle &= \frac{1}{2}, \quad \langle 10 | \rho_{\text{Lnk}} | 01 \rangle = \frac{1}{2}. \end{aligned} \quad (\text{D.128})$$

From appendix D.7, the off-diagonal matrix elements of $\overline{\mathcal{O}}_{\alpha, \beta}$ read

$$\begin{aligned} o'_{01} &= \langle 0 | \hat{\sigma}_\alpha^{m_A} | 1 \rangle \cdot \langle 1 | \hat{\sigma}_\beta^{m_B} | 0 \rangle = 4\kappa_A \kappa_B \cdot \eta_A \eta_B \cdot \alpha^* \beta \cdot e^{-\eta_A |\alpha|^2 - \eta_B |\beta|^2}, \\ o'_{10} &= \langle 1 | \hat{\sigma}_\alpha^{m_A} | 0 \rangle \cdot \langle 0 | \hat{\sigma}_\beta^{m_B} | 1 \rangle = 4\kappa_A \kappa_B \cdot \eta_A \eta_B \cdot \alpha \beta^* \cdot e^{-\eta_A |\alpha|^2 - \eta_B |\beta|^2}, \end{aligned} \quad (\text{D.129})$$

with the sum

$$o'_{01} \langle 10 | \rho_{\text{Lnk}} | 01 \rangle + o'_{10} \langle 01 | \rho_{\text{Lnk}} | 10 \rangle = \frac{1}{2} (o'_{01} + o'_{10}) = 4\kappa_A \kappa_B \cdot \eta_A \eta_B \cdot \text{Re}[\alpha^* \beta] \cdot e^{-\eta_A |\alpha|^2 - \eta_B |\beta|^2}. \quad (\text{D.130})$$

Substituting these expressions into equation (D.127), the in-qubit witness reads

$$w = 4\kappa_A \kappa_B \cdot \eta_A \eta_B \cdot \text{Re}[\alpha^* \beta] \cdot e^{-\eta_A |\alpha|^2 - \eta_B |\beta|^2} - 2|o'_{01}| \cdot \min \left(\sqrt{\mathbb{P}_{\text{Loc}}^{00} \mathbb{P}_{\text{Loc}}^{11}}, \sqrt{\mathbb{P}_{\text{Loc}}^{01} \mathbb{P}_{\text{Loc}}^{10}} \right). \quad (\text{D.131})$$

The simple situation with no dark-count and symmetric efficiencier $\eta_A = \eta_B = \eta_{\text{AB}}$ leads to

$$w = 4\eta_{\text{AB}}^2 \cdot \text{Re}[\alpha^* \beta] \cdot e^{-\eta_{\text{AB}}(|\alpha|^2 + |\beta|^2)}, \quad (\text{D.132})$$

since $\mathbb{P}_{\text{Loc}}^{11} = 0$. In this context, the witness is positive whenever $\text{Re}[\alpha^* \beta] > 0$, which is the only sign condition on (α, β) required for the witness to detect entanglement. The factor η_{AB} exposes the second-order sensitivity of the qubit-restricted witness to mode losses, while the prefactor $\kappa_A \kappa_B$ captures its first-order degradation under dark counts. The exponential prefactor sets the optimal displacement amplitudes around $\eta_A |\alpha|^2 \approx \eta_B |\beta|^2 \approx 1/2$.

D.9 .Entanglement witness beyond the qubit subspace

This appendix extends the entanglement witness w derived in appendix D.8 to the whole Fock space, taking into account contribution of arbitrary number of excitations. The qubit-restricted derivation assumes ρ_{Lnk} lives in $\mathcal{H}^{\text{In}} = \text{Span}\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$, which holds only in the low-pump limit. Beyond this limit, ρ_{Lnk} has non-zero support in $\mathcal{H}^{\text{Out}}$ involving multi-photon states. The contributions from $\mathcal{H}^{\text{Out}}$ brings a correction to the PPT inequality, yielding in a valid witness over the full Hilbert space $\mathcal{H} = \mathcal{H}^{\text{In}} \oplus \mathcal{H}^{\text{Out}}$.

Setup and block decomposition

Let Π_{In} and Π_{Out} denote the orthogonal projectors onto $\mathcal{H}^{\text{In}}$ and $\mathcal{H}^{\text{Out}}$. Any operator X on $\mathcal{H}$ decomposes into four blocks

$$X = \Pi_{\text{In}} X \Pi_{\text{In}} + \Pi_{\text{In}} X \Pi_{\text{Out}} + \Pi_{\text{Out}} X \Pi_{\text{In}} + \Pi_{\text{Out}} X \Pi_{\text{Out}}, \quad (\text{D.133})$$

which we write in matrix form on $\mathcal{H}^{\text{In}} \oplus \mathcal{H}^{\text{Out}}$ as

$$\rho_{\text{Lnk}} = \begin{pmatrix} \rho_{\text{Lnk}}^{\text{In}} & \rho_{\text{Lnk}}^{\star} \\ \rho_{\text{Lnk}}^{\star \dagger} & \rho_{\text{Lnk}}^{\text{Out}} \end{pmatrix}, \quad \overline{\mathcal{O}} = \begin{pmatrix} \mathcal{O}^{\text{In}} & \mathcal{O}^{\star} \\ \mathcal{O}^{\star \dagger} & \mathcal{O}^{\text{Out}} \end{pmatrix}, \quad (\text{D.134})$$

where the diagonal blocks $\rho_{\text{Lnk}}^{\text{In}}$, $\mathcal{O}^{\text{In}}$ and $\rho_{\text{Lnk}}^{\text{Out}}$, $\mathcal{O}^{\text{Out}}$ are respectively 4×4 and infinite-dimensional block matrices acting on $\mathcal{H}^{\text{In}}$ and $\mathcal{H}^{\text{Out}}$. The off-diagonal blocks $\rho_{\text{Lnk}}^{\star}$, $\mathcal{O}^{\star}$ are the $4 \times \infty$ coupling blocks. The expectation value of $\overline{\mathcal{O}}$ splits accordingly

$$\langle \overline{\mathcal{O}} \rangle_{\rho_{\text{Lnk}}} = \underbrace{\text{tr}(\rho_{\text{Lnk}}^{\text{In}} \mathcal{O}^{\text{In}})}_{\text{in-qubit}} + \underbrace{\text{tr}(\rho_{\text{Lnk}}^{\star} \mathcal{O}^{\star \dagger}) + \text{tr}(\rho_{\text{Lnk}}^{\star \dagger} \mathcal{O}^{\star})}_{\text{coupling}} + \underbrace{\text{tr}(\rho_{\text{Lnk}}^{\text{Out}} \mathcal{O}^{\text{Out}})}_{\text{out-qubit}}. \quad (\text{D.135})$$

The first term is bounded by the in-qubit construction of appendix D.8, called $^{\text{In}}w_{\text{PPT}}$. The remaining two terms are bounded separately below, leading to a correction $^{\text{Out}}w_{\text{PPT}}$ from the contribution outside of the qubit space. Finally the total separability boundary reads $w_{\text{PPT}} = ^{\text{In}}w_{\text{PPT}} + ^{\text{Out}}w_{\text{PPT}}$.

Out-qubit diagonal contribution

The block $\rho_{\text{Lnk}}^{\text{Out}}$ is positive semidefinite as a diagonal block of a positive operator, and $\mathcal{O}^{\text{Out}}$ is Hermitian. The trace of a product of a positive matrix and a Hermitian gives the bound

$$\text{tr}(\rho_{\text{Lnk}}^{\text{Out}} \mathcal{O}^{\text{Out}}) \leq \|\mathcal{O}^{\text{Out}}\|_{\infty} \cdot \text{tr}(\rho_{\text{Lnk}}^{\text{Out}}) \leq \|\overline{\mathcal{O}}\|_{\infty} \cdot \text{tr}(\rho_{\text{Lnk}}^{\text{Out}}), \quad (\text{D.136})$$

where $\|\cdot\|_{\infty}$ denotes the largest eigenvalue in absolute value, and the second inequality uses the fact that the spectral norm of a diagonal block cannot exceed the one of the full

matrix. Since $\overline{\mathcal{O}}$ is a product of unitaries and operators with spectrum bounded by 1, we have $\|\overline{\mathcal{O}}\|_\infty \leq 1$, thus

$$\text{tr}(\rho_{\text{Lnk}}^{\text{Out}} \mathcal{O}^{\text{Out}}) \leq T_c. \quad (\text{D.137})$$

with $T_c = \text{tr}(\rho_{\text{Lnk}}^{\text{Out}})$. Now let's bound T_c with experimental quantities. Expanding the trace over Fock states $|i, j\rangle$ with $i + j \geq 2$ and either $i \geq 2$ or $j \geq 2$,

$$\begin{aligned} T_c &= \sum_{i \geq 2, j} \mathbb{P}_{\text{Loc}}^{ij} + \sum_{i, j \geq 2} \mathbb{P}_{\text{Loc}}^{ij} - \sum_{i \geq 2, j \geq 2} \mathbb{P}_{\text{Loc}}^{ij} \\ &= \mathbb{P}(A > 1) + \mathbb{P}(B > 1) - \mathbb{P}(A > 1, B > 1) \\ &\leq \mathbb{P}(A > 1) + \mathbb{P}(B > 1), \end{aligned} \quad (\text{D.138})$$

where $\mathbb{P}(A > 1)$ and $\mathbb{P}(B > 1)$ are respectively the marginal probability of having strictly more than one photon on Alice's m_A mode as dBob's mode m_B . From appendix D.10, we show both marginals probabilities $\mathbb{P}(A > 1)$ and $\mathbb{P}(B > 1)$ are bounded by coincidence detection after a 50/50 beam splitter, which yields in a bound $\mathbb{P}(A > 1) + \mathbb{P}(B > 1) \leq p_A^* + p_B^* = p^*$ such that

$$\text{tr}(\rho_{\text{Lnk}}^{\text{Out}} \mathcal{O}^{\text{Out}}) \leq p^*, \quad (\text{D.139})$$

with $p_A^* = 2 \mathbb{P}_{\text{Coinc}}^A$ and $p_B^* = 2 \mathbb{P}_{\text{Coinc}}^B$, the coincidence detections for Alice and Bob's arm.

Coupling contribution

The coupling term $\text{tr}(\rho_{\text{Lnk}}^* \mathcal{O}^{\star\dagger}) + \text{tr}(\rho_{\text{Lnk}}^{\star\dagger} \mathcal{O}^*) = 2 \text{Re}[\text{tr}(\rho_{\text{Lnk}}^* \mathcal{O}^{\star\dagger})]$ requires a finer treatment. The argument combines three steps: a singular value decomposition of ρ_{Lnk}^* to reduce the trace to a sum over singular values, a bound on the singular values of $\mathcal{O}^{\star\dagger}$ from the phase-averaging structure of $\overline{\mathcal{O}}$, and a final bound on the sum of singular values of ρ_{Lnk}^* using positivity of ρ_{Lnk} .

Reduction via SVD of ρ_{Lnk}^* The coupling block ρ_{Lnk}^* is a $4 \times \infty$ matrix with at most 4 non-zero singular values. Its singular value decomposition reads

$$\rho_{\text{Lnk}}^* = \sum_{i=1}^4 s_i |u_i\rangle \langle v_i|, \quad (\text{D.140})$$

with $s_i \geq 0$ the singular values, $\{|u_i\rangle\} \subset \mathcal{H}^{\text{In}}$ and $\{|v_i\rangle\} \subset \mathcal{H}^{\text{Out}}$ orthonormal singular vectors. Substituting into the trace,

$$\left| \text{tr}(\rho_{\text{Lnk}}^* \mathcal{O}^{\star\dagger}) \right| = \left| \sum_{i=1}^4 s_i \langle v_i | \mathcal{O}^{\star\dagger} | u_i \rangle \right| \leq \left(\sum_{i=1}^4 s_i \right) \cdot S_{\max}, \quad (\text{D.141})$$

where $S_{\max}$ is the largest singular value of $\mathcal{O}^*$, the s_i are positive by construction, and the final inequality uses $|\langle v_i | \mathcal{O}^{\star\dagger} | u_i \rangle| \leq \|\mathcal{O}^*\|_\infty = S_{\max}$.

Largest singular value of $\mathcal{O}^*$

The phase averaging of $\overline{\mathcal{O}}$ kills all matrix elements coupling blocks with different total photon number, $\langle ij | \overline{\mathcal{O}} | kl \rangle = 0$ for $i + j \neq k + l$. The block $\mathcal{O}^*$ connects $\mathcal{H}^{\text{In}}$ composed of blocks with a total photon number 0, 1, or 2 from $|11\rangle$, and $\mathcal{H}^{\text{Out}}$ composed of blocks with a total photon number ≥ 2 excluding $|11\rangle$. Therefore, the only non-zero matrix elements in $\mathcal{O}^*$ are those coupling $|11\rangle$ to $|02\rangle$ and $|20\rangle$, both in the 2-photon sector. The matrix $\mathcal{O}^*$ thus has rank at most 1 and its largest singular value reads

$$S_{\max} = \sqrt{|\langle 11 | \overline{\mathcal{O}} | 02 \rangle|^2 + |\langle 11 | \overline{\mathcal{O}} | 20 \rangle|^2}. \quad (\text{D.142})$$

The explicit computation of $\langle 11 | \overline{\mathcal{O}} | 02 \rangle$ and $\langle 11 | \overline{\mathcal{O}} | 20 \rangle$ extends the matrix-element calculation of appendix D.7 to the two-photon Fock state. The final expression reads

$$S_{\max} = 2\kappa_A\kappa_B \cdot \eta_A\eta_B \cdot e^{-\eta_A|\alpha|^2 - \eta_B|\beta|^2} \cdot \sqrt{2|\lambda_{\eta_A}(\alpha)|^2 + 2|\lambda_{\eta_B}(\beta)|^2}, \quad (\text{D.143})$$

with the auxiliary function

$$\lambda_\eta(x) = 2(1 - \eta) + |x|^2(2 - \eta^2). \quad (\text{D.144})$$

In the lossless and noiseless limit $\eta \rightarrow 1$, $\text{dc} \rightarrow 0$, $\lambda_\eta(x)$ reduces to $|x|^2$ and $S_{\max}$ recovers the expression of [122].

Bound on $\sum_i s_i$ To bound the sum of singular values of ρ_{Lnk}^* , we exploit the positivity of the full state ρ_{Lnk} . For a parameter $\theta \in (0, \pi/2)$ and each $i \in \{1, \dots, 4\}$, we define the mixed orthonormal vector

$$|\psi_i^\theta\rangle = \cos \theta |u_i\rangle + \sin \theta |v_i\rangle, \quad (\text{D.145})$$

where $|u_i\rangle, |v_i\rangle$ are the singular vectors from (D.140). The orthonormality of $\{|u_i\rangle\}$ and $\{|v_i\rangle\}$ ensures $\langle \psi_i^\theta | \psi_j^\theta \rangle = \delta_{ij}$, thus $\{|\psi_i^\theta\rangle\}_{i=1}^4$ is a non-generatrice orthonormal family in the total space $\mathcal{H}$, hence $\sum_i |\psi_i^\theta\rangle \langle \psi_i^\theta| \preceq \mathbb{1}$. Combined with $\rho_{\text{Lnk}} \succeq 0$, we get

$$\sum_{i=1}^4 \langle \psi_i^\theta | \rho_{\text{Lnk}} | \psi_i^\theta \rangle \leq \text{tr}(\rho_{\text{Lnk}}) = 1. \quad (\text{D.146})$$

Expanding each term using the block structure (D.134),

$$\langle \psi_i^\theta | \rho_{\text{Lnk}} | \psi_i^\theta \rangle = \cos^2 \theta \langle u_i | \rho_{\text{Lnk}}^{\text{In}} | u_i \rangle + \sin^2 \theta \langle v_i | \rho_{\text{Lnk}}^{\text{Out}} | v_i \rangle + 2 \cos \theta \sin \theta \text{Re}[\langle u_i | \rho_{\text{Lnk}}^* | v_i \rangle]. \quad (\text{D.147})$$

Summing over i and using $\langle u_i | \rho_{\text{Lnk}}^* | v_i \rangle = s_i$ from the SVD,

$$\cos^2 \theta \sum_i \langle u_i | \rho_{\text{Lnk}}^{\text{In}} | u_i \rangle + \sin^2 \theta \sum_i \langle v_i | \rho_{\text{Lnk}}^{\text{Out}} | v_i \rangle + 2 \cos \theta \sin \theta \sum_i s_i \leq 1. \quad (\text{D.148})$$

The first two sums are bounded by $\text{tr}(\rho_{\text{Lnk}}^{\text{In}}) = 1 - T_c$ and $\text{tr}(\rho_{\text{Lnk}}^{\text{Out}}) = T_c$ respectively, since $\{|u_i\rangle\}$ and $\{|v_i\rangle\}$ are orthonormal but incomplete in their respective subspaces. Rearranging,

$$\sum_{i=1}^4 s_i \leq \frac{1 - \cos^2 \theta \cdot (1 - T_c) - \sin^2 \theta \cdot T_c}{2 \cos \theta \sin \theta} \equiv f_{T_c}(\theta). \quad (\text{D.149})$$

Minimising $f_{T_c}(\theta)$ over θ yields the optimal value $\theta_0 = \arctan\left(\sqrt{T_c/(1 - T_c)}\right)$, giving

$$\sum_{i=1}^4 s_i \leq \sqrt{T_c(1 - T_c)}. \quad (\text{D.150})$$

Since $x \mapsto \sqrt{x(1 - x)}$ is monotonically increasing on $[0, 1/2]$, and assuming $p^* \leq 1/2$, which is the case in practice, the bound $T_c \leq p^*$ gives the final inequality

$$\sum_{i=1}^4 s_i \leq \sqrt{p^*(1 - p^*)}. \quad (\text{D.151})$$

Combining the two bounds

$$\text{tr}(\rho_{\text{Lnk}}^* \mathcal{O}^{\dagger}) + \text{tr}(\rho_{\text{Lnk}}^{\dagger} \mathcal{O}^*) \leq 4S_{\max} \sqrt{\mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}} (tfrac{1}{2} - \mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}})}. \quad (\text{D.152})$$

with $p_{\text{A}}^* = 2\mathbb{P}_{\text{Coinc}}^{\text{A}}$ and $p_{\text{B}}^* = 2\mathbb{P}_{\text{Coinc}}^{\text{B}}$, the coincidence detections for Alice and Bob's arm as developed in appendix D.10.

Out-qubit separability boundary and witness

Combining the on-diagonal bound (D.139) and off-diagonal coupling bound (D.152), the out-qubit contribution reads

$$\boxed{{}^{\text{Out}}w_{\text{PPT}} = 2(\mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}}) + 4S_{\max} \sqrt{\mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}} (tfrac{1}{2} - \mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}})}}. \quad (\text{D.153})$$

standing for state ρ_{Lnk} on $\mathcal{H}$ provided $\mathbb{P}_{\text{Coinc}}^{\text{A}} \leq 1/4$, which is respected in practice. The out-qubit correction ${}^{\text{Out}}w_{\text{PPT}}$ vanishes when $\mathbb{P}_{\text{Coinc}}^{\text{A}} = \mathbb{P}_{\text{Coinc}}^{\text{B}} = 0$, recovering the in-qubit witness as expected.

Total witness

Finally, the witness reads

$$w = \langle \overline{\mathcal{O}} \rangle_{\rho_{\text{Lnk}}} - {}^{\text{In}}w_{\text{PPT}} - {}^{\text{Out}}w_{\text{PPT}}, \quad (\text{D.154})$$

with the contirbutions

$${}^{\text{In}}w_{\text{PPT}} = \sum_{ij} o_{ij} \mathbb{P}_{\text{Loc}}^{ij} + 2|o'_{01}| \cdot \min \left(\sqrt{\mathbb{P}_{\text{Loc}}^{00} \mathbb{P}_{\text{Loc}}^{11}}, \sqrt{\mathbb{P}_{\text{Loc}}^{01} \mathbb{P}_{\text{Loc}}^{10}} \right) \quad (\text{D.155})$$

$${}^{\text{Out}}w_{\text{PPT}} = 2(\mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}}) + 4S_{\text{max}} \sqrt{\mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}} \left(\frac{1}{2} - \mathbb{P}_{\text{Coinc}}^{\text{A}} + \mathbb{P}_{\text{Coinc}}^{\text{B}} \right)},$$

and the matrix elements

$$\begin{aligned} \langle 0 | \hat{o}_\alpha | 0 \rangle &= 2\kappa \cdot e^{-\eta|\alpha|^2} - 1, \\ \langle 0 | \hat{o}_\alpha | 1 \rangle &= -2\kappa \cdot \eta \alpha^* e^{-\eta|\alpha|^2}, \\ \langle 1 | \hat{o}_\alpha | 1 \rangle &= 2\kappa \cdot \left[1 + \eta \left(\eta|\alpha|^2 - 1 \right) \right] e^{-\eta|\alpha|^2} - 1, \end{aligned} \quad (\text{D.156})$$

The total PPT boundary on the full Hilbert space follows by additivity

$$w_{\text{PPT}} = {}^{\text{In}}w_{\text{PPT}} + {}^{\text{Out}}w_{\text{PPT}}, \quad (\text{D.157})$$

with ${}^{\text{In}}w_{\text{PPT}}$ from equation (D.126).

D.10 .Bounding multi-photon marginals via twofold coincidences

This appendix derives an experimentally accessible upper bound on the multi-photon marginal probabilities $\mathbb{P}(A > 1)$ and $\mathbb{P}(B > 1)$ used in the out-qubit witness of appendix D.9. The bound relies on splitting the incoming mode on a 50/50 beam splitter and detecting twofold coincidences with two non-resolving photodetectors. The argument is presented for Alice's mode, the derivation for Bob's mode is identical by symmetry.

Setup

Consider an arbitrary state ρ on Alice's mode m_A with photon-number distribution $p_n = \langle n | \rho | n \rangle$. The mode is split on a 50/50 beam splitter into two output modes a_1 and a_2 , each connected to a non-resolving photodetector. We compute the probability $\mathbb{P}_{\text{Coinc}}^A$ of a twofold coincidence, that is, both detectors clicking.

The beam splitter maps the input creation operator to $a^\dagger \rightarrow (a_1^\dagger + a_2^\dagger)/\sqrt{2}$, so the n -photon Fock state transforms as

$$|n\rangle \longrightarrow \frac{(a_1^\dagger + a_2^\dagger)^n}{\sqrt{2^n} n!} |0, 0\rangle = \frac{1}{\sqrt{2^n}} \sum_{k=0}^n \binom{n}{k}^{1/2} |k, n-k\rangle, \quad (\text{D.158})$$

where $|k, n-k\rangle$ denotes k photons in a_1 and $n-k$ in a_2 . The diagonal terms of ρ thus transform as

$$|n\rangle \langle n| \longrightarrow \frac{1}{2^n} \sum_{k,l=0}^n \binom{n}{k}^{1/2} \binom{n}{l}^{1/2} |k, n-k\rangle \langle l, n-l|. \quad (\text{D.159})$$

Since we focus on diagonal click probabilities, only the $k = l$ contributions survive after projection on the coincidence POVM derived below, so the off-diagonal terms of ρ in the Fock basis play no role.

Coincidence POVM

A twofold coincidence corresponds to at least one photon in each output mode. The associated POVM element on the two output modes reads

$$\Pi_{\text{Coinc}} = \sum_{k \geq 1, l \geq 1} |k, l\rangle \langle k, l|, \quad (\text{D.160})$$

which excludes the events where one of the two modes is empty.

Coincidence probability

Combining the beam-splitter transformation with the POVM, the coincidence probability reads

$$\mathbb{P}_{\text{Coinc}}^A = \text{tr}(\Pi_{\text{Coinc}} \rho) = \sum_{n \geq 0} p_n \cdot \frac{1}{2^n} \sum_{k=1}^{n-1} \binom{n}{k}, \quad (\text{D.161})$$

where the inner sum runs over $k \in \{1, \dots, n-1\}$ to enforce that both modes receive at least one photon. Using the binomial identity $\sum_{k=0}^n \binom{n}{k} = 2^n$ and isolating the boundary terms $k=0$ and $k=n$, both equal to 1,

$$\sum_{k=1}^{n-1} \binom{n}{k} = 2^n - 2, \quad (\text{D.162})$$

so

$$\mathbb{P}_{\text{Coinc}}^A = \sum_{n \geq 0} p_n \cdot \frac{2^n - 2}{2^n} = \sum_{n \geq 2} p_n \cdot (1 - 2^{1-n}), \quad (\text{D.163})$$

where the sum starts at $n=2$ since the prefactor vanishes for $n \in \{0, 1\}$.

Lower bound and final inequality

The prefactor $1 - 2^{1-n}$ is monotonically increasing in n and equals $1/2$ at $n=2$, hence

$$1 - 2^{1-n} \geq \frac{1}{2} \quad \forall n \geq 2. \quad (\text{D.164})$$

Substituting into (D.163),

$$\mathbb{P}_{\text{Coinc}}^A \geq \frac{1}{2} \sum_{n \geq 2} p_n = \frac{1}{2} \mathbb{P}(A > 1). \quad (\text{D.165})$$

Rearranging gives the desired bound

$$\boxed{\mathbb{P}(A > 1) \leq 2 \mathbb{P}_{\text{Coinc}}^A}, \quad (\text{D.166})$$

and similarly $\mathbb{P}(B > 1) \leq 2 \mathbb{P}_{\text{Coinc}}^B$ for Bob's mode. This inequality provides an experimentally accessible upper bound on the multi-photon marginals, using only twofold coincidence counts on a 50/50 beam splitter and two non-resolving detectors, without requiring photon-number-resolving detection.

Part IV

Appendix - Quantum Circuit Synthesis

D.11 .From Markov Decision Problem to Reinforcement Learning

Markov Decision Problem

Markov Decision Processes take root in a philosophical debate on free will between Markov and Nekrasov [175]. Nekrasov claimed that independence between random variables was a necessary assumption for the law of large numbers. Formally

$$\text{Law of large numbers} \implies \text{Independence} \quad (\text{D.167})$$

Observing that this law was empirically valid in society, he concluded that human actions were independent, supporting the existence of free will. Markov refute his claim with an empirical counterexample. In [176], he built 200 samples of 100 letters from Pushkin's *Eugene Onegin* poem, and classified each letter in two categories: vowel or consonant. With this binary representation, he showed strong correlations between successive letters: a vowel is followed by a consonant 82% of the time, and a consonant by a vowel 39% of the time. Despite these correlations, he showed the proportions of vowels and consonants over 200 samples converged to fixed values, confirming that the law of large numbers extends beyond independent assumption. Beyond the debate, this result paved the way to what is now known as a *Markov chain*: a memoryless stochastic process that iteratively transitions between states. Here this corresponds to a two-state chain where the current letter, vowel or consonant, fully determines the distribution of the next.

Nowadays, *Markov Decision Processes* (MDP), the modern extension of Markov Chain, are widely studied in mathematics and optimization, with a broad scope of applications. In the original book reference [138], Howard introduces MDPs through a simple toy model: a frog jumping on a set of lily pads. At any time, the frog stands on a lily pad, jumps, and lands on a new one, repeating this process indefinitely. This setup captures the two fundamental notions of MDPs: states and transitions. Labelling the pads from 1 to N , the frog is characterized by a state $s \in \{1, \dots, N\}$ corresponding to the pad it occupies. At each jump step t , the system undergoes a transition, leading to a new state s_{t+1} , drawn from a transition probability distribution $\mathbb{P}(s_{t+1}|s_t)$. This measure $\mathbb{P}$ derives from a transition function $\mathcal{P} : \mathcal{S} \rightarrow \Delta(\mathcal{S})$, which sets the game dynamics. Calling t the time step, we index with t the current state, thus we write $s_t \rightarrow s_{t+1}$. Note that the identity action, which leaves the frog on its current pad, is a valid transition, it can be disabled by setting $\mathbb{P}(s_{t+1} = s_t|s_t) = 0$. This toy model also highlights a central feature of MDPs: the memoryless property. The transition distribution depends only on the current state regardless of the past trajectory:

$$\mathbb{P}(s_{t+1}|s_t, s_{t-1}, \dots, s_0) = \mathbb{P}(s_{t+1}|s_t). \quad (\text{D.168})$$

Therefore, at each step, the state alone fully determines the future dynamics of the process.

While capturing the fundamental notions of MDPs, the frog toy model remains purely descriptive. Once the number of pads N , the transition function $\mathcal{P}$, and the initial pad s_0 are set, the frog evolves without any external intervention. However, Howard's original motivation was to develop a computationally feasible model for engineers who needed to act on systems, not only describe them, a goal the simple frog toy model cannot fulfill [138]. The author therefore gradually introduces complexity to the process. First, a reward mechanism is introduced, assigning a reward $r_{ij} \in \mathbb{R}$ to the system for each transition from a state i to a state j . Thus, at every time step t , the system received a reward R_t called *immediate reward*. This reward defines an objective for the process, which is to optimize the earnings over time. Second, an external entity is introduced, the *decision maker*, which is allowed to pick an *action* at each step rather than passively observing the process. From this interactive behavior, an instance of an MDP is sometimes called a *game*. Consequently, the transition distribution is no longer fixed: each action determines a specific transition probability distribution with a specific expected immediate reward. More formally, the transition probability is now expressed as $\mathbb{P}(s_{t+1}|s_t, a_t)$. The objective of the decision maker is to find the optimal sequence of action to maximizes the long-term returns. Finally, the author illustrates this framework with practical cases such as baseball strategies or automobile replacement planning. Building upon this progressive complexification, the community converges to a modern formulation of the MDP [177].

Mathematically, an MDP is defined as a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, R, \gamma \rangle$, whose elements are:

- **State space** $\mathcal{S}$: the set of all possible states.
- **Action space** $\mathcal{A}$: the set of all possible actions.
- **Transition function** $\mathcal{P}$: the function $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{A})$ encoding the environment dynamics. For a state-action pair (s, a) , it returns the probability distribution $\mathbb{P}(\cdot|s, a)$ over the next state, with $\mathbb{P}(s'|s, a)$ denoting the probability of transitioning to s' .
- **Transition probability** $\mathcal{P}$: the transition function, distribution representing the environment dynamics. The probability of transitioning to state s' given that the system is in state s and action a is taken is expressed $\mathbb{P}(s'|s, a)$
- **Reward function** R : the function defines the immediate scalar feedback returned to the agent from a transition. In its most general form, this expected reward depends on the current state s , the chosen action a , and the resulting state s' . Thus, a transition $s \xrightarrow{a} s'$ yields an immediate reward $r = R(s, a, s')$. While capturing the most general framework, simpler forms such as $R(s, a)$ or $R(s')$ are widely used in practice. By providing local signals, the reward function actually defines the global objective of the decision process. Concretely, for each step t , the decision maker's

goal is to find a strategy that maximizes the expected cumulative reward, or gain, denoted $G_{t:\infty}$. For a finite sequence of steps, the gain is defined as

$$G_{t_1:t_2} = \sum_{t=t_1}^{t_2} r_t, \quad (\text{D.169})$$

where r_t is the immediate reward collected at time step t , i.e. $r_t = R(s_{t-1}, a_{t-1}, s_{t+1})$ which is also a standard notation. Calling $\sigma = (s_{t_1}, \dots, s_{t_2})$ the path from s_{t_1} to s_{t_2} , we also denote G_σ the cumulative gain along this path.

- **Discount factor** γ : a real parameter $\gamma \in [0, 1)$ that weights the importance of future rewards. The quantity actually maximized by the decision maker is the expected return

$$G_{t:\infty} = \sum_{k=t}^{\infty} \gamma^k \cdot r_k, \quad (\text{D.170})$$

initially used as a tool to ensure the convergence of infinite games with asymptotic behavior, it also controls how greedy the process is: with $\gamma \rightarrow 0$ favoring immediate rewards and $\gamma \rightarrow 1$ favoring long-term returns.

To build intuition, let us return to the Rubik's cube analogy. To frame the Rubik's cube as an MDP, we must specify the tuple $\langle \mathcal{S}, \mathcal{A}, \mathbb{P}, \mathcal{R}, \gamma \rangle$. As detailed below, these elements can be conceptually divided into two families [114].

Game mechanism: The first family, rigid and inherently defined by the game, corresponds to the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P} \rangle$. The state space $\mathcal{S}$ is the set of all possible cube configurations, thus a state $s \in \mathcal{S}$ corresponds to one cube configuration $C \in \mathcal{C}$. We denote by C_{id} the solved configuration. The action space $\mathcal{A}$ is the set of all possible rotations of the cube. In the case of the Rubik's cube transitions are deterministic: a rotation applied to a given configuration leads to a single configuration. Thus, for every pair (s, a) the resulting state s' is uniquely defined, and the transition probability gives:

$$\forall \tilde{s} \in \mathcal{S}, \quad \mathbb{P}(\tilde{s}|s, a) = \begin{cases} 1 & \text{if } s' = \tilde{s} \\ 0 & \text{else} \end{cases} \quad (\text{D.171})$$

Game objective: The second family is more flexible and defines the goal of the game $\langle R, \gamma \rangle$. While the previous family encodes the rules of the game, the reward function R specifies what the decision maker is asked to achieve. This flexibility comes with a direct impact on the performance, thus the design of R deserves particular care. For instance, one could think of providing intermediate rewards to guide the agent toward the solution, by rewarding action leading to configurations similar colors faces. However, such a reward would mislead the agent: solving paths typically go through disordered looking configurations, thus intermediate reward would penalize these trajectories. An

alternative is to assign a single reward when the cube reaches the solved configuration and zero otherwise:

$$R(s) = \begin{cases} 1 & \text{if } s = C_{\text{ld}} \\ 0 & \text{else} \end{cases} \quad (\text{D.172})$$

With a discount factor $\gamma = 1$, this configuration gives equal weight to all steps and, with a finite budget of T_{max} moves, the cumulative reward reduces to

$$G_{0:\infty} = \sum_{t=1}^{\infty} \gamma^t \cdot r_t = r_{t_{\text{end}}} = \begin{cases} 1 & \text{if } t_{\text{end}} = C_{\text{ld}} \\ 0 & \text{else} \end{cases} \quad (\text{D.173})$$

where $t_{\text{end}} \leq t_{\text{max}}$ denotes the final step, reached either when the cube is solved $t_{\text{end}} < t_{\text{max}}$ or when the move budget is reached $t_{\text{end}} = t_{\text{max}}$.

While providing a formal framework of the process, the MDP formalism does not specify how to solve it. From this purpose, the community introduces two central tools: the policy and the value function.

- **Policy:** this quantity defines probability distribution for the decision maker of picking an action a conditioned of being on the current state s . Formally it is called π and defined such that

$$\begin{aligned} \pi: \mathcal{A} \times \mathcal{S} &\rightarrow [0, 1] \\ (a, s) &\mapsto \pi(a \mid s) \end{aligned} \quad (\text{D.174})$$

This distribution defines the behavior of the decision maker: at each state s an action is drawn from $\pi(\cdot \mid s)$. Thus, the policy can be seen as the intelligence of the decision maker, driving the path through the decision process

- **State Value function:** quantifies the expected cumulative reward obtained from a state s when following the policy π . Formally, at a time step t , it reads

$$V^{\pi}(s) = \mathbb{E}_{\pi} \left[\sum_{k=0}^{\infty} \gamma^k \cdot r_{k+1} \mid s_0 = s \right]. \quad (\text{D.175})$$

- **Action Value function:** quantifies the expected cumulative reward obtained while picking an action a from a state s when following the policy π . Formally, at a time step t , it reads

$$Q^{\pi}(s, a) = \mathbb{E}_{\pi} \left[\sum_{k=0}^{\infty} \gamma^k \cdot r_{k+1} \mid s_0 = s, a_0 = a \right]. \quad (\text{D.176})$$

Note that because the Markov property ensures that the value function depends only on the current state and not on the time step, we conventionally set $s_0 = s$. Since the goal of the decision maker is to find a sequence of actions that maximizes the cumulative reward, solving MDP actually means finding a policy π^* , called the *optimal policy*, that maximizes

the value function over all states. Defining the partial order relation

$$\pi' \leq \pi \iff \forall s \in \mathcal{S}, V_{\pi'}(s) \leq V_{\pi}(s), \quad (\text{D.177})$$

an optimal policy π^* is defined as an upper bound:

$$\forall \pi, \quad \pi \leq \pi^*. \quad (\text{D.178})$$

The existence of such a policy was demonstrated in [178].

Seeking for MDPs resolution, a pioneering contribution lies with Bellman work in [178] where he developed dynamic programming. The core idea is to solve the Bellman equations, whose treatment lies beyond the scope of this thesis. While providing convergence to an optimal policy π^* with strong stability, this algorithm relies on strong prior knowledge of the process. Concretely, dynamic programming requires an explicit access to the transition function $\mathcal{P}$ and to the reward function R , which may not be available in practice. Built on iterative estimations of the value function, the algorithm also requires to span the entire state space $\mathcal{S}$, which becomes impossible for large spaces such as quantum circuit synthesis. The community has developed a wide range of algorithms, each adapted to specific assumptions on available information on the MDP. Among these, *Reinforcement Learning* offers an attractive framework: based on state space exploration, it solves MDP without prior knowledge of the game dynamics, by progressively learning the optimal policy through environment interaction.

The reinforcement learning framework

Single-agent reinforcement learning, abbreviated RL, is an area of machine learning that solves MDPs through interaction. Framed as a Markov Decision Process (MDP), the system is described by a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, R, \gamma \rangle$ called *environment*. The decision maker, called *agent*, aims to learn an optimal policy π^* that maximizes the expected cumulative reward G . Unlike dynamic programming, an RL agent has no prior access to the reward function $R(s, a, s')$, the transition function $\mathcal{P}$, nor the state space $\mathcal{S}$. Thus, the agent must explore the environment through iterative interaction, discovering the dynamic of the game, collecting reward feedback, and progressively converging toward an optimal policy π^* .

Concretely, the RL framework can be described as an iterative loop. For a time step t , the environment is in a state $s_t \in \mathcal{S}$, with a set of available actions $\mathcal{A}_{s_t} \subseteq \mathcal{A}$ accessible to the agent. From via environment observation, the agent picks an action $a_t \in \mathcal{A}_{s_t}$ according to its current policy π_t , the action leads the environment to a new state s_{t+1} drawn from the transition probability $\mathbb{P}(s_{t+1}|s_t, a_t)$. This transition returns a reward r_t to the agent, following the reward function $R(s_{t-1}, a_{t-1}, s_t)$. The agent uses this feedback to update its policy $\pi_t \rightarrow \pi_{t+1}$, refining its strategy toward the optimal policy π^* . The process is then repeated until either a terminal state s_{end} is reached or a maximum number of

steps $t_{\max}$ is exceeded. Within this loop, the agent accumulates knowledge on the game, progressively refining its policy. This mechanism is illustrated in appendix D.11.

In a more geometric way, RL can be framed as a path-finding problem in a *decision tree*. Each node stores the state history v_h with $h = (s_0, \dots, s_t)$, and each branch represents the labeled transition $s_t \xrightarrow{a_t} s_{t+1}$ weighted by the transition probability $\mathbb{P}(s_{t+1}|s_t, a_t)$. Each transition yields an immediate reward r_t , and a cumulative reward $G_\sigma = \sum_{t=0}^{t_{\text{end}}} r_t$ from a path $\sigma = (s_0, \dots, s_{t_{\text{end}}})$ is associated with each leaf. In this representation, the problem reduces to finding the path leading to the most promising leaf. Note that, under the memoryless Markov assumption, the tree naturally folds into a graph by collapsing the history $(s_0, \dots, s_t)$ to its endpoint s_t for each node. This decision tree representation is at the heart of AlphaZero and will be studied in detail in the following sections.

The previous paragraph introduced the RL framework as a bare agent-environment interaction mechanism, leaving the resolution part to the agent. Indeed, while the environment tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, R, \gamma \rangle$ sets the game mechanism, rules, and objective, environment itself does not bring any specific algorithmic strategy. The resolution lies in the agent's learning policy $\pi \rightarrow \pi^*$: the agent collects knowledge of the environment by iteratively interacting with it, and uses a learning algorithm to refine its policy. Several learning algorithms are available, to name a few: Q-learning [179], REINFORCE [180], or PPO [181]. The choice of this learning algorithm depends on the characteristics of the environment, such as the state space, the transition mechanism, or the available information. Thus, for the same game $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, R, \gamma \rangle$, one can train different agents with different learning capabilities. This separation draws a parallel with the structure of dynamic programming: the framework reduces the problem to solving the Bellman equations, while the choice of method to solve them constitutes the resolution strategy.

Finally, let's add two brief clarifications. First, the previous paragraph implicitly introduced the notion of observation, which deserves a brief clarification. The RL framework actually fits within a broader formalism called Partially Observable Markov Decision Process [182]. While the MDP formalism assumes direct access of the decision maker to the system current state s , an RL agent only receives an *observation* $o \in \mathcal{O}$, with $\mathcal{O}$ the observation space. This observation may be a partial view of the environment state s , which the agent must interpret. It is actually a practical common case, for instance a self-driven car agent has access to a modelization of car's environment from sensors. In the scope of this thesis, the observation coincides with the state $o_t = s_t$, and we use the two terms interchangeably. Secondly, this thesis assumes deterministic transitions, meaning a state-action pair (s_t, a_t) leads to a unique state s_{t+1} . Formally,

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad \exists! s' \in \mathcal{S} / \mathbb{P}(s' | s, a) = 1. \quad (\text{D.179})$$




Figure D.1: Minimalist illustration of the reinforcement learning interaction mechanics. At each time step t , the agent picks an action a_t from $\mathcal{A}$ according to the policy $\pi(\cdot | s_t)$. This action drives the environment transition $s_t \xrightarrow{a_t} s_{t+1}$, producing a reward feedback $R(s_t, a_t, s_{t+1}) = r_t$. The interaction loop stops when the environment reaches the target state s_{targ} or the maximum step count t_{max} .

Bibliography

- [1] 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**, 1895 (1993).
- [2] A. K. Ekert, "Quantum cryptography based on Bell's theorem", *Physical Review Letters* **67**, 661 (1991).
- [3] R. Renner and R. Wolf, *The debate over QKD: A rebuttal to the NSA's objections*, (July 27, 2023) <http://arxiv.org/abs/2307.15116>, pre-published.
- [4] P. W. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer", *SIAM Journal on Computing* **26**, 1484 (1997).
- [5] C. Gidney and M. Ekerå, "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits", *Quantum* **5**, 433 (2021).
- [6] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, "Fundamental limits of repeaterless quantum communications", *Nature Communications* **8**, 15043 (2017).
- [7] B. Yurke and D. Stoler, "Bell's-inequality experiments using independent-particle sources", *Physical Review A* **46**, 2229 (1992).
- [8] M. Żukowski, A. Zeilinger, M. A. Horne, and A. K. Ekert, "Event-ready-detectors" Bell experiment via entanglement swapping", *Physical Review Letters* **71**, 4287 (1993).
- [9] S. Muralidharan, L. Li, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, "Optimal architectures for long distance quantum communication", *Scientific Reports* **6**, 20463 (2016).
- [10] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, "Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication", *Physical Review Letters* **81**, 5932 (1998).
- [11] W. Dür, H.-J. Briegel, J. I. Cirac, and P. Zoller, "Quantum repeaters based on entanglement purification", *Physical Review A* **59**, 169 (1999).
- [12] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, "Long-distance quantum communication with atomic ensembles and linear optics", *Nature* **414**, 413 (2001).

- [13] N. Sangouard, C. Simon, H. de Riedmatten, and N. Gisin, "Quantum repeaters based on atomic ensembles and linear optics", *Reviews of Modern Physics* **83**, 33 (2011).
- [14] L. Jiang, J. M. Taylor, K. Nemoto, W. J. Munro, R. Van Meter, and M. D. Lukin, "Quantum repeater with encoding", *Physical Review A* **79**, 032325 (2009).
- [15] W. J. Munro, K. A. Harrison, A. M. Stephens, S. J. Devitt, and K. Nemoto, "From quantum multiplexing to high-performance quantum networking", *Nature Photonics* **4**, 792 (2010).
- [16] E. Knill and R. Laflamme, *Concatenated Quantum Codes*, (Aug. 8, 1996) <http://arxiv.org/abs/quant-ph/9608012>, pre-published.
- [17] A. G. Fowler, D. S. Wang, C. D. Hill, T. D. Ladd, R. Van Meter, and L. C. L. Hollenberg, "Surface Code Quantum Communication", *Physical Review Letters* **104**, 180503 (2010).
- [18] W. J. Munro, A. M. Stephens, S. J. Devitt, K. A. Harrison, and K. Nemoto, "Quantum communication without the necessity of quantum memories", *Nature Photonics* **6**, 777 (2012).
- [19] S. Muralidharan, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, "Ultra-fast and Fault-Tolerant Quantum Communication across Long Distances", *Physical Review Letters* **112**, 250501 (2014).
- [20] K. Azuma, S. E. Economou, D. Elkouss, P. Hilaire, L. Jiang, H.-K. Lo, and I. Tzitrin, "Quantum repeaters: From quantum networks to the quantum internet", *Reviews of Modern Physics* **95**, 045006 (2023).
- [21] M. Fleischhauer and M. D. Lukin, "Dark-State Polaritons in Electromagnetically Induced Transparency", *Physical Review Letters* **84**, 5094 (2000).
- [22] B. Kraus, W. Tittel, N. Gisin, M. Nilsson, S. Kröll, and J. I. Cirac, "Quantum memory for nonstationary light fields based on controlled reversible inhomogeneous broadening", *Physical Review A* **73**, 020302 (2006).
- [23] C. Simon, H. de Riedmatten, M. Afzelius, N. Sangouard, H. Zbinden, and N. Gisin, "Quantum Repeater with Photon Pair Sources and Multimode Memories", *Physical Review Letters* **98**, 190503 (2007).
- [24] M. Afzelius, C. Simon, H. de Riedmatten, and N. Gisin, "Multimode quantum memory based on atomic frequency combs", *Physical Review A* **79**, 052329 (2009).
- [25] V. Damon, M. Bonarota, A. Louchet-Chauvet, T. Chanelière, and J.-L. Le Gouët, "Revival of silenced echo and quantum memory for light", *New Journal of Physics* **13**, 093031 (2011).

- [26] C. W. Chou, H. de Riedmatten, D. Felinto, S. V. Polyakov, S. J. van Enk, and H. J. Kimble, "Measurement-induced entanglement for excitation stored in remote atomic ensembles", *Nature* **438**, 828 (2005).
- [27] C.-W. Chou, J. Laurat, H. Deng, K. S. Choi, H. de Riedmatten, D. Felinto, and H. J. Kimble, "Functional quantum nodes for entanglement distribution over scalable quantum networks", *Science* **316**, 1316 (2007).
- [28] J.-L. Liu, X.-Y. Luo, Y. Yu, C.-Y. Wang, B. Wang, Y. Hu, J. Li, M.-Y. Zheng, B. Yao, Z. Yan, D. Teng, J.-W. Jiang, X.-B. Liu, X.-P. Xie, J. Zhang, Q.-H. Mao, X. Jiang, Q. Zhang, X.-H. Bao, and J.-W. Pan, "Creation of memory-memory entanglement in a metropolitan quantum network", *Nature* **629**, 579 (2024).
- [29] Y. Yu, F. Ma, X.-Y. Luo, B. Jing, P.-F. Sun, R.-Z. Fang, C.-W. Yang, H. Liu, M.-Y. Zheng, X.-P. Xie, W.-J. Zhang, L.-X. You, Z. Wang, T.-Y. Chen, Q. Zhang, X.-H. Bao, and J.-W. Pan, "Entanglement of two quantum memories via fibres over dozens of kilometres", *Nature* **578**, 240 (2020).
- [30] M. Businger, L. Nicolas, T. S. Mejia, A. Ferrier, P. Goldner, and M. Afzelius, "Non-classical correlations over 1250 modes between telecom photons and 979-nm photons stored in 171Yb3+:Y2SiO5", *Nature Communications* **13**, 6438 (2022).
- [31] M. Gündoğan, P. M. Ledingham, K. Kutluer, M. Mazzerà, and H. de Riedmatten, "Solid State Spin-Wave Quantum Memory for Time-Bin Qubits", *Physical Review Letters* **114**, 230501 (2015).
- [32] A. Seri, A. Lenhard, D. Rieländer, M. Gündoğan, P. M. Ledingham, M. Mazzerà, and H. de Riedmatten, "Quantum Correlations between Single Telecom Photons and a Multimode On-Demand Solid-State Quantum Memory", *Physical Review X* **7**, 021028 (2017).
- [33] L. Feldmann, S. Wengerowsky, A. Das, S. Duranti, J. Hänni, S. Grandi, and H. de Riedmatten, "Cavity-Enhanced Spin-Wave Solid-State Quantum Memory", *Physical Review Letters* **135**, 120801 (2025).
- [34] D. Lago-Rivera, S. Grandi, J. V. Rakonjac, A. Seri, and H. de Riedmatten, "Telecom-heralded entanglement between multimode solid-state quantum memories", *Nature* **594**, 37 (2021).
- [35] J. Hänni, A. E. Rodríguez-Moldes, F. Appas, S. Wengerowsky, D. Lago-Rivera, M. Teller, S. Grandi, and H. de Riedmatten, "Heralded Entanglement of On-Demand Spin-Wave Solid-State Quantum Memories for Multiplexed Quantum Network Links", *Physical Review X* **15**, 041003 (2025).
- [36] N. Sangouard, R. Dubessy, and C. Simon, "Quantum repeaters based on single trapped ions", *Physical Review A* **79**, 042340 (2009).

- [37] N. Sangouard, C. Simon, J. Minář, H. Zbinden, H. de Riedmatten, and N. Gisin, “Long-distance entanglement distribution with single-photon sources”, *Physical Review A* **76**, 050301 (2007).
- [38] H. J. Kimble, “The quantum internet”, *Nature* **453**, 1023 (2008).
- [39] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, “Heralded entanglement between solid-state qubits separated by three metres”, *Nature* **497**, 86 (2013).
- [40] A. J. Stolk, K. L. van der Enden, M.-C. Slater, I. te Raa-Derckx, P. Botma, J. van Rantwijk, J. J. B. Biemond, R. A. J. Hagen, R. W. Herfst, W. D. Koek, A. J. H. Meskers, R. Vollmer, E. J. van Zwet, M. Markham, A. M. Edmonds, J. F. Geus, F. Elsen, B. Jungbluth, C. Haefner, C. Tresp, J. Stuhler, S. Ritter, and R. Hanson, “Metropolitan-scale heralded entanglement of solid-state qubits”, *Science Advances* **10**, eadp6442 (2024).
- [41] C. M. Knaut, A. Suleymanzade, Y.-C. Wei, D. R. Assumpcao, P.-J. Stas, Y. Q. Huan, B. Machielse, E. N. Knall, M. Sutula, G. Baranes, N. Sinclair, C. De-Eknamkul, D. S. Levonian, M. K. Bhaskar, H. Park, M. Lončar, and M. D. Lukin, “Entanglement of nanophotonic quantum memory nodes in a telecom network”, *Nature* **629**, 573 (2024).
- [42] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, “Entanglement of single-atom quantum bits at a distance”, *Nature* **449**, 68 (2007).
- [43] V. Krutyanskiy, M. Meraner, J. Schupp, V. Krcmarsky, H. Hainzer, and B. P. Lanyon, “Light-matter entanglement over 50 km of optical fibre”, *npj Quantum Information* **5**, 72 (2019).
- [44] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, “High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network”, *Physical Review Letters* **124**, 110501 (2020).
- [45] V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D. A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B. P. Lanyon, and T. E. Northup, “Entanglement of Trapped-Ion Qubits Separated by 230 Meters”, *Physical Review Letters* **130**, 050803 (2023).
- [46] T. van Leent, M. Bock, F. Fertig, R. Garthoff, S. Eppelt, Y. Zhou, P. Malik, M. Seubert, T. Bauer, W. Rosenfeld, W. Zhang, C. Becher, and H. Weinfurter, “Entangling single atoms over 33 km telecom fibre”, *Nature* **607**, 69 (2022).

- [47] B.-W. Lu, C.-W. Yang, R.-Q. Wang, B.-F. Gao, Y.-Z. Zhen, Z.-G. Wang, J.-K. Shi, Z.-Q. Ren, T. A. Hahn, E. Y.-Z. Tan, X.-P. Xie, M.-Y. Zheng, X. Jiang, J. Zhang, F. Xu, Q. Zhang, X.-H. Bao, and J.-W. Pan, "Device-independent quantum key distribution over 100 km with single atoms", *Science* **391**, 592 (2026).
- [48] S. Wehner, D. Elkouss, and R. Hanson, "Quantum internet: A vision for the road ahead", *Science* **362**, eaam9288 (2018).
- [49] H. Jacinto, É. Gouzien, and N. Sangouard, "Network requirements for distributed quantum computation", *Physical Review Research* **8**, 013205 (2026).
- [50] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, "A quantum network of clocks", *Nature Physics* **10**, 582 (2014).
- [51] Z. Zhang and Q. Zhuang, "Distributed quantum sensing", *Quantum Science and Technology* **6**, 043001 (2021).
- [52] D. Gottesman, T. Jennewein, and S. Croke, "Longer-Baseline Telescopes Using Quantum Repeaters", *Physical Review Letters* **109**, 070503 (2012).
- [53] *Quantum Internet Alliance | Building a Global Quantum Internet Made in Europe*, Quantum Internet Alliance, (2023) <https://quantuminternetalliance.org/>.
- [54] R. Baumann, "Radiation-induced soft errors in advanced semiconductor technologies", *IEEE Transactions on Device and Materials Reliability* **5**, 305 (2005).
- [55] G. Dong, N. Xie, and T. Zhang, "On the Use of Soft-Decision Error-Correction Codes in nand Flash Memory", *IEEE Transactions on Circuits and Systems I: Regular Papers* **58**, 429 (2011).
- [56] C. E. Shannon, "A mathematical theory of communication", *The Bell System Technical Journal* **27**, 379 (1948).
- [57] R. W. Hamming, "Error detecting and error correcting codes", *The Bell System Technical Journal* **29**, 147 (1950).
- [58] Y. Cai, E. F. Haratsch, O. Mutlu, and K. Mai, "Error patterns in MLC NAND flash memory: Measurement, characterization, and analysis", in *2012 Design, Automation & Test in Europe Conference & Exhibition (DATE)* (Mar. 2012), pp. 521–526.
- [59] A. Berthiaume, D. Deutsch, and R. Jozsa, "The stabilisation of quantum computations", in *Proceedings Workshop on Physics and Computation. PhysComp '94* (Nov. 1994), pp. 60–62.
- [60] P. W. Shor, "Scheme for reducing decoherence in quantum computer memory", *Physical Review A* **52**, R2493 (1995).

- [61] A. M. Steane, "Error Correcting Codes in Quantum Theory", *Physical Review Letters* **77**, 793 (1996).
- [62] A. R. Calderbank and P. W. Shor, "Good quantum error-correcting codes exist", *Physical Review A* **54**, 1098 (1996).
- [63] P. W. Shor, *Fault-tolerant quantum computation*, (Mar. 5, 1997) <http://arxiv.org/abs/quant-ph/9605011>, pre-published.
- [64] S. B. Bravyi and A. Y. Kitaev, *Quantum codes on a lattice with boundary*, (Nov. 20, 1998) <https://arxiv.org/abs/quant-ph/9811052v1>, pre-published.
- [65] A. Yu. Kitaev, "Fault-tolerant quantum computation by anyons", *Annals of Physics* **303**, 2 (2003).
- [66] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, "Topological quantum memory", *Journal of Mathematical Physics* **43**, 4452 (2002).
- [67] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation", *Physical Review A* **86**, 032324 (2012).
- [68] D. Litinski, "A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery", *Quantum* **3**, 128 (2019).
- [69] A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta, *Open Quantum Assembly Language*, (July 13, 2017) <http://arxiv.org/abs/1707.03429>, pre-published.
- [70] R. S. Smith, M. J. Curtis, and W. J. Zeng, *A Practical Quantum Instruction Set Architecture*, (Feb. 17, 2017) <http://arxiv.org/abs/1608.03355>, pre-published.
- [71] M. E. Beverland, P. Murali, M. Troyer, K. M. Svore, T. Hoefler, V. Kliuchnikov, G. H. Low, M. Soeken, A. Sundaram, and A. Vashillo, *Assessing requirements to scale to practical quantum advantage*, (Nov. 14, 2022) <http://arxiv.org/abs/2211.07629>, pre-published.
- [72] D. Gottesman, "Stabilizer Codes and Quantum Error Correction" (arXiv, May 28, 1997).
- [73] D. Gottesman, "Theory of fault-tolerant quantum computation", *Physical Review A* **57**, 127 (1998).
- [74] P. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan, "On universal and fault-tolerant quantum computing: a novel basis and a new constructive proof of universality for Shor's basis", in *40th Annual Symposium on Foundations of Computer Science (Cat. No.99CB37039)* (Oct. 1999), pp. 486–494.

- [75] B. Eastin and E. Knill, “Restrictions on Transversal Encoded Quantum Gate Sets”, *Physical Review Letters* **102**, 110502 (2009).
- [76] D. Gottesman and I. L. Chuang, “Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations”, *Nature* **402**, 390 (1999).
- [77] S. Bravyi and A. Kitaev, “Universal quantum computation with ideal Clifford gates and noisy ancillas”, *Physical Review A* **71**, 022316 (2005).
- [78] C. Gidney, N. Shutty, and C. Jones, *Magic state cultivation: growing T states as cheap as CNOT gates*, (Sept. 26, 2024) <http://arxiv.org/abs/2409.17595>, pre-published.
- [79] E. Rosenfeld, C. Gidney, G. Roberts, A. Morvan, N. Lacroix, D. Kafri, J. Marshall, M. Li, V. Sivak, D. Abanin, A. Abbas, R. Acharya, L. A. Beni, G. Aigeldinger, R. Alcaraz, S. Alcaraz, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, W. Askew, N. Astrakhantsev, J. Atalaya, R. Babbush, B. Ballard, J. C. Bardin, H. Bates, A. Bengtsson, M. B. Karimi, A. Bilmes, S. Bilodeau, F. Borjans, J. Bovaird, D. Bowers, L. Brill, P. Brooks, M. Broughton, D. A. Browne, B. Buchea, B. B. Buckley, T. Burger, B. Burkett, N. Bushnell, J. Busnaina, A. Cabrera, J. Campero, H.-S. Chang, S. Chen, Z. Chen, B. Chiaro, L.-Y. Chih, A. Y. Cleland, B. Cochrane, M. Cockrell, J. Cogan, P. Conner, H. Cook, R. G. Cortiñas, W. Courtney, A. L. Crook, B. Curtin, M. Damyanov, S. Das, D. M. Debroy, S. Demura, P. Donohoe, I. Drozdov, A. Dunsworth, V. Ehimhen, A. Eickbusch, A. M. Elbag, L. Ella, M. Elzouka, D. Enriquez, C. Erickson, L. Faoro, V. S. Ferreira, M. Flores, L. F. Burgos, S. Fontes, E. Forati, J. Ford, B. Foxen, M. Fukami, A. W. L. Fung, L. Fuste, S. Ganjam, G. Garcia, C. Garrick, R. Gasca, H. Gehring, R. Geiger, É. Genois, W. Giang, D. Gilboa, J. E. Goeters, E. C. Gonzales, R. Gosula, S. J. de Graaf, A. G. Dau, D. Graumann, J. Grebel, A. Greene, J. A. Gross, J. Guerrero, L. L. Guevel, T. Ha, S. Habegger, T. Hadick, A. Hadjikhani, M. C. Hamilton, M. Hansen, M. P. Harrigan, S. D. Harrington, J. Hartshorn, S. Heslin, P. Heu, O. Higgott, R. Hiltermann, J. Hilton, H.-Y. Huang, M. Hucka, C. Hudspeth, A. Huff, W. J. Huggins, L. B. Ioffe, E. Jeffrey, S. Jevons, Z. Jiang, X. Jin, C. Joshi, P. Juhas, A. Kabel, H. Kang, K. Kang, A. H. Karamlou, R. Kaufman, K. Kechedzhi, T. Khattar, M. Khezri, S. Kim, P. V. Klimov, C. M. Knaut, B. Kobrin, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, R. Kudo, B. Kueffler, A. Kumar, V. D. Kurilovich, V. Kutsko, T. Lange-Dei, B. W. Langley, P. Laptev, K.-M. Lau, E. Leavell, J. Ledford, J. Lee, K. Lee, B. J. Lester, W. Leung, L. Li, W. Y. Li, A. T. Lill, W. P. Livingston, M. T. Lloyd, A. Locharla, L. D. Lorenzo, E. Lucero, D. Lundahl, A. Lunt, S. Madhuk, A. Maiti, A. Maloney, S. Mandrà, L. S. Martin, O. Martin, E. Mascot, P. M. Das, D. Maslov, M. Mathews, C. Maxfield,

- J. R. McClean, M. McEwen, S. Meeks, A. Megrant, K. C. Miao, Z. K. Mineev, R. Molavi, S. Molina, S. Montazeri, C. Neill, M. Newman, A. Nguyen, M. Nguyen, C.-H. Ni, M. Y. Niu, N. Noll, L. Oas, W. D. Oliver, R. Orosco, K. Ottosson, A. Pagano, A. D. Paolo, S. Peek, D. Peterson, A. Pizzuto, E. Portoles, R. Potter, O. Pritchard, M. Qian, C. Quintana, G. Ramachandran, A. Ranadive, M. J. Reagor, R. Resnick, D. M. Rhodes, D. Riley, R. Rodriguez, E. Ropes, L. B. D. Rose, E. Rosenberg, D. Rosenstock, E. Rossi, P. Roushan, D. A. Rower, R. Salazar, K. Sankaragomathi, M. C. Sarihan, M. Schaefer, S. Schroeder, H. F. Schurkus, A. Shahingohar, M. J. Shearn, A. Shorter, N. Shutty, V. Shvarts, S. Small, W. C. Smith, D. A. Sobel, B. Spells, S. Springer, G. Sterling, J. Suchard, A. Szasz, A. Szein, M. Taylor, J. P. Thiruraman, D. Thor, D. Timucin, E. Tomita, A. Torres, M. M. Torunbalci, H. Tran, A. Vaishnav, J. Vargas, S. Vdovichev, G. Vidal, B. Villalonga, C. V. Heidweiller, M. Voorhees, S. Waltman, J. Waltz, S. X. Wang, D. Wang, B. Ware, J. D. Watson, Y. Wei, T. Weidel, T. White, K. Wong, B. W. K. Woo, C. J. Wood, M. Woodson, C. Xing, Z. J. Yao, P. Yeh, B. Ying, J. Yoo, N. Yosri, E. Young, G. Young, A. Zalcman, R. Zhang, Y. Zhang, N. Zhu, N. Zobrist, Z. Zou, H. Neven, S. Boixo, C. Jones, J. Kelly, A. Bourassa, and K. J. Satzinger, *Magic state cultivation on a superconducting quantum processor*, (Dec. 15, 2025) <http://arxiv.org/abs/2512.13908>, pre-published.
- [80] A. J. Landahl and C. Cesare, *Complex instruction set computing architecture for performing accurate quantum Z rotations with less magic*, (Oct. 15, 2013) <http://arxiv.org/abs/1302.3240>, pre-published.
- [81] E. T. Campbell, B. M. Terhal, and C. Vuillot, "Roads towards fault-tolerant universal quantum computation", *Nature* **549**, 172 (2017).
- [82] F. T. Chong, D. Franklin, and M. Martonosi, "Programming languages and compiler design for realistic quantum hardware", *Nature* **549**, 180 (2017).
- [83] D. Kremer, A. Javadi-Abhari, and P. Mukhopadhyay, *Optimizing the non-Clifford-count in unitary synthesis using Reinforcement Learning*, (Dec. 11, 2025) <http://arxiv.org/abs/2509.21709>, pre-published.
- [84] H. Nagarajan, O. Lockwood, and C. Coffrin, *QuantumCircuitOpt: An Open-source Framework for Provably Optimal Quantum Circuit Design*, (Nov. 23, 2021) <http://arxiv.org/abs/2111.11674>, pre-published.
- [85] É. Gouzien and N. Sangouard, *Provably optimal exact gate synthesis from a discrete gate set*, (Mar. 19, 2025) <http://arxiv.org/abs/2503.15452>, pre-published.
- [86] É. Gouzien and N. Sangouard, "Factoring 2048-bit RSA Integers in 177 Days with 13 436 Qubits and a Multimode Memory", *Physical Review Letters* **127**, 140503 (2021).

- [87] M. Remaud and V. Vandaele, *Ancilla-free Quantum Adder with Sublinear Depth*, (2025) <http://arxiv.org/abs/2501.16802>, pre-published.
- [88] M. Remaud, *Quantum adders: on the structural link between the ripple-carry and carry-lookahead techniques*, (Oct. 1, 2025) <http://arxiv.org/abs/2510.00840>, pre-published.
- [89] A. Y. Kitaev, "Quantum computations: algorithms and error correction", *Russian Mathematical Surveys* **52**, 1191 (1997).
- [90] C. M. Dawson and M. A. Nielsen, *The Solovay-Kitaev algorithm*, (Aug. 23, 2005) <http://arxiv.org/abs/quant-ph/0505030>, pre-published.
- [91] M. Amy, D. Maslov, M. Mosca, and M. Roetteler, "A Meet-in-the-Middle Algorithm for Fast Synthesis of Depth-Optimal Quantum Circuits", *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **32**, 818 (2013).
- [92] H. Nagarajan, *Harshangrjn/QuantumCircuitOpt.jl*, Jan. 5, 2026.
- [93] M. Mosca and P. Mukhopadhyay, "A polynomial time and space heuristic algorithm for T-count", *Quantum Science and Technology* **7**, 015003 (2021).
- [94] V. Gheorghiu, M. Mosca, and P. Mukhopadhyay, "A (quasi-)polynomial time heuristic algorithm for synthesizing T-depth optimal circuits", *npj Quantum Information* **8**, 110 (2022).
- [95] C. Lancu, M. Davis, and E. Smith, *Quantum Search Compiler (Qsearch) v2.0*, Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA (United States), 2020.
- [96] E. Younis, K. Sen, K. Yelick, and C. Lancu, *QFAST: Conflating Search and Numerical Optimization for Scalable Quantum Circuit Synthesis*, (Mar. 12, 2021) <http://arxiv.org/abs/2103.07093>, pre-published.
- [97] E. Smith, M. G. Davis, J. Larson, E. Younis, L. B. Otfelie, W. Lavrijsen, and C. Lancu, "LEAP: Scaling Numerical Optimization Based Synthesis Using an Incremental Approach", *ACM Transactions on Quantum Computing* **4**, 5:1 (2023).
- [98] A. Paradis, J. Dekoninck, B. Bichsel, and M. Vechev, "Synthetiq: Fast and Versatile Quantum Circuit Synthesis", *Reproduction Package for the Article "Synthetiq: Fast and Versatile Quantum Circuit Synthesis"* **8**, 96:55 (2024).
- [99] *Eth-sri/synthetiq*, SRI Lab, ETH Zurich, Nov. 19, 2025.
- [100] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, "Machine learning and the physical sciences", *Reviews of Modern Physics* **91**, 045002 (2019).

- [101] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning”, *Nature* **521**, 436 (2015).
- [102] A. Fawzi, M. Balog, A. Huang, T. Hubert, B. Romera-Paredes, M. Barekatain, A. Novikov, F. J. R. Ruiz, J. Schrittwieser, G. Swirszcz, D. Silver, D. Hassabis, and P. Kohli, “Discovering faster matrix multiplication algorithms with reinforcement learning”, *Nature* **610**, 47 (2022).
- [103] J. Abramson, J. Adler, J. Dunger, R. Evans, T. Green, A. Pritzel, O. Ronneberger, L. Willmore, A. J. Ballard, J. Bambrick, S. W. Bodenstein, D. A. Evans, C.-C. Hung, M. O’Neill, D. Reiman, K. Tunyasuvunakool, Z. Wu, A. Žemgulytė, E. Arvaniti, C. Beattie, O. Bertolli, A. Bridgland, A. Cherepanov, M. Congreve, A. I. Cowen-Rivers, A. Cowie, M. Figurnov, F. B. Fuchs, H. Gladman, R. Jain, Y. A. Khan, C. M. R. Low, K. Perlin, A. Potapenko, P. Savy, S. Singh, A. Stecula, A. Thillaisundaram, C. Tong, S. Yakneen, E. D. Zhong, M. Zielinski, A. Židek, V. Bapst, P. Kohli, M. Jaderberg, D. Hassabis, and J. M. Jumper, “Accurate structure prediction of biomolecular interactions with AlphaFold 3”, *Nature* **630**, 493 (2024).
- [104] A. Merchant, S. Batzner, S. S. Schoenholz, M. Aykol, G. Cheon, and E. D. Cubuk, “Scaling deep learning for materials discovery”, *Nature* **624**, 80 (2023).
- [105] A. Goldie, A. Mirhoseini, M. Yazgan, J. W. Jiang, E. Songhori, S. Wang, Y.-J. Lee, E. Johnson, O. Pathak, A. Nova, J. Pak, A. Tong, K. Srinivasa, W. Hang, E. Tuncer, Q. V. Le, J. Laudon, R. Ho, R. Carpenter, and J. Dean, “Addendum: A graph placement methodology for fast chip design”, *Nature* **634**, E10 (2024).
- [106] M. Krenn, J. Landgraf, T. Foesel, and F. Marquardt, “Artificial intelligence and machine learning for quantum technologies”, *Physical Review A* **107**, 010101 (2023).
- [107] J. Bausch, A. W. Senior, F. J. H. Heras, T. Edlich, A. Davies, M. Newman, C. Jones, K. Satzinger, M. Y. Niu, S. Blackwell, G. Holland, D. Kafri, J. Atalaya, C. Gidney, D. Hassabis, S. Boixo, H. Neven, and P. Kohli, “Learning high-accuracy error decoding for quantum processors”, *Nature* **635**, 834 (2024).
- [108] J. Olle, R. Zen, M. Puviani, and F. Marquardt, “Simultaneous discovery of quantum error correction codes and encoders with a noise-aware reinforcement learning agent”, *npj Quantum Information* **10**, 126 (2024).
- [109] G. Acampora, A. Ambainis, N. Ares, L. Banchi, P. Bhardwaj, D. Binosi, G. A. D. Briggs, T. Calarco, V. Dunjko, J. Eisert, O. Ezratty, P. Erker, F. Fedele, E. Gil-Fuster, M. Gärttner, M. Granath, M. Heyl, I. Kerenidis, M. Klusch, A. F. Kockum, R. Kueng, M. Krenn, J. Lässig, A. Macaluso, S. Maniscalco, F. Marquardt, K. Michielsen, G. Muñoz-Gil, D. Müssig, H. P. Nautrup, S. A.

- Neubauer, E. van Nieuwenburg, R. Orus, J. Schmiedmayer, M. Schmitt, P. Slusallek, F. Vicentini, C. Weitenberg, and F. K. Wilhelm, *Quantum computing and artificial intelligence: status and perspectives*, (June 30, 2025) <http://arxiv.org/abs/2505.23860>, pre-published.
- [110] L. Moro, M. G. A. Paris, M. Restelli, and E. Prati, “Quantum compiling by deep reinforcement learning”, *Communications Physics* **4**, 178 (2021).
 - [111] S. Rietsch, A. Y. Dubey, C. Ufrecht, M. Periyasamy, A. Plinge, C. Mutschler, and D. D. Scherer, “Unitary Synthesis of Clifford+T Circuits with Reinforcement Learning”, in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, Vol. 01 (Sept. 2024), pp. 824–835.
 - [112] F. Furrutter, G. Muñoz-Gil, and H. J. Briegel, “Quantum circuit synthesis with diffusion models”, *Nature Machine Intelligence* **6**, 515 (2024).
 - [113] R. Zen, J. Olle, L. Colmenarez, M. Puviani, M. Müller, and F. Marquardt, “Quantum Circuit Discovery for Fault-Tolerant Logical State Preparation with Reinforcement Learning”, *Physical Review X* **15**, 041012 (2025).
 - [114] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, red. by F. Bach, 2nd ed., Adaptive Computation and Machine Learning Series (MIT Press, Cambridge, MA, USA, 2018), 552 pp.
 - [115] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, “Mastering the game of Go with deep neural networks and tree search”, *Nature* **529**, 484 (2016).
 - [116] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. Lillicrap, K. Simonyan, and D. Hassabis, *Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm*, (Dec. 5, 2017) <http://arxiv.org/abs/1712.01815>, pre-published.
 - [117] M. H. S. Segler, M. Preuss, and M. P. Waller, “Planning chemical syntheses with deep neural networks and symbolic AI”, *Nature* **555**, 604 (2018).
 - [118] C. Jones, “Low-overhead constructions for the fault-tolerant Toffoli gate”, *Physical Review A* **87**, 022328 (2013).
 - [119] C. Gidney, “Halving the cost of quantum addition”, *Quantum* **2**, 74 (2018).
 - [120] T. Khattar and C. Gidney, “Rise of conditionally clean ancillae for efficient quantum circuit constructions”, *Quantum* **9**, 1752 (2025).

- [121] P. Cussenot, B. Grivet, L. Feldmann, S. Wengerowsky, B. P. Lanyon, T. E. Northup, H. de Riedmatten, A. S. Sørensen, and N. Sangouard, "Uniting Quantum Processing Nodes of Cavity-Coupled Ions with Rare-Earth Quantum Repeaters Using Single-Photon Pulse Shaping Based on Atomic Frequency Comb", *Physical Review Letters* **135**, 240803 (2025).
- [122] P. Caspar, E. Verbanis, E. Oudot, N. Maring, F. Samara, M. Caloz, M. Perrenoud, P. Sekatski, A. Martin, N. Sangouard, H. Zbinden, and R. T. Thew, "Heralded Distribution of Single-Photon Path Entanglement", *Physical Review Letters* **125**, 110506 (2020).
- [123] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden, "Quantum cryptography", *Reviews of Modern Physics* **74**, 145 (2002).
- [124] F. Marsili, V. B. Verma, J. A. Stern, S. Harrington, A. E. Lita, T. Gerrits, I. Vayshenker, B. Baek, M. D. Shaw, R. P. Mirin, and S. W. Nam, "Detecting single infrared photons with 93% system efficiency", *Nature Photonics* **7**, 210 (2013).
- [125] D. V. Reddy, R. R. Nerem, S. W. Nam, R. Mirin, and V. Verma, "Superconducting nanowire single-photon detectors with 98% system detection efficiency at 1550 nm", *Optica* **7**, 10.1364/OPTICA.400751 (2020).
- [126] K. Laiho, M. Avenhaus, K. N. Cassemiro, and C. Silberhorn, "Direct probing of the Wigner function by time-multiplexed detection of photon statistics", *New Journal of Physics* **11**, 043012 (2009).
- [127] R. Kumar, A. Wonfor, R. Penty, T. Spiller, and I. White, *Experimental demonstration of single-shot quantum and classical signal transmission on single wavelength optical pulse*, (July 22, 2019) <http://arxiv.org/abs/1907.09346>, pre-published.
- [128] A. Christ, K. Laiho, A. Eckstein, K. N. Cassemiro, and C. Silberhorn, "Probing multimode squeezing with correlation functions", *New Journal of Physics* **13**, 033027 (2011).
- [129] E. Wigner, "On the Quantum Correction For Thermodynamic Equilibrium", *Physical Review* **40**, 749 (1932).
- [130] E. C. G. Sudarshan, "Equivalence of Semiclassical and Quantum Mechanical Descriptions of Statistical Light Beams", *Physical Review Letters* **10**, 277 (1963).
- [131] R. J. Glauber, "Coherent and Incoherent States of the Radiation Field", *Physical Review* **131**, 2766 (1963).
- [132] P. Sekatski, N. Sangouard, F. Bussières, C. Clausen, N. Gisin, and H. Zbinden, "Detector imperfections in photon-pair source characterization", *Journal of Physics B: Atomic, Molecular and Optical Physics* **45**, 124016 (2012).

- [133] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, "Quantum entanglement", *Reviews of Modern Physics* **81**, 865 (2009).
- [134] V. C. Vivoli, P. Sekatski, J.-D. Bancal, C. C. W. Lim, A. Martin, R. T. Thew, H. Zbinden, N. Gisin, and N. Sangouard, "Comparing different approaches for generating random numbers device-independently using a photon pair source", *New Journal of Physics* **17**, 023023 (2015).
- [135] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, "Proposed Experiment to Test Local Hidden-Variable Theories", *Physical Review Letters* **23**, 880 (1969).
- [136] A. Peres, "Separability criterion for density matrices", *Physical Review Letters* **77**, 1413 (1996).
- [137] M. Horodecki, P. Horodecki, and R. Horodecki, "Separability of mixed states: necessary and sufficient conditions", *Physics Letters A* **223**, 1 (1996).
- [138] R. A. Howard, *Dynamic programming and Markov processes* (Technology Press of Massachusetts Institute of Technology, 1960), 152 pp.
- [139] V. Kliuchnikov, D. Maslov, and M. Mosca, *Fast and efficient exact synthesis of single qubit unitaries generated by Clifford and T gates*, (Feb. 27, 2013) <http://arxiv.org/abs/1206.5236>, pre-published.
- [140] K. Matsumoto and K. Amano, *Representation of Quantum Circuits with Clifford and $\pi/8$ Gates*, (June 24, 2008) <http://arxiv.org/abs/0806.3834>, pre-published.
- [141] Thompson and William, "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples", *Biometrika* **25**, 285 (1933).
- [142] W. R. Thompson, "On the Theory of Apportionment", *American Journal of Mathematics* **57**, 450 (1935).
- [143] H. Robbins, "Some aspects of the sequential design of experiments", *Bulletin of the American Mathematical Society* **58**, 527 (1952).
- [144] D. A. Berry and B. Fristedt, *Bandit problems* (Springer Netherlands, Dordrecht, 1985).
- [145] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules", *Advances in Applied Mathematics* **6**, 4 (1985).
- [146] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time Analysis of the Multi-armed Bandit Problem", *Machine Learning* **47**, 235 (2002).
- [147] R. Coulom, "Efficient Selectivity and Backup Operators in Monte-Carlo Tree Search", in (2006).

- [148] L. Kocsis and C. Szepesvári, “Bandit Based Monte-Carlo Planning”, in *Machine Learning: ECML 2006*, edited by J. Fürnkranz, T. Scheffer, and M. Spiliopoulou (2006), pp. 282–293.
- [149] S. Gelly, Y. Wang, R. Munos, and O. Teytaud, *Modification of UCT with Patterns in Monte-Carlo Go*, Research Report RR-6062 (INRIA, 2006).
- [150] S. Gelly and D. Silver, “Combining Online and Offline Knowledge in UCT”, *ACM International Conference Proceeding Series* **227**, 10.1145/1273496.1273531 (2007).
- [151] C. D. Rosin, “Multi-armed bandits with episode context”, *Annals of Mathematics and Artificial Intelligence* **61**, 203 (2011).
- [152] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, “Curriculum learning”, in *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML ’09 (June 14, 2009), pp. 41–48.
- [153] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, (Jan. 30, 2017) <http://arxiv.org/abs/1412.6980>, pre-published.
- [154] K. He, X. Zhang, S. Ren, and J. Sun, *Deep Residual Learning for Image Recognition*, (Dec. 10, 2015) <http://arxiv.org/abs/1512.03385>, pre-published.
- [155] R. Frostig, M. Johnson, and C. Leary, “Compiling machine learning programs via high-level tracing”, in (2018).
- [156] J. Christensen, *Jackc/pgx*, May 4, 2026.
- [157] J. Marshall, *Turbozero: fast + parallel AlphaZero*, Apr. 8, 2026.
- [158] C. P. Williams, *Explorations in quantum computing* (Springer London, 2011).
- [159] B. Giles and P. Selinger, “Exact synthesis of multiqubit Clifford+T circuits”, *Physical Review A* **87**, 032332 (2013).
- [160] D. Gosset, V. Kliuchnikov, M. Mosca, and V. Russo, *An algorithm for the T-count*, (Aug. 19, 2013) <http://arxiv.org/abs/1308.4134>, pre-published.
- [161] R. Haenel, X. Luo, and C. Zhao, *Tsim: Fast Universal Simulator for Quantum Error Correction*, (Apr. 1, 2026) <http://arxiv.org/abs/2604.01059>, pre-published.
- [162] B. Grvt, *BastienGrvt/q-mlink*, May 26, 2026.
- [163] B. Grvt, *BastienGrvt/q-fullrep*, May 23, 2026.
- [164] X. Valcarce, B. Grivet, and N. Sangouard, *Unitary Synthesis with AlphaZero via Dynamic Circuits*, (Aug. 28, 2025) <http://arxiv.org/abs/2508.21217>, pre-published.
- [165] X. Valcarce, *Xvalcarce/az_eus*, Dec. 4, 2025.

- [166] X. Valcarce, *Xvalcarce/quantum_compilation*, Sept. 2, 2025.
- [167] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. Sales Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin, “Logical quantum processor based on reconfigurable atom arrays”, *Nature* **626**, 58 (2024).
- [168] R. Acharya, D. A. Abanin, L. Aghababaie-Beni, I. Aleiner, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, N. Astrakhantsev, J. Atalaya, R. Babbush, D. Bacon, B. Ballard, J. C. Bardin, J. Bausch, A. Bengtsson, A. Bilmes, S. Blackwell, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, D. A. Browne, B. Buchea, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, A. Cabrera, J. Campero, H.-S. Chang, Y. Chen, Z. Chen, B. Chiaro, D. Chik, C. Chou, J. Claes, A. Y. Cleland, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, S. Das, A. Davies, L. De Lorenzo, D. M. Debroy, S. Demura, M. Devoret, A. Di Paolo, P. Donohoe, I. Drozdov, A. Dunsworth, C. Earle, T. Edlich, A. Eickbusch, A. M. Elbag, M. Elzouka, C. Erickson, L. Faoro, E. Farhi, V. S. Ferreira, L. F. Burgos, E. Forati, A. G. Fowler, B. Foxen, S. Ganjam, G. Garcia, R. Gasca, É. Genois, W. Giang, C. Gidney, D. Gilboa, R. Gosula, A. G. Dau, D. Graumann, A. Greene, J. A. Gross, S. Habegger, J. Hall, M. C. Hamilton, M. Hansen, M. P. Harrigan, S. D. Harrington, F. J. H. Heras, S. Heslin, P. Heu, O. Higgott, G. Hill, J. Hilton, G. Holland, S. Hong, H.-Y. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, S. Jordan, C. Joshi, P. Juhas, D. Kafri, H. Kang, A. H. Karamlou, K. Kechedzhi, J. Kelly, T. Khair, T. Khattar, M. Khezri, S. Kim, P. V. Klimov, A. R. Klotz, B. Kobrin, P. Kohli, A. N. Korotkov, F. Kostritsa, R. Kothari, B. Kozlovskii, J. M. Kreikebaum, V. D. Kurilovich, N. Lacroix, D. Landhuis, T. Lange-Dei, B. W. Langley, P. Laptev, K.-M. Lau, L. Le Guevel, J. Ledford, J. Lee, K. Lee, Y. D. Lensky, S. Leon, B. J. Lester, W. Y. Li, Y. Li, A. T. Lill, W. Liu, W. P. Livingston, A. Locharla, E. Lucero, D. Lundahl, A. Lunt, S. Madhuk, F. D. Malone, A. Maloney, S. Mandrà, J. Manyika, L. S. Martin, O. Martin, S. Martin, C. Maxfield, J. R. McClean, M. McEwen, S. Meeks, A. Megrant, X. Mi, K. C. Miao, A. Mieszala, R. Molavi, S. Molina, S. Montazeri, A. Morvan, R. Movassagh, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, H. Neven, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, C.-H. Ni, M. Y. Niu, T. E. O’Brien, W. D. Oliver, A. Opremcak, K. Ottosson, A. Petukhov, A. Pizzuto, J. Platt, R. Potter, O. Pritchard, L. P. Pryadko, C. Quintana, G. Ramachandran, M. J. Reagor, J. Redding, D. M. Rhodes, G. Roberts, E. Rosenberg, E. Rosenfeld, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster,

- A. W. Senior, M. J. Shearn, A. Shorter, N. Shutty, V. Shvarts, S. Singh, V. Sivak, J. Skruzny, S. Small, V. Smelyanskiy, W. C. Smith, R. D. Somma, S. Springer, G. Sterling, D. Strain, J. Suchard, A. Szasz, A. Sztein, D. Thor, A. Torres, M. M. Torunbalci, A. Vaishnav, J. Vargas, S. Vdovichev, G. Vidal, B. Villalonga, C. V. Heidweiller, S. Waltman, S. X. Wang, B. Ware, K. Weber, T. Weidel, T. White, K. Wong, B. W. K. Woo, C. Xing, Z. J. Yao, P. Yeh, B. Ying, J. Yoo, N. Yosri, G. Young, A. Zalcman, Y. Zhang, N. Zhu, N. Zobrist, and Google Quantum AI and Collaborators, "Quantum error correction below the surface code threshold", *Nature* **638**, 920 (2025).
- [169] C. Gidney, *How to factor 2048 bit RSA integers with less than a million noisy qubits*, (May 21, 2025) <http://arxiv.org/abs/2505.15917>, pre-published.
- [170] N. de Beaudrap and D. Horsman, "The ZX calculus is a language for surface code lattice surgery", *Quantum* **4**, 218 (2020).
- [171] D. B. Tan, M. Y. Niu, and C. Gidney, "A SAT Scalpel for Lattice Surgery: Representation and Synthesis of Subroutines for Surface-Code Fault-Tolerant Quantum Computing", in *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)* (June 29, 2024), pp. 325–339.
- [172] J. Zhou, Y. Liu, E. Decker, J. Kalloor, M. Weiden, K. Chen, C. Iancu, and G. Li, *TopoLS: Lattice Surgery Compilation via Topological Program Transformations*, (Mar. 28, 2026) <http://arxiv.org/abs/2601.23109>, pre-published.
- [173] R. Caruana, "Multitask Learning", *Machine Learning* **28**, 41 (1997).
- [174] M. Vasmer and D. E. Browne, "Three-dimensional surface codes: Transversal gates and fault-tolerant architectures", *Physical Review A* **100**, 012312 (2019).
- [175] B. Hayes, "First Links in the Markov Chain", *American Scientist* **101**, 92 (2013).
- [176] A. A. Марковъ, "Примѣръ статистическаго изслѣдованія надъ текстомъ „Евгенія Онѣгина“ иллюстрирующий связь испытаній въ цѣпь", *Извѣстія Императорской Академіи Наукъ* (Bulletin de l'Académie Impériale des Sciences de St.-Pétersbourg), 153 (1913).
- [177] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, 1st ed., Wiley Series in Probability and Statistics (Wiley, 1994).
- [178] R. I. C. H. A. R. D. BELLMAN, "A markovian decision process", *Journal of Mathematics and Mechanics* **6**, 679 (1957).
- [179] C. J. C. H. Watkins and P. Dayan, "Q-learning", *Machine Learning* **8**, 279 (1992).

- [180] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning", [Machine Learning](#) **8**, 229 (1992).
- [181] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms", 2017.
- [182] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains", [Artificial Intelligence](#) **101**, 99 (1998).