universität wien

DISSERTATION | DOCTORAL THESIS

Titel | Title

Phonon-assisted driving of quantum dot single-photon sources and its applications in quantum cryptography

verfasst von | submitted by

Lennart Louis Jehle M.Sc.

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Doktor der Naturwissenschaften (Dr.rer.nat.)

Wien | Vienna, 2026

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 796 605 411
Dissertationsgebiet lt. Studienblatt | Field of study as it appears on the student record sheet:	Physik
Betreut von | Supervisor:	Univ.-Prof. Dipl.-Ing. Dr. Philip Walther

Statement of own contribution

Robust excitation of C-band quantum dots for quantum communication

My contribution: With support from Michal Vyvlečka, I built and designed the pulse shaper to consistently tune the excitation wavelength and the analysis setup. Together with Michal Vyvlečka, we characterized the sample to choose a quantum dot (QD) for the experiment. The project was developed by Michal Vyvlečka, Mathieu Bozzio, and myself. Together with Michal Vyvlečka, I performed measurements and, with his input, processed and visualized the experimental data. I implemented the theoretical extrapolation of our results for quantum key distribution (QKD) under the guidance of Mathieu Bozzio.

The quantum dot sample was provided by our collaborators from Peter Michler's group at the University of Stuttgart. All experimental measurements were performed at the University of Vienna.

The manuscript [1] was published on 23.10.2023 by the American Institute of Physics in *Applied Physics Letters* under a CC-BY-4.0 license.

Asymmetric two-photon response of an incoherently excited quantum emitter

My contribution: The research plan and project goals were conceived and developed by myself. While the original QD excitation and collection setup used in the above work required only recalibration for the new QD sample and minor modifications for the laser pulse preparation, I had to design and build the analysis setup for the temporal and spectral correlation measurements. I collected the experimental data, with participation of Thomas W. Sandø, and processed it with input from Lena M. Hansen and Patrik I. Sund.

The quantum dot sample was provided by our collaborators from Peter Michler's group at the University of Stuttgart. All experimental measurements were performed at the University of Vienna.

The manuscript [2] is online available under a CC-BY-4.0 license and was submitted to a peer-reviewed journal on 01.10.2025.

Experimental single-photon limited-knowledge proofs

My contribution: The single-photon polarization encoding setup was designed by Michal Vyvlecka and myself, and I built, optimized, and characterized its performance. With assistance from Francesco Giorgino, I conceived and implemented electrical synchronization within the QD lab. I built the receiver's station, implemented automatic polarization compensation, and wrote code to synchronize the experiment across both labs. The experimental data was collected and processed by myself. I supported Rebecca Hofer and Mathieu Bozzio in including experimental imperfection in the security analysis of the protocol. With input from Tobias Guggemos, Rebecca Hofer and myself implemented the computational experiments and analyzed the results.

The quantum dot sample was provided by our collaborators from Peter Michler's group at the University of Stuttgart. All experimental measurements were performed at the University of Vienna.

The manuscript was submitted to a peer-reviewed journal on 02.02.2026 and is yet unpublished.

Statement regarding the usage of artificial intelligence

Throughout the research phase, AI-enhanced tools were used only for literature searches. All research questions and experimental designs were developed without AI-enhanced tools. All results were generated, interpreted, and analyzed without AI-enhanced tools. The writing of this thesis was assisted by Grammarly, a tool for grammar and spelling checks.

Abstract

In the past two decades, QDs have become essential sources of quantum light for a wide range of applications in quantum information and a promising platform for fundamental experiments in quantum optics. The integration of these artificial atoms into a solid-state environment offers exceptional control over their emission properties, and their simultaneously high single-photon purity and efficiency make QDs highly appealing for quantum technologies, including optical quantum computing and quantum-secured communication.

The scope of the thesis is twofold. At the fundamental level, it studies the dynamics of light–matter interactions that emerge when QDs are probed by laser pulses, placing a special focus on the influence of phonons; specifically, on longitudinal acoustic (LA) phonon-assisted driving, which leverages the QD's coupling to its surrounding solid-state environment to off-resonantly excite the system via phonon-mediated relaxation. Exploring the method in depth, we analyze how the photon-number statistics of the emission are modulated by the laser's power and wavelength. Furthermore, we show that the emission frequency is dynamically shifted by the instantaneous intensity of the driving laser field, an effect known as AC-Stark shift. For the first time, we demonstrate that, due to this shift, a two-photon response arising from re-excitation exhibits an asymmetric spectrum, and we present a technique exploiting the asymmetry to restore high single-photon purity even in the presence of re-excitation.

The second emphasis is on applying the above findings to single-photon quantum cryptography. Here, the phonon-assisted scheme offers significant advantages, including its intrinsically phase-randomized photonic output, straightforward laser suppression, and robust implementation. Throughout this thesis, we use QD sources that directly emit in the telecom C-band, thereby ensuring minimal transmission losses and compatibility with existing communication infrastructure. While weak coherent state (WCS) are fully characterized by their mean photon number, solid-state single-photon sources may exhibit a more complex relationship between their efficiency and purity, as shown by our first study. Interestingly, we find that there is no global optimum but that the ideal excitation parameters depend on the cryptographic protocol and its required security guarantees. In the final work, the insights from the fundamental investigations are applied to experimentally realize a loss-tolerant quantum bit commitment (QBC) scheme using the C-band emission of our QD. Since it remains unclear how to adapt the decoy-state technique commonly used with WCS in QKD to primitives in which the sender and receiver mutually mistrust one another (which is the case in our QBC), the high single-photon purity provided by QDs is particularly beneficial for these mistrustful protocols, enabling

QDs to outperform WCS across all realistic scenarios. We use the QBC as a key building block for a new cryptographic primitive, which we term limited-knowledge proofs, and prove its information-theoretic security in the bounded-storage model. The protocol is implemented over a deployed fiber link connecting two university buildings and using actively polarization-encoded BB84 states. The polarization drift of the link is compensated by an automated routine, achieving an average quantum bit error rate of less than 0.5 %.

Zusammenfassung

In den letzten zwei Jahrzehnten haben sich Quantenpunkte (QP) zu einer unverzichtbaren Quelle von Quantenlicht entwickelt, die in einer Vielzahl von Anwendungen im Bereich der Quanteninformation eingesetzt werden, und sich zugleich als vielversprechende Plattform zur Umsetzung grundlegender Experimente in der Quantenoptik etabliert haben. Die Integration dieser künstlichen Atome in eine Festkörperumgebung ermöglicht eine außergewöhnliche Kontrolle ihrer Emissionseigenschaften. Die gleichzeitig hohe Einzelphotonenreinheit und Effizienz machen QP für eine Reihe von Quantentechnologien, darunter optisches Quantencomputing und quantengesicherte Kommunikation, äußerst attraktiv.

Die Fragestellung dieser Dissertation ist zweigeteilter Natur. Auf fundamentaler Ebene untersucht sie die Dynamik der Licht-Materie-Wechselwirkungen durch Laserimpulse angeregter QP. Hierbei liegt ein besonderer Schwerpunkt auf dem Einfluss von longitudinalen akustischen Phononen, da die Kopplung des QP an seine umgebende Festkörperumgebung es erlaubt, diesen mittels eines nicht resonanten Laserpulses unter Emission eines Phonons in den optisch aktiven Zustand zu überführen. Die umfassende Analyse dieser Anregungsmethode zeigt, wie die Photonenzahlstatistik der Emission durch Leistung und Wellenlänge des Lasers moduliert wird. Darüber hinaus zeigen wir, dass bei der phononunterstützten Anregung die Emissionswellenlänge dynamisch und proportional zur instantanen Intensität des treibenden Laserfeldes verschoben wird; ein Effekt, der auch als AC-Stark-Verschiebung bekannt ist. Dabei zeigen wir erstmals, dass das Zwei-Photonen-Spektrum, das durch Reanregung entsteht, asymmetrisch ist, und stellen eine Technik vor, die diese Asymmetrie nutzt, um trotz Reanregung eine hohe Einzelphotonenreinheit wiederherzustellen.

Der zweite Schwerpunkt der Arbeit liegt auf der Anwendung der oben genannten Erkenntnisse für die Einzelphotonen-Quantenkryptografie. Hier bietet die phononen-assistierte Anregung erhebliche Vorteile: eine intrinsisch zufällig verteilte Phasenbeziehung der erzeugten Lichtzustände, eine einfach umzusetzende Laserunterdrückung und eine robuste Implementierung. In dieser Arbeit verwenden wir QP, die direkt im Telekommunikations-C-Band emittieren, wodurch minimale Übertragungsverluste und die Kompatibilität mit bestehender Infrastruktur gewährleistet werden. Während abgeschwächte kohärente Zustände (WCS) vollständig durch ihre mittlere Photonenzahl charakterisiert sind, weisen Festkörper-Einzelphotonenquellen eine komplexere Beziehung zwischen Effizienz und Reinheit auf, wie unsere erste Studie gezeigt hat. Interessanterweise stellen wir fest, dass es kein globales Optimum gibt, sondern die idealen Anregungsparameter vom jeweiligen kryptografischen Protokoll und dessen Sicherheitsgarantien abhängen. In der letzten Arbeit werden

die Erkenntnisse aus diesen grundlegenden Untersuchungen angewendet, um ein verlusttolerantes Quantenbit-Commitment (QBC) mithilfe der C-Band-Photonen unseres QPs experimentell zu demonstrieren. Da bisher unklar ist, wie die bei WCS üblicherweise verwendete Decoy-State-Methode an kryptographische Primitive angepasst werden kann, bei denen Sender und Empfänger einander misstrauen (wie der Fall bei unserem QBC), ist die hohe Einzelphotonenreinheit, die QPs bieten, für diese Klasse von Protokollen besonders vorteilhaft, und QP übertreffen WCS in allen realistischen Szenarien. Wir verwenden das QBC als wichtigen Baustein in einem neuen kryptografischen Primitiv, welches wir als Limited-Knowledge-Proof bezeichnen, und beweisen dessen informationstheoretische Sicherheit unter der Annahme eines Quantenspeichers begrenzter Größe (Bounded-Storage-Modell). Das Protokoll wird über eine Glasfaserverbindung zwischen zwei Universitätsgebäuden und mithilfe aktiv polarisationscodierter BB84-Zustände implementiert. Der Polarisationsdrift der Glasfaserverbindung wird durch eine automatisierte Routine kompensiert, wodurch eine durchschnittliche Quantenbitfehlerrate von weniger als 0,5 % erreicht wird.

Publications

Published

Robust excitation of C-band quantum dots for quantum communication
Michal Vyvlecka, Lennart Jehle, Cornelius Nawrath, Francesco Giorgino, Mathieu Bozzio, Robert Sittig, Michael Jetter, Simone L. Portalupi, Peter Michler, and Philip Walther
published on 23.10.2023, Applied Physics Letters, vol. 123, 174001, 2023
Chapter 4

Under peer review

Asymmetric two-photon response of an incoherently excited quantum emitter
Lennart Jehle, Lena M. Hansen, Patrik I. Sund, Thomas W. Sandø, Raphael Joos, Michael Jetter, Mathieu Bozzio, Simone L. Portalupi, Peter Michler, and Philip Walther
submitted on 01.10.2025, online available arXiv.2507.07082
Chapter 5

Experimental single-photon limited-knowledge proofs
Lennart Jehle, Rebecca Hofer, Tobias Guggemos, Francesco Giorgino, Danijela Milojevic, M. Vyvlecka, Raphael Joos, Michael Jetter, Simone L. Portalupi, Peter Michler, Mathieu Bozzio and Philip Walther
submitted on 02.02.2026
Chapter 6

Not directly related to this dissertation

Single-active-element demultiplexed multi-photon source
Lena M. Hansen, Lorenzo Carosini, Lennart Jehle, Francesco Giorgino, Romane Houvenaghel, Michal Vyvlecka, Juan C. Loredo, and Philip Walther
Optica Quantum 1, 1-5, 2023

Multi-photon emission from a resonantly pumped quantum dot
Francesco Giorgino, Patrik Zahálka, Lennart Jehle, Lorenzo Carosini, Lena M. Hansen, Juan C. Loredo, and Philip Walther
Optica Quantum 3, 402-407, 2025

Non-classical excitation of a solid-state quantum emitter
Lena M. Hansen, Francesco Giorgino, Lennart Jehle, Lorenzo Carosini, Juan Camilo López Carreño, Iñigo Arrazola, Philip Walther and Juan C. Loredo
arXiv.2407.20936

Contents

1 Introduction **1**

2 Preliminaries **5**
2.1 Modern cryptography 5
2.1.1 Introduction to cryptography 5
2.1.2 Quantum information in the context of quantum cryptography 7
2.1.3 Quantum cryptography using single photons 12
2.1.4 Quantum cryptographic primitives 18
2.2 Quantum dots as quantum light sources 32
2.2.1 Quantum dots in a nutshell 32
2.2.2 Figures of merit 35
2.2.3 Light–matter interaction 41
2.2.4 Optical excitation 48
2.2.5 Photonic extraction structures 59

3 Experimental methods **69**
3.1 Quantum dot excitation and collection 69
3.1.1 Quantum dot samples 69
3.1.2 Optical setup 70
3.2 Analyzing temporal information in quantum optics 72
3.2.1 Basic measurements 73
3.2.2 Time-gated measurements 76

4 Robust excitation of a C-band quantum dot for quantum cryptography **83**
4.1 Methods 84
4.1.1 Experimental setup 85
4.1.2 Data processing 85
4.1.3 Extrapolation of the SKR 87
4.1.4 Maximum QKD-distance approximation from experimental measures 87
4.2 Manuscript reprint 90
4.2.1 Introduction 90
4.2.2 Results 91
4.2.3 Discussion 96
4.3 Complementary material 98
4.3.1 Photon-number optimization via variable attenuation . . . 98
4.3.2 Time filtering 99

5 Asymmetric two-photon response of an incoherently driven quantum emitter 103
5.1 Methods . . . 104
5.1.1 Experimental setup . . . 105
5.1.2 Temporal & spectral two-photon measurements . . . 106
5.2 Manuscript reprint . . . 109
5.2.1 Introduction . . . 109
5.2.2 Results . . . 111
5.2.3 Discussion . . . 117
5.3 Complementary material . . . 119
5.3.1 Laser–QD detuning variations . . . 119
5.3.2 Frequency-filtered time traces . . . 120

6 Experimental single-photon limited-knowledge proofs 123
6.1 Methods . . . 124
6.1.1 Experimental setup . . . 125
6.1.2 Problem instance . . . 133
6.2 Manuscript reprint . . . 135
6.2.1 Introduction . . . 135
6.2.2 Limited-knowledge proofs . . . 137
6.2.3 Experimental demonstration . . . 139
6.2.4 Discussion . . . 144
6.3 Complementary material . . . 146
6.3.1 Experimental data . . . 146
6.3.2 Security analysis . . . 147
6.3.3 Extended parameter variations . . . 154
6.3.4 3-5 CNF generation and benchmarking . . . 157

7 Conclusion 163
7.1 Phonon-assisted excitation for single-photon applications . . . 163
7.2 Single-photon cryptography . . . 164
7.3 Challenges and perspectives . . . 165

List of Acronyms 167

A APPENDIX: Correlation functions 197

B APPENDIX: Security analysis for limited-knowledge proofs 199
B.1 Technical preliminaries . . . 199
B.1.1 Polynomial-time reductions . . . 199
B.1.2 Calculation of min-entropies . . . 201
B.2 Security against P* . . . 202
B.2.1 Length notations and finite-size effects . . . 202
B.2.2 Security bounds . . . 204

B.2.3 String definitions 205
B.2.4 Uncertainty about X_{v} 206
B.2.5 Depolarizing channel 208
B.3 Security against V* 209
B.3.1 Commitment phase 209
B.3.2 Verification phase 210

List of Figures

2.1 Single qubit representation of a Bloch sphere. 8
2.2 Bloch sphere representation of linear retarder states 17
2.3 Secure key rate for QKD over channel distance. 23
2.4 Phase discrimination for a single-photon state 29
2.5 Illustration of a QD and its energy structure 33
2.6 Excitonic complexes in a QD 34
2.7 Energy spectrum of QD-laser system interaction potential 42
2.8 QD eigenstates depending on laser coupling 44
2.9 Spontaneous emission from a dressed TLS 45
2.10 Spectral density of the electron-phonon coupling $J(\omega)$ 47
2.11 Dynamic LA phonon coupling 52
2.12 Schematic time series of a re-excitation event 54
2.13 Multiphoton probability of an ideal 2LS under RF 55

3.1 Characteristics of the CBG QD sample 70
3.2 Optical setup for QD excitation and collection 71
3.3 Examples of a one- and two-dimensional histogram 74
3.4 Reduction from $G^{(2)}(t_1, t_2)$ to $G^{(2)}(\tau)$ 75
3.5 Illustration of a reduced integration window for $G^{(2)}(\tau)$ 76
3.6 Comparison of gating strategies 78
3.7 Temporal distribution of the multiphoton component 79
3.8 Correlation logic using virtual channels 80
3.9 Individual photon time profiles from a 2d histogram 81
3.10 Photon profiles caused by re-excitation of a 2LS 82

4.1 Scheme of the experimental setup 85
4.2 Time-resolved fluorescence measurement 86
4.3 Characterization of the positively-charged exciton transition under pulsed LA excitation 92
4.4 Measured photon-number statistics and extrapolated QKD secure key bits per pulse for LA excitation 93
4.5 Secure key rates of BB84–QKD for varying communication distances 95
4.6 Secure key bits per pulse over distance for variable signal attenuation 99
4.7 Comparing time filtering and pump power tuning for QKD 101

5.1 Experimental setup including QD excitation and analysis stage . . 105
5.2 Schematic spectro-temporal analysis setup 106

5.3 Schematic time series of a re-excitation event 110
5.4 Re-excitation under phonon-assisted driving 111
5.5 Re-excitation spectrum for varying pump power 113
5.6 Reducing the multiphoton contribution by frequency filtering . . . 116
5.7 QD spectrum for blue- and red-detuned laser at varying powers . . 120
5.8 Two-photon spectrum for varying laser–QD detunings 121
5.9 Histogram of frequency filtered QD emission 122

6.1 Scheme of the experimental setup 125
6.2 Characteristics of the QD emission 127
6.3 Autocorrelation measurements of the QD 128
6.4 Optical channel gating 129
6.5 Characterization of the polarization encoding 130
6.6 Schematic layout of the synchronization setup 132
6.7 Limited-knowledge proof using quantum commitments 136
6.8 Experimental implementation of the quantum commitment phase . 140
6.9 Maximum leakage rate for V^* 141
6.10 Limited-knowledge performance 143
6.11 Analysis of the full detected data 146
6.12 Maximally tolerable quantum bit error rate 155
6.13 Verification performance under parameter variations 156
6.14 Performance of V^* 159
6.15 Conditions for a secure protocol 160

1 Introduction

Contrary to a common misconception that quantum cryptography was developed in response to the security risks posed by quantum algorithms, many date the birth of the field back to Wiesner's seminal paper in 1983 [3], thus preceding groundbreaking works in quantum computing, such as those by Shor and Grover, at least by a decade. However, these advances have undoubtedly attracted interest from a broader audience, including public funding and government agencies, and also added a sense of urgency to the challenge of designing cryptosystems that prevail even in the presence of a fault-tolerant quantum computer. Three decades and immense investments in quantum computing later, widely-used encryptions including RSA and Diffie-Hellman are still considered secure today, but the progress across various platforms has been impressive [4–9], and systems sufficiently large to break those encryptions are getting ever closer [10, 11]. Hence, the need to design practical cryptographic protocols that guarantee information-theoretic security—i.e., without relying on any computational assumptions—in real-world implementations remains as strong as ever.

In quantum cryptography, such guarantees can be provided solely by assuming the validity of quantum mechanics; in particular, the no-cloning theorem [12], stating that no quantum state can be perfectly copied, and conjugate coding [3], inspired by the uncertainty principle, which dictates that two conjugate observables cannot be measured simultaneously with arbitrary precision.

Interestingly, since quantum cryptography and quantum computing share much of their theoretical foundations in quantum information and, in parts, even the same hardware, they often benefit from one another. For example, a recent and still ongoing endeavor with applications in both fields is the implementation of a highly efficient and noise-tolerant quantum memory [13–18]. Similarly, the development of new quantum light sources has long been driven by both fundamental questions and the advancement of quantum technologies. The experimental realization of bipartite entanglement was originally pursued to settle the long-standing foundational debate about local hidden variable theory [19–23], but the early experiments eventually led to multi-photon experiments demonstrating quantum teleportation [24, 25], entanglement swapping [26], as well as the generation of three-photon entangled states[27, 28] and their non-locality [29]; all of which are now essential to quantum information technologies.

While to date the race for the best physical qubit for building a quantum computer remains undecided, and likely multiple specialized platforms will emerge, the very nature of communication defines photons as the ideal choice for transmitting quantum information over long distances. Early implementations employed strongly attenuated and actively encoded laser pulses [30–32], but with the evolution of quantum photonics, heralded single-photon and entanglement sources followed soon after. This first wave of progress was driven by spontaneous parametric down-conversion (SPDC), creating entangled photon pairs, which allowed passive state encoding, stronger security notions using (measurement) device independence [33–36], and first demonstrations of a quantum repeater [37, 38] to extend the communication ranges.

However, the photon distribution of an SPDC source—just like for a WCS from an attenuated laser—is governed by Poisson statistics such that the probabilities of creating one, two, and more photons are intrinsically linked. While the principles of no-cloning and conjugate coding ensure security for information imprinted on a single-photon pulse, sending more than one photon clearly compromises this security [39]. And although the decoy-state technique [40, 41] mitigates the multiphoton risk of a Poisson-distributed source (PDS) remarkably well for some protocols, researchers began to pursue the development of a "true" single-photon source (SPS), which emits exactly one photon when triggered (i.e. exhibits sub-Poissonian statistics) [42, 43]. An ideal SPS must have high system efficiency and single-photon purity; for some protocols, the photon indistinguishability must be maximized simultaneously. Today, such quantum emitters have been shown to outperform decoy-state QKD with PDS [44], and even more substantial advantages are expected for protocols that are more sensitive to multiphoton noise. For example, multiphoton contributions limit the fidelity of repeater operations [38] and the performance of many *mistrustful* cryptographic protocols [45]; for strong coin-flipping, the advantage of SPSs over PDSs was also recently experimentally demonstrated [46].

The first breakthroughs in the realization of an SPS [47–49] were quickly followed by significant improvements in emission rates, coherence, indistinguishability, and the generation of entangled photon pairs [50–54]. Today, epitaxially-grown quantum dots have established themselves as the leading platform for efficient generation of high-quality single photons[55, 56], and more recently also as a candidate for creating multipartite entangled states [57–61]. They enable foundational studies of quantum optics [62–66] and are used for a wide range of quantum information technologies [67–73]. Depending on the application and material composition, which determines the emission wavelength, single photons may be used in the near-infrared, frequency converted [74–77], or generated directly in the telecom O- or C-band [73, 78–80].

Compared to other emerging platforms for SPSs, QDs profit from pre-existing, highly advanced fabrication techniques for III-V semiconductors, which have enabled their integration into microstructures to manipulate the electronic and photonic

environment early on. Most notable are tunnel-barrier diodes [81] that enhance indistinguishability, control the charge state, and suppress fluctuations therein (i.e., blinking) [82–84], and cavities [55, 68, 85, 86] that funnel radiative emission into a specific mode, thereby creating Purcell enhancement [87] and improving collection efficiencies. Their success as a highly efficient light–matter interface has motivated the community to extend the range of functionality beyond that of an SPS, and QDs have been demonstrated to serve as quantum memories [88, 89] or for spin-mediated entanglement generation [60, 61, 90].

Returning to their application for cryptography, in particular practical implementations, several more aspects must be considered. Since no realistic source emits a pure single-photon state, phase randomization in the number state basis, as known from PDS [91], must be ensured for most protocols. From an experimental point of view, the operation of the SPS must be simple and efficient, but also reasonably robust to enable real-world demonstrations. These characteristics are influenced not only by the features of the QD sample, such as the presence of an electric diode or photonic cavity, but also strongly depend on the optical driving scheme. Over the past two decades, numerous excitation methods have been demonstrated [71, 92–94], and no single technique fits all needs; instead, the advantages must be balanced in light of a given application.

This work takes a holistic approach to studying and optimizing the potential of QDs for applications in practical quantum cryptography. It spans the full spectrum from fundamental investigation of the light-matter interaction specific to solid-state emitters, to tailoring the QD emission properties to primitives and the characteristics of the quantum channel, to designing and demonstrating a new protocol that leverages the strengths of these quantum light sources. Given the focus on quantum communication, all experiments are performed using QDs that emit directly in the telecom C-band to minimize channel loss. They have been kindly provided by our collaborators from Prof. Peter Michler's group at the University of Stuttgart. Furthermore, the experiments exclusively employ LA phonon-assisted excitation, which is highly suitable for quantum cryptography due to its robustness against laser fluctuations (as we will show [1]) and the inherently phase-randomized emission.

In detail, Chapter 4 explores the excitation parameter space that becomes available when using phonon-assisted driving, explicitly the laser wavelength and pulse area. Initiated by a theoretical study suggesting that the single-photon efficiency and purity may not be optimized simultaneously [95], we measure single-photon and multiphoton number probabilities, p_1 and p_m respectively, to infer the ideal operation point. Interestingly, we find that the optimal excitation parameters depend on the cryptographic primitive and additional experimental parameters such as channel loss, detection efficiency, and error rate [1]. Together, they determine how p_1 and p_m must be balanced to maximize performance.

In this first work, the laser pulse length was kept "short" (the pulse length

Δt_L is much smaller than the radiative decay time τ_{rad}) to limit multiphoton noise caused by re-excitation [96, 97]. However, even under these circumstances p_{m} showed a non-trivial dependence on the pulse energy, which was attributed to the phonon coupling and motivated an in-depth analysis of the dynamics. Varying the ratio $\Delta t_L/\tau_{\mathrm{rad}}$, the scaling of the re-excitation probability is reported for the first time for LA excitation and, even more importantly, a redshift in the emission wavelength of the first photon is exposed [2]. This asymmetry in the two-photon response is explained in accordance with the dynamical dressed-state model. It is then also experimentally exploited to significantly reduce the impact of re-excitation using spectral filtering, demonstrating a multiphoton component independent of the laser pulse length. These experiments, described in Chapter 5, were enabled by a significantly improved QD source based on a circular Bragg grating (CBG) cavity, which provides a much higher signal-to-noise ratio (SNR) and a reduced radiative lifetime. Notably, the results inspired further investigations that tackle the impact various cavity designs have on the re-excitation process for a phonon-assisted driven QD .

In parallel, we developed a new *mistrustful* cryptographic protocol that builds on an *i.t.*-secure quantum bit commitment (in the bounded-storage model) to enable what we term *limited-knowledge* proofs. These allow one party to convince another that they know the solution to a hard problem while disclosing only a small, constant fraction of that solution, so the problem can be reused multiple times. We implemented the protocol over a deployed fiber link using active polarization state encoding, demonstrating not only its robustness and feasibility in a realistic environment but also a significant rate and distance advantage over PDS, since the underlying quantum bit commitment is particularly sensitive to multiphoton noise. These unpublished results are detailed in Chapter 6.

2 Preliminaries

Reflecting the two-sided nature of this work, the preliminaries are divided into two main sections. First, we will introduce concepts of modern cryptography in 2.1. After recalling standard techniques of quantum information in Sec. 2.1.2, the focus shifts to the photonic aspects of quantum cryptography (Sec. 2.1.3), mainly the different states of light and how a qubit may be realized using photons. In Sec. 2.1.4, we combine elements from the previous sections to introduce an example of each a trustful (Sec. 2.1.4.1) and mistrustful (Sec. 2.1.4.2) cryptographic primitive. This part of the preliminaries concludes by discussing attacks exploiting imperfections present in practical implementations (Sec. 2.1.4.3).

In the second half (Sec. 2.2), we turn toward the photonic hardware used in this work: epitaxially grown semiconductor QDs serving as quantum light sources. In particular, their energy-level structure and desired emission properties (Sec. 2.2.2) are presented before detailing the light-matter interaction that underlies their single-photon emission (Sec. 2.2.3). Sec. 2.2.4 is dedicated to the realm of optical excitation schemes, due to their large variety and importance for cryptographic protocols. Finally, Sec. 2.2.5 provides an overview of micro- and nanostructures that may be employed to enhance the light-matter interaction as well as photonic collection efficiencies.

2.1 Modern cryptography

2.1.1 Introduction to cryptography

Cryptography is the science of securing information and communication, ensuring that sensitive data remains protected from unauthorized access or manipulation. It has been an essential tool for safeguarding secrets throughout history, evolving from simple substitution ciphers to sophisticated mathematical algorithms that underpin modern digital security. At its core, cryptography provides security guarantees such as confidentiality, ensuring that only intended recipients can access the information; integrity, protecting data from unauthorized alterations; authentication, verifying the identity of communicating parties; and non-repudiation, preventing individuals from denying their actions. These guarantees are fundamental to the functioning of modern digital systems, from online banking and e-commerce to secure messaging and digital governance.

Starting from only a few cryptographic primitives, the building blocks of modern cryptography, complex protocols can be constructed through an artful combination of these primitives. For example, almost every cryptosystem uses an encryption algorithm to ensure confidentiality by transforming plaintext into ciphertext that can only be decrypted by the rightful recipient. Other widely used primitives include hash functions, message authentication codes, digital signatures, commitments, or zero-knowledge proofs. Together, these primitives form the backbone of cryptographic protocols, enabling secure communication and data protection in diverse applications.

Until about 50 years ago, all encryption schemes used a secret key–shared only between the sender and receiver—to first encrypt the plaintext and then decrypt the ciphertext on the other side. When combined with one-time pad encryption [98]—encoding every bit of information using one bit of the shared key—such a scheme offers perfect security as long as the key is used only once. This symmetric-key cryptography is computationally efficient and is still widely used to encrypt large volumes of data. However, it requires the communication parties to hold a secret shared key—a premise that is sometimes difficult to fulfill in the modern world. The invention of asymmetric cryptography—using a public key for encryption and a private key for decryption—gave rise to large open networks like the internet, where far-separated parties can converse securely without ever meeting beforehand or sharing a secret. Unfortunately, the security of this approach relies on the assumption that inverting the underlying encryption function is extremely inefficient, rendering the exchange of messages secure for the relevant timescales. A prominent example of such systems is RSA encryption [99], which is based on the mathematical hardness of prime factorization.

Comparing strategies that realize the same primitive under differing constraints highlights an important concept in cryptography: at the cost of additional assumptions, tasks that would otherwise be impossible may become possible. Consequently, cryptographic systems are classified according to their security promises: information-theoretic security (*i.t.-security*) provides protection even against adversaries with unlimited computational power, while *computational security* assumes a hard problem that requires solving times that are non-polynomial in input size. When safeguarding an application with cryptographic tools, one must carefully balance the functionalities offered by computational security against the associated assumptions, and crucially, ascertain that these assumptions hold for the specific use case. In particular, with the onset of quantum algorithms, the hardness of the underlying problem must be revisited. Famously, Shor's algorithm based on quantum Fourier transformation [100] could break both asymmetric key pairs established using RSA and symmetric key schemes (such as Diffie-Hellman) [101], if executed on a quantum computer with enough error-corrected qubits.

Even before the advance of quantum algorithms, researchers began exploring the potential quantum information might offer to the field of cryptography, recognizing

that classical cryptography, despite its achievements, has fundamental limitations. Quantum cryptography, then, is the field of cryptography that uses quantum resources to ensure protection against both classical and quantum adversaries, aiming to achieve higher security guarantees based on the validity of quantum mechanics. The key principles that shape quantum information and benefit cryptography are quantum superposition and interference, the no-cloning theorem, and entanglement.

Although quantum cryptography also enables fundamentally new primitives (which cannot be realized classically, not even with computational security), such as public-key quantum tokens [3], the primary interest remains in developing *i.t.*-secure counterparts to classical primitives that rely on computational assumptions. As mentioned above, classically, remote key exchange cannot be performed in an *i.t.*-secure way, but the task becomes possible using quantum resources (assuming a classical authenticated channel). Today, 40 years after the foundational work of Wiesner [3] and Bennett & Brassard [30], QKD is a highly mature technology with impressive demonstrations over large distances on land [102] and in space [103], and a substantial commercial industry has grown around it.

However, while secure key exchange is indispensable, designing complex cryptosystems requires also other primitives. These may be considered *trustful*, if at least one party acts honestly (e.g., key distribution or quantum tokens), or *mistrustful*, if the parties have competing interests, which incentivize them to behave dishonestly (e.g., digital signatures, bit commitment, and others). The latter class also features primitives that have been proven to be impossible even with quantum resources [104, 105], unless new assumptions are included. For this thesis, a particularly relevant assumption is the noisy storage model [106], which protects against an adversary with imperfect and/or bounded quantum memory, and is commonly used to enable perfect and unbiased quantum bit commitment. Another popular choice is to restrict the adversary's power using constraints arising from special relativity [107, 108].

2.1.2 Quantum information in the context of quantum cryptography

2.1.2.1 The qubit

The qubit (**qu**antum **bit**) is the quantum analog of the classical bit and serves as the fundamental unit of quantum information. A qubit is represented by a quantum state in a two-dimensional complex vector space (Hilbert space), typically denoted as $\mathbb{C}^2$. The state of a qubit can be expressed as a linear combination (or coherent superposition) of two basis states, which span its Hilbert space. The generic qubit state is commonly written in terms of the computational basis states using Dirac notation, $|0\rangle$ and $|1\rangle$ corresponding to the logical 0 or 1,

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle\,, \tag{2.1-1}$$

where α and β are complex numbers satisfying the normalization condition $|\alpha|^2 + |\beta|^2 = 1$. The coefficients α and β represent the probability amplitudes of the qubit

being in the states $|0\rangle$ and $|1\rangle$, respectively. Unlike classical bits, which are either 0 or 1, a qubit can exist in a superposition of both states simultaneously, enabling quantum systems to encode and process information in a fundamentally new way.

Alternatively, using the normalization condition, the complex weights can be rephrased as angles

$$|\Psi\rangle = \cos(\theta/2)\,|0\rangle + \mathrm{e}^{i\varphi}\sin(\theta/2)\,|1\rangle\ , \tag{2.1-2}$$

where a global phase was dropped. The mathematical description using two angles θ and φ lends itself to the graphical representation of any qubit state as a point on a sphere with the computational basis states on either pole. This so-called Bloch sphere is shown in Fig. 2.1.

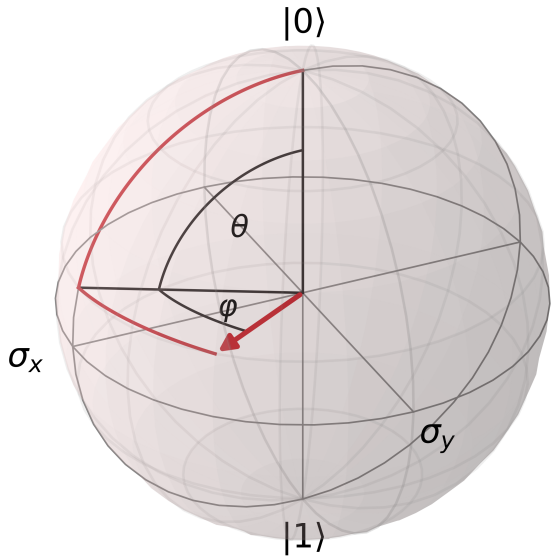


Figure 2.1: **Single qubit representation of a Bloch sphere.** An exemplary qubit state according to Eq. 2.1-2 with $\theta = 75°$ and $\varphi = 45°$ is shown on the Bloch sphere. The three Cartesian axes indicate the qubit's projection on the Pauli operators (see Eq.2.1-8), where the eigenstates of the Pauli operator σ_z are associated with the computational basis and are commonly denoted as $|0\rangle$ and $|1\rangle$.

2.1.2.2 Multiple qubits

When multiple qubits are combined, they form a multi-qubit system, whose state is described by the tensor product of individual qubit states. For example, the state of a two-qubit system is represented in a four-dimensional Hilbert space spanned by the computational basis states $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$[1]. An arbitrary two-qubit

[1]A short comment on notation: depending on the context, multi-qubit states composed of qubit A and qubit B may be written with various levels of abbreviation $|00\rangle = |0_A 0_B\rangle = |0\rangle_A \otimes |0\rangle_B$.

state can then be written as:

$$|\psi\rangle = \alpha_{00}\,|00\rangle + \alpha_{01}\,|01\rangle + \alpha_{10}\,|10\rangle + \alpha_{11}\,|11\rangle\,, \tag{2.1-3}$$

where α_{ij} are complex weights satisfying the normalization condition $\sum_{i,j}|\alpha_{ij}|^2 = 1$.

Multi-qubit states can exhibit entanglement, a phenomenon unique to quantum systems, where the state of one qubit becomes intrinsically correlated with the state of another, regardless of the physical distance between them. John Bell showed that these correlations are stronger than possible for classical systems [20], and the loophole-free violation of the Bell inequality became a key metric to verify entanglement experimentally [109]. For example, the Bell states are maximally entangled two-qubit states, such as

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle), \tag{2.1-4}$$

which cannot be decomposed into a tensor product of individual qubit states. Entanglement is a key resource in quantum information, enabling applications such as quantum teleportation and quantum computation.

2.1.2.3 Density operator

An alternative but mathematically equivalent formalism to the state-based description of quantum mechanics uses so-called density operators. This framework is particularly powerful when dealing with systems that are not in a *pure state*—states that are fully described by a single wavefunction—but are part of a statistical ensemble. The latter are called *mixed* states and arise when a system is in a probabilistic mixture of different pure states or when it interacts with an environment, leading to decoherence[2]. After introducing the density operator, the definitions of *pure* and *mixed* states are also mathematically formalized.

For an ensemble of pure states $\{p_i, |\psi_i\rangle\}$, the density operator is defined as

$$\rho = \sum_i p_i\,|\psi_i\rangle\,\langle\psi_i|\,, \tag{2.1-5}$$

where $\langle\psi|$ is the dual state vector (the state vector's conjugate transpose). Since p_i are real probabilities > 0, the density operator satisfies the following properties

(i) ρ is Hermitian: $\rho^\dagger = \rho$

(ii) ρ is positive semi-definite: $\langle\phi|\,\rho\,|\phi\rangle \geq 0$ for all $|\phi\rangle$

(iii) $\mathrm{Tr}(\rho) = 1$, following from the normalization $\sum_i p_i = 1$.

[2]Effectively, the interaction entangles our system with the environment, but since the environment is not tracked by the density matrix of our system, it appears as decoherence. In the next section, using the partial trace, we see how the 1-qubit subsystem of a Bell state is in the marginal—that is, fully mixed—state.

For a pure state, Eq. 2.1-5 reduces to a rank-1 projection operator $\rho = |\psi\rangle\langle\psi|$ and the density matrix is idempotent $\rho^2 = \rho$. Defining the purity of a state as

$$\mathcal{P} = \mathrm{Tr}[\rho^2], \tag{2.1-6}$$

provides a mathematical tool to distinguish *pure* and *mixed* states. From the idempotency and property iii) of the density operator, $\mathcal{P} = 1$ follows for pure states, whereas for mixed states $\mathcal{P} < 1$.

Importantly, linear combinations of pure states are *coherent* superpositions and should not be confused with statistical combinations of states:

$$\begin{aligned} |\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \rightarrow \quad & \rho_{\text{pure}} = \frac{1}{2}\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \rightarrow \quad \mathcal{P} = 1 \\ \rho_{\text{mixed}} = \frac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|) \rightarrow \quad & \rho_{\text{mixed}} = \frac{1}{2}\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \rightarrow \quad \mathcal{P} = \frac{1}{2} \end{aligned} \tag{2.1-7}$$

The coherence of pure states is signified by the off-diagonal elements of the density operator, and it is this coherence that enables effects like entanglement or quantum interference. Mixed states, on the other hand, lack quantum coherence, as they represent classical probability distributions rather than quantum superpositions.

For single-qubit states, the density matrix is also often expressed as a decomposition of the Pauli matrices

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{2.1-8}$$

weighted by $\boldsymbol{r} = (r_x, r_y, r_z)$ such that using $\sigma = (\sigma_x, \sigma_y, \sigma_z)$

$$\rho = \frac{1}{2}(I + \boldsymbol{r} \cdot \sigma) \ . \tag{2.1-9}$$

Here, $\boldsymbol{r}$ is called the Bloch vector, as it indicates the point at which the state would reside in a Bloch representation using Cartesian coordinates. Notably, pure states are located on the sphere ($|\boldsymbol{r}| = 1$), while mixed states live within ($|\boldsymbol{r}| \leq 1$). For example, $\boldsymbol{r} = (0, 0, 1)$ corresponds to the pure state $|1\rangle$ and $\boldsymbol{r} = (0, 0, 0)$ represents the fully mixed state $1/2\,(|0\rangle\langle 0| + |1\rangle\langle 1|)$.

2.1.2.4 Reduced density operator

It is when describing a subsystem of a larger quantum system that the density operator formalism demonstrates its versatility best. Consider a composite quantum system consisting of two subsystems, A and B, whose state is described by a density operator ρ_{AB}.

If we are only interested in the state of subsystem A, we can obtain its *reduced density operator* ρ_A by performing the *partial trace* over the degrees of freedom of subsystem B. The reduced density operator ρ_A is defined as

$$\rho_A = \mathrm{Tr}_B(\rho_{AB})\,, \tag{2.1-10}$$

where Tr_B denotes the partial trace over subsystem B. The reduced density operator ρ_A contains only all the information accessible from measurements performed on the subsystems alone.

Mathematically, this "tracing out" of a subsystem is a generalization of the trace operation and is particularly useful when dealing with entangled states or when a system interacts with an environment, as it allows one to describe the state of a subsystem without requiring knowledge of the entire system. The partial trace is defined in terms of orthonormal basis states $|i\rangle_B$ of the subsystem B being traced out

$$\mathrm{Tr}_B(\rho_{AB}) = \sum_i \langle i|_B \, \rho_{AB} \, |i\rangle_B \ . \tag{2.1-11}$$

To illustrate the partial trace and its effect on entangled systems, consider the $|\Psi^+\rangle$ Bell state from Eq. 2.1-4. Its density matrix is

$$\rho_{AB}^{\Psi^+} = |\Psi^+\rangle_{AB} \langle\Psi^+|_{BA} = \frac{1}{2} \left(|00\rangle \langle 00| + |00\rangle \langle 11| + |11\rangle \langle 00| + |11\rangle \langle 11|\right) \ . \tag{2.1-12}$$

Tracing over subsystem B yields

$$\rho_A = \frac{1}{2} \left(|0\rangle \langle 0| + |1\rangle \langle 1|\right) \ ,$$

which shows that subsystem A is in a fully mixed state, provided no side information about B is available. In photonic implementations, this can easily occur when a photon is lost or its state is measured, but the outcome is not shared. This result also indicates why faster-than-light signaling is not possible, even using entanglement. To signal a bit without error, one must not only project one of the photons but also communicate the measurement basis and outcome.

2.1.2.5 Measurements and quantum state tomography

The previous sections have already used the term *measurement* in a way that is familiar to most, but in quantum mechanics, the act of measuring plays a central role and can have unexpected outcomes and consequences. Unlike classical measurements, which reveal the exact state of a system, quantum measurements are probabilistic and fundamentally alter the state being measured. The simplest, and yet important, type of quantum measurement is a *projective measurement*, which is described by a set of orthonormal projection operators $\{P_m = |\varphi_m\rangle \langle\varphi_m|\}$ that satisfy the completeness relation $\sum_m P_m = \mathbb{I}$. The observable operator is then defined as

$$M = \sum_m \lambda_m \, P_m \, , \tag{2.1-13}$$

where λ_m are the corresponding eigenvalues, and the probability to find any state $|\psi\rangle$ in state $|\varphi_m\rangle$ is given by Born's Rule

$$\langle\psi| \, P_m \, |\psi\rangle = |\langle\psi|\varphi_m\rangle|^2 \equiv p(\lambda_m) \, . \tag{2.1-14}$$

The outcome of a projective measurement must be one of the eigenvalues of the measurement operator and occurs with probability $p(\lambda_m)$. Immediately after the measurement, the quantum system is in the corresponding eigenstate.

Since the outcome of a quantum measurement is probabilistic, it is useful to define the expectation value of an observable. If a measurement is repeated on many copies of the same state, the outcome will tend to this expectation value

$$E(M) = \langle M \rangle = \sum_m \lambda_m \, p(\lambda_m) \, . \tag{2.1-15}$$

It can be understood as a weighted average of possible outcomes and is also linked to the standard deviation (that is, the variance) of the observable

$$[\Delta M]^2 = \langle (M - \langle M \rangle)^2 \rangle = \langle M^2 \rangle - \langle M \rangle^2 \, . \tag{2.1-16}$$

In density operator formalism, the expectation value is simply calculated as

$$\langle M \rangle_\rho = \mathrm{Tr} \left[M \rho \right] \, . \tag{2.1-17}$$

Since quantum states cannot be directly observed and every projection measurement destroys the initially prepared state, the complete density matrix ρ can only be reconstructed from a set of measurements performed on identical and independently prepared copies of the same state (*i.i.d. assumption*). This technique is called *quantum state tomography* [110] and is fundamental to quantum information science.

To estimate a single-qubit state through measurements, we need to learn the expectation value of the orthogonal observables. Using the density operator description of Eq. 2.1-9, that translates to retrieving components of the Bloch vector $\boldsymbol{r}$ by appropriate projective measurements, where the observables are the Pauli matrices

$$\langle \sigma_x \rangle_\rho = \mathrm{Tr} \left[\sigma_x \rho \right] = r_x, \quad \langle \sigma_y \rangle_\rho = \mathrm{Tr} \left[\sigma_y \rho \right] = r_y, \quad \langle \sigma_z \rangle_\rho = \mathrm{Tr} \left[\sigma_z \rho \right] = r_z. \tag{2.1-18}$$

For larger systems, the measurement complexity scales exponentially ($4^n - 1$, where n is the number of qubits), and full state tomography becomes impractical even for a few-qubit state. Alternative approaches use so-called *witness* measures, which require many fewer copies but cannot reconstruct the density matrix. Instead, they quantify how close the experimental state is to a target state, typically expressed via the *fidelity*

$$\mathcal{F}(\rho_{\mathrm{exp}}, \rho_{\mathrm{target}}) = \left(\mathrm{Tr} \sqrt{\rho_{\mathrm{exp}}^{1/2} \rho_{\mathrm{target}} \rho_{\mathrm{exp}}^{1/2}} \right)^2 \, . \tag{2.1-19}$$

2.1.3 Quantum cryptography using single photons

The basic ideas behind quantum cryptography are straightforward. According to the measurement postulate, a quantum state can only be measured once, as the measurement collapses the wave function by projecting the output state onto one

of its eigenstates. In addition, the no-cloning theorem guarantees that—without changing the state of the original qubit—no one can make a perfect copy to effectively measure the state multiple times [12, 111]. And finally, the encoding method known as *conjugate encoding*, which relies solely on the validity of quantum mechanics to protect the encoded information. The final ingredient for an *i.t.-secure* communication relying only on the validity of quantum mechanics is an encoding method known as *conjugate encoding*. Introduced by Wiesner in 1983 [3], it is often regarded as the starting point for the field of quantum cryptography. The concept will be explained in more detail in Sec. 2.1.3.2, but in short, the information is randomly encrypted in one of two conjugate encoding bases, and since the eigenvectors of conjugate bases have a non-zero overlap, they cannot be perfectly discriminated (Sec. 2.1.4.3). Consequently, the information could only be retrieved perfectly if the encoding basis were known before the measurement.

When transmitting quantum information over long distances, photonic qubits are the obvious choice for several reasons. First, photons travel at the speed of light; second, they interact very little with their environment, which keeps decoherence at bay; and third, there is already extensive existing infrastructure and rich experience with classical optical communication. However, to exchange quantum information securely using photons, one must ensure that there is only a single copy of each qubit; that is, the information must be encoded into a single photon. Therefore, this section will briefly review different states of light and their quantum mechanical description (Sec. 2.1.3.1) before detailing how information may be encoded onto a photon (Sec. 2.1.3.2).

2.1.3.1 States of light

In quantum mechanics, the electromagnetic field is quantized through an auxiliary cavity of large volume V that hosts only discrete modes, which may be populated by any number of excitations. Mathematically, these excitations can be described by the bosonic operators $a^\dagger_{\boldsymbol{k}\mu}$ and $a_{\boldsymbol{k}\mu}$ that create or annihilate a photon in mode $\boldsymbol{k}\mu$, respectively ($\boldsymbol{k}$ denotes the wave vector and μ the polarization.

The derivation of the quantized electric field found in any quantum optics textbook (for example, see [112]) yields

$$\boldsymbol{E}(\boldsymbol{r},t) = \sum_{\boldsymbol{k},\mu} \boldsymbol{e}_{\boldsymbol{k},\mu} \sqrt{\frac{\omega_{\boldsymbol{k}}\hbar}{2\varepsilon_0 V}} \left(a_{\boldsymbol{k}\mu} \mathrm{e}^{-i\varphi_{\boldsymbol{k}}(\boldsymbol{r},t)} + a^\dagger_{\boldsymbol{k}\mu} \mathrm{e}^{i\varphi_{\boldsymbol{k}}(\boldsymbol{r},t)} \right) , \tag{2.1-20}$$

where $\boldsymbol{e}_{\boldsymbol{k},\mu}$ is the unit vector, $\omega_{\boldsymbol{k}}$ the angular frequency, and $\varphi_{\boldsymbol{k}}$ the phase of each mode. From now on, the description will be limited to a single-mode field, and for comparing different types of light, the normalization prefactor can also be dropped.

Fock states are eigenstates of the photon number operator $\hat{n} = \hat{a}^\dagger \hat{a}$, so that a Fock state[3] $|n\rangle$ contains a well-defined number of photons, and is mathematically expressed as:

$$|n\rangle = \frac{(\hat{a}^\dagger)^n}{\sqrt{n!}} |0\rangle \,, \tag{2.1-21}$$

where $|0\rangle$ is the vacuum state, and n is a non-negative integer.

Since every state has exactly n photons, the probability p_n to find n photons, the mean (that is, expectation value) of the number operator, and its variance are simply

$$p_n = \delta_{n,\bar{n}} \tag{2.1-22}$$

$$\bar{n} = \langle \hat{n} \rangle = \langle n | \, \hat{n} \, | n \rangle = n \tag{2.1-23}$$

$$[\Delta \hat{n}]^2 = \langle \hat{n}^2 \rangle - \langle \hat{n} \rangle^2 = 0 \,. \tag{2.1-24}$$

This state of light cannot be captured using the classical description of electromagnetism and is therefore often referred to as nonclassical or quantum light.

Coherent states were introduced by Glauber in 1963 [113] (they are also known as Glauber states) and are strongly related to his work on quantum coherence and correlation functions [114, 115]. They exhibit perfect phase-coherence and are fully described by a single parameter α. These states are the quantum states of light that most closely resemble classical electromagnetic waves. They are eigenstates of the annihilation operator $\hat{a}$

$$\hat{a} \, |\alpha\rangle = \alpha \, |\alpha\rangle \ . \tag{2.1-25}$$

Expressed in photon number basis, the coherent state reads as

$$|\alpha\rangle = \mathrm{e}^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle \ . \tag{2.1-26}$$

The photon number distribution $\{p_n\}$, mean photon number and variance

$$p_n = \frac{|\alpha|^{2n}}{n!} \mathrm{e}^{-|\alpha|^2} \tag{2.1-27}$$

$$\bar{n} = \langle \alpha | \, \hat{n} \, | \alpha \rangle = |\alpha|^2 \tag{2.1-28}$$

$$[\Delta \hat{n}]^2 = \langle \hat{n}^2 \rangle - \langle \hat{n} \rangle^2 = |\alpha|^2 \tag{2.1-29}$$

follow a Poisson distribution. A light source producing these photonic states is therefore called a PDS and the most prominent example is a single-mode laser

[3] In the literature, Fock state and photon number state are sometimes used interchangeably. However, while Fock states refer to a single-mode state with exactly n photons, the concept of a photon number state is broader, as it can also describe a coherent superposition of Fock states, including the vacuum state. In that sense, no real device can emit pure Fock states, and a Fock state does not have a phase by itself. In contrast, photon number states may have a well-defined phase between different number components, which might be exploited by an eavesdropper as shown later (Sec 2.1.4.3).

above its lasing threshold. Similarly, sources generating light with a lower variance (such as Fock states) are referred to as sub-Poissonian sources.

Thermal states describe light in thermal equilibrium, such as blackbody radiation or light emitted by a thermal source (e.g., a light bulb). They are statistical mixtures of Fock states, with a photon number distribution given by the Bose-Einstein distribution

$$p_n = \frac{\bar{n}^n}{(1+\bar{n})^{n+1}}\,, \tag{2.1-30}$$

where the mean photon number $\bar{n}$ depends on the temperature of the thermal source and the frequency of the light. Since the light is in a mixed state, it cannot be written in the state vector formalism, but only using the density operator

$$\rho_{\text{thermal}} = \sum_{n=0}^{\infty} p_n \left|n\right\rangle \left\langle n\right| \,. \tag{2.1-31}$$

Their variance is larger than the mean photon number

$$[\Delta\hat{n}]^2 = \bar{n}(\bar{n}+1) \tag{2.1-32}$$

which makes them super-Poissonian distributed.

Considerations for quantum cryptography

Recalling that discrete-variable (DV) quantum cryptography requires exactly one single photon to encode every qubit, a Fock-1 state is clearly the ideal choice. This simple observation motivates the pursuit of an ideal SPS[4] and many physical systems have been proposed and studied over the past years [68, 118].

However, so far, there exists no source generating purely single-photon states, and one must carefully evaluate which photon-number distribution best suits the application. For this purpose, the photon-number distribution $\{p_n\}$ is often divided into three contributions: single-photon target states; vacuum states occurring with p_0 that by themselves cannot transmit information; and multiphoton states that are assumed to fully reveal the encoded information and are generated with $p_{\text{m}} = \sum_{n>1} p_n$. Although at first glance coherent states appear to be ill-suited for the requirements of quantum cryptography, choosing α such that $\bar{n} \ll 1$ can provide a multiphoton probability that is sufficiently small to be secure. Due to their low mean photon number, these states are known as WCSs, and they can be implemented simply by attenuating a laser pulse to the desired power level.

Later, in Sec. 2.1.4, security properties are introduced, and we will analyze which light source may prove advantageous under which circumstances. In addition, Chapter 4 presents a comparison between a realistic SPS and PDS based on the example protocol of QKD.

[4]Naturally, there are many more fields in physics that strive for a SPS for different reasons. Most closely related to the topic of this thesis are applications in quantum computing [55, 59] or metrology [116], as well as fundamental sciences [116, 117].

2.1.3.2 Encoding information onto photons

Once the single-photon state has been prepared, the information must be encoded in a way that ensures security. A powerful method that underlies most quantum cryptography protocols up to today was introduced by Wiesner in 1983 and is known as conjugate coding [3]. It harnesses the uncertainty principle of quantum mechanics, which states that two conjugate variables of a quantum state, such as momentum and position, energy and time, or phase and photon number, can only be measured with limited precision. Mathematically speaking, the product of the standard deviation of both variables is bounded by some constant.

Working with DV encoding of a qubit state, the Pauli operators form a set of conjugate variables. A qubit may be implemented in various degrees of freedom of the single photon; most frequently used are its polarization, time, and sometimes spatial modes. What follows holds for any encoding choice but for simplicity the notation of polarization encoding, which is used in the experiment presented in Chapter 6, is adopted.

The horizontally and vertically linearly polarized states $|H\rangle$ and $|V\rangle$ are mapped to $|0\rangle$ and $|1\rangle$, respectively, and form the *rectilinear* basis. The basis spanned by the eigenstates of σ_x

$$|D\rangle = \frac{1}{\sqrt{2}}\left(|H\rangle + |V\rangle\right) \qquad |A\rangle = \frac{1}{\sqrt{2}}\left(|H\rangle - |V\rangle\right) \tag{2.1-33}$$

is called *diagonal* and the basis of σ_y

$$|L\rangle = \frac{1}{\sqrt{2}}\left(|H\rangle + i\,|V\rangle\right) \qquad |R\rangle = \frac{1}{\sqrt{2}}\left(|H\rangle - i\,|V\rangle\right) \tag{2.1-34}$$

circular. To implement conjugate coding, the parties agree on two of these bases. Using a polarization qubit, the set formed by the four eigenstates of two conjugate bases is in reference to the first proposal of polarization-based cryptography [30] (see also Sec. 2.1.4.1) known as BB84 states.

During the protocol, the sender randomly picks one basis and generates one of the eigenstates that corresponds to the logical bit they want to encode. For example, using σ_z and σ_y as conjugate bases, each logic 0 is either encoded as $|H\rangle$ or $|L\rangle$, and each logic 1 as $|V\rangle$ or $|R\rangle$. Since $|H\rangle$ and $|R\rangle$ (likewise, $|V\rangle$ and $|L\rangle$) are non-orthogonal, it is impossible to distinguish them perfectly. If a logic 0 was encoded using the circular basis but is projected onto the rectilinear basis, the outcome is with equal probability $|H\rangle$ or $|V\rangle$. Hence, one cannot reliably retrieve the encoded logic bit without prior knowledge of the encoding basis.

From a practical point of view, there are generally two techniques for encoding information in the polarization of a photon. Either the sender first prepares a stream of single photons with identical polarization and then applies some unitary to adjust the polarization (*prepare & measure scheme*), or they generate a polarization Bell state and measure one of the two photons themselves (*entanglement-based scheme*). In the latter case, the remaining photon is in a perfectly mixed state unless the

basis and outcome of the projection of the first photon are known (Sec. 2.1.2.4). Although the strategies may appear quite different at first glance, it can be shown that they reduce to an equivalent security analysis [119]. Experimentally, there may be several reasons to favor one approach over the other. Particularly relevant are the rate at which one can generate the polarization-encoded photons, the error probability of the encoding, and the probability of inadvertently leaking information through photon-number coherence (PNC) or a high multiphoton component (see Sec. 2.1.4.3).

The experiment in Chapter 6 employs the *prepare & measure scheme* (also known as *active encoding*). Suppose the generated stream of N single photons is in state $|H\rangle^{\otimes N}$ and the rectilinear and circular bases are chosen as a conjugate pair. Up to a global phase, the unitary of a general linear retarder with its principal axis

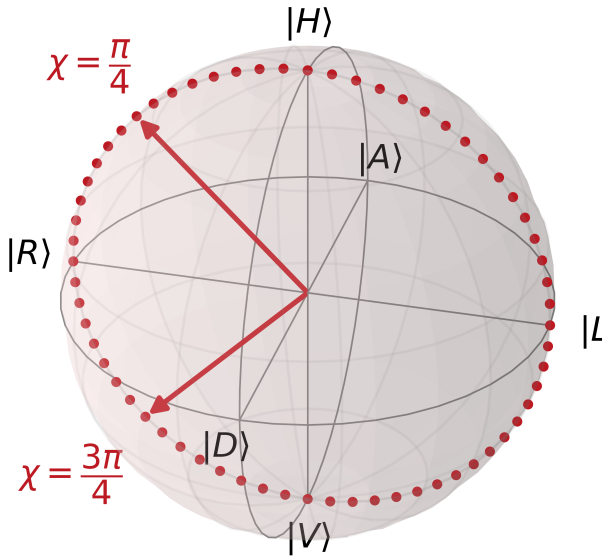


Figure 2.2: **Bloch sphere representation of linear retarder states.** Using a linear retarder that is tilted by $\alpha = \pi/4$ with respect to the input polarization axis of $|H\rangle$, all states on the great circle defined by $\{|H\rangle, |R\rangle\}$ can be encoded by varying the retardance χ according to Eq. 2.1-36. In particular, the set of BB84 states $\{|H\rangle, |R\rangle, |V\rangle, |L\rangle\}$ may be generated with $\chi = \{0, \pi/2, \pi, 3\pi/2\}$.

rotated by α with respect to the polarization axis of the incoming light is

$$U_{\text{ret}}(\chi, \alpha) = \begin{bmatrix} \cos^2(\alpha) + e^{i\chi}\sin^2(\alpha) & (1 - e^{i\chi})\cos(\alpha)\sin(\alpha) \\ (1 - e^{i\chi})\cos(\alpha)\sin(\alpha) & \cos^2(\alpha) + e^{i\chi}\sin^2(\alpha) \end{bmatrix}, \tag{2.1-35}$$

where the retardance is given by χ. Fixing the tilt angle to $\alpha = \pi/4$, the unitary simplifies to

$$U_{\text{ret}}(\chi, \alpha = \pi/4) = U_{\text{enc}}(\chi) = \frac{1}{2}\begin{bmatrix} 1 + e^{i\chi} & 1 - e^{i\chi} \\ 1 - e^{i\chi} & 1 + e^{i\chi} \end{bmatrix}, \tag{2.1-36}$$

which acts like a quarter wave plate for $\chi = \pi/2$ and like a half wave plate for $\chi = \pi$. It is easy to see that applying $U_{\text{enc}}(\chi)$ to $|H\rangle$, all states on the great circle of the Bloch sphere shown in Fig. 2.2 can be encoded. Specifically, for $\chi \in \{0, \pi/2, \pi, 3\pi/2\}$ the four basis states of the rectilinear and circular basis $\{|H\rangle, |R\rangle, |V\rangle, |L\rangle\}$ can be generated.

The unitary $U_{\text{ret}}(\chi, \alpha)$ can be implemented by applying a strong electric field to an anisotropic crystal, resulting in anisotropic electro-optic modulation of the refractive index [120]. Such a setup is known as Pockels Cell and allows to control the polarization of each successive photon individually (for more details, see methods of Chapter 6).

2.1.4 Quantum cryptographic primitives

While most researchers working in the field of quantum technologies will be familiar with QKD because of its pioneering role, it is just one primitive of quantum cryptography. As outlined in Sec. 2.1.1, cryptographic primitives form the backbone of each cryptosystem and are often described as building blocks. Clearly, to remain with this linguistic picture, more than just one fundamental block is necessary to achieve variety or design more complex structures. A detailed description of all existing primitives, their subtleties, and variants would far exceed the scope of these Preliminaries, so only a brief classification and overview will be given. For an extensive and recent review of quantum cryptography that goes beyond QKD see [121].

One of the most important and distinctive characteristics of a primitive is whether its legitimate parties trust one another and collaborate, or not. Naturally, they are called either *trustful* or *mistrustful*. The first set includes QKD [102, 103], unforgeable tokens [3], unclonable encryption [122, 123], position verification, covert communication [124, 125], fingerprinting [126], and data locking, whereas the second covers digital signatures [127], bit commitment [104, 128], oblivious transfer [127], strong and weak coin flipping [129, 130], and one-time programs [131]. Although they are all treated in their own right, some of the above are not fundamental blocks. For example, using unbiased weak coin flipping allows one to perform biased strong coin flipping, oblivious transfer, or bit commitment. Given a perfect bit commitment, one can construct an oblivious transfer [132]. Recently, it has also been proven that, in the quantum setting, one-way functions enable oblivious transfer [133]. Analyzing all these relations, a hierarchy of primitives may be established, and interestingly, in the quantum world, this hierarchy differs significantly from the one known classically (although not all connections have been proven for quantum primitives).

When designing a quantum cryptographic protocol, its security guarantees must be analyzed in the context of a clearly defined model. In the first stage, experimental imperfections are usually neglected to understand if a protocol can be secure under ideal conditions. While the adversary is typically assumed all-

powerful[5], some primitives require additional assumptions weakening the adversary model to provide *i.t.*-security (see Sec. 2.1.4.2, also [121] and references therein).

Once the protocol is provably secure under ideal conditions, one must account for inevitable imperfections of the experiment, and this step generally restricts the parameter space for which the protocols security is guaranteed. In photonic demonstrations, these limitations primarily arise from imperfect qubit preparation, multiphoton contributions, photon loss in the quantum channel, or detector dark counts. To evaluate how a specific imperfection impairs the protocol's performance, the optimal attack strategy of the adversary is analyzed, which may include measurements that are difficult to realize but are theoretically possible, such as photon-number splitting (PNS) [39], arbitrary unambiguous state discrimination (USD), and quantum non-demolition measurements. Furthermore, a dishonest party always has the power to "replace" the lossy public quantum channel with a lossless one.

Sec. 2.1.4.1 revisits QKD and presents an overview of the protocol but also introduces some general concepts to analyze the security of a primitive. Sec. 2.1.4.2 then illustrates for QBC how to establish a protocol when the parties mutually mistrust one another, and how *i.t.*-security can still be achieved in such a setting if additional physical assumptions are included. Finally, Sec. 2.1.4.3 turns the focus to vulnerabilities in photonic implementations, how a practical attack may extract more information, and which countermeasures can prevent exploits of experimental imperfections.

2.1.4.1 Quantum key distribution

QKD allows two parties who trust one another, commonly referred to as Alice (sender) and Bob (receiver), to generate a shared symmetric key pair with *i.t.*-security using quantum resources and an authenticated classical channel. The latter ensures that a given message actually originates from the legitimate sender and is required during classical postprocessing to establish a secure key (SK) as shown below. This key may later be used to encrypt a message, provide a digital signature, etc.

We will first present the protocol in a condensed and simple form to convey the main concept before proving further details and insights about the individual steps involved. For an in-depth review of state-of-the-art QKD including security analysis and experimental imperfections, see [35]. A formal security proof is given, for example, in [119, 134].

The most common variant of this protocol uses a set of four quantum states chosen from the basis states of two pre-agreed conjugate bases; these states are known as BB84 states. In a prepare & measure implementation using polarization

[5]In this model, the adversary has unlimited classical and quantum resources (including a perfect quantum memory) and is only restricted by the laws of quantum mechanics. However, it is generally required that they cannot access the lab of the legitimate parties (*secure lab assumption*).

qubits, Alice generates a stream of single photons, randomly prepares their polarization state (keeping a local record of each encoding), and sends them via a public quantum channel to Bob. Projecting the states randomly to either of the two bases, Bob measures the polarization and arrival time of each photon, concluding the quantum phase of the protocol.

Next, they compare their preparation and measurement bases for each qubit over the authenticated but public classical channel. This step is known as *basis reconciliation* or *key sifting*, and statistically, half of the events must be discarded due to basis mismatch. In the absence of any photon loss, errors, or tampering from an adversary, Alice and Bob would now hold an identical, random bit string (raw key). But to verify that the key is indeed secret, they must sacrifice a randomly selected subset and publicly compare it to look for unexpected irregularities. If the estimated error rate exceeds a certain threshold, Alice and Bob conclude that they have been attacked; up to this point, no secret data has been leaked, and they simply discard the raw key and repeat the process. If the error estimate satisfies their requirements, they proceed with *error correction*, where Alice and Bob interactively identify and correct bit flips. Since this step leaks information about the raw key to the public, they perform *privacy amplification* to remove any side information an adversary Eve might have gained. Finally, Alice and Bob hold a symmetric SK only known to them.

Intuitively, the security of this protocol is based on conjugate coding (see 2.1.3.2) that prevents Eve from measuring a qubit in both bases with high precision simultaneously and on the no-cloning theorem, which guarantees that she cannot make a perfect copy beforehand to measure each in one of the two bases. Any attempt to learn the raw key by measuring some of the qubits introduces disturbances that Bob can detect[6].

To analyze the protocol more formally, we must define security parameters. The key obtained through QKD must be *correct*, i.e. Alice and Bob hold an identical bit string, and it must be *secret*, i.e. uniformly distributed to anyone but them. Taking small deviations from the ideal case into account, ε_{cor} is introduced and the protocol is said to be correct if

$$P(k_A \neq k_B) \leq \varepsilon_{\text{cor}}\,, \tag{2.1-37}$$

where $P(k_A, k_B)$ denotes the probability distribution that a classical bit value of Alice k_A matches Bob's k_B. To measure how secret a key is, the joint state is

[6] If Eve simply measures all the photons, Alice and Bob cannot establish a raw key and the measurement outcomes are useless for Eve. Hence, after measuring the intercepted photon (and destroying its state), Eve must generate a new photon which is ideally in the same quantum state (*intercept & resend attack*). However, since Eve cannot learn the full density matrix from a single-shot measurement, she can only try to approximate the original state and thus minimize the errors her state will cause in Alice's and Bob's correlations. In an ideal experiment, Alice and Bob can always detect tampering, but Eve might exploit the imperfections present in practical implementations as discussed in Sec. 2.1.4.3.

described as a classical-classical-quantum state shared between Alice, Bob and Eve

$$\rho_{ABE}^{\text{ideal}} = \sum_{k_A,k_B} P(k_A,k_B)\,|k\rangle_A\langle k| \otimes |k\rangle_B\langle k| \otimes \rho_E^{(k_A,k_B)}\,. \tag{2.1-38}$$

Ideally, all bits of Alice and Bob are identical, $k_A = k_B = k$, and uncorrelated with Eve's subsystem

$$\rho_{ABE}^{\text{ideal}} = 2^{-m}\sum_{k} |k\rangle_A\langle k| \otimes |k\rangle_B\langle k| \otimes \rho_E\,, \tag{2.1-39}$$

where m is the length of the bit string. Then, the protocol is said to be ε-secure if the final distilled state[7] ρ_{ABE} satisfies [135]

$$\min_{\rho_E} \frac{1}{2}\,(1-p_{\text{abort}})|\rho_{ABE} - \rho_{ABE}^{\text{ideal}}| \leq \varepsilon\,, \tag{2.1-40}$$

where p_{abort} is the probability that the protocol aborts and $|\rho|$ is the trace distance. Eq. 2.1-40 states that the distilled state is ε-close to the ideal state, such that Eve has at most a chance of ε of guessing the key.

Eqs. 2.1-37-2.1-40 indicate that the length of the final SK depends both on the desired security parameter ε and, implicitly, on the experimental performance. Any practical implementation will exhibit some errors stemming from imperfect state preparation, polarization drifts in the quantum channel, or an imperfect detection stage. To understand how much SK (if any) the parties can extract from a given raw key, they must complete a step known as *parameter estimation* just before the error correction. Hence, we now derive an explicit expression that limits the number of secure bits that can be extracted per state in a practical implementation with imperfections.

Describing the probability that Bob detects an i-photon state

$$Y_i = Y_0 + (1-Y_0)(1-(1-\eta_d\eta_t)^i) \tag{2.1-41}$$

where Y_0 is the probability of a dark count click, η_d Bob's detector efficiency, and η_t the transmission of the quantum channel, the gain is defined as joint probability of Alice generating a phase-randomized i-photon state and Bob detecting it

$$Q_i = p_i Y_i\,. \tag{2.1-42}$$

We also define the error probability of an i-photon state

$$e_i = \frac{e_0 Y_0 + e_{\text{d}}\left[1-(1-\eta_{\text{d}}\eta_{\text{ch}})^i\right]}{Y_i}\,, \tag{2.1-43}$$

[7]The final state is sometimes referred to as *distilled*, since most security proof frameworks use entanglement distillation to "purify" the state to remove any knowledge Eve might have about the key. This step is closely related to privacy amplification mentioned above.

where the parameter e_d characterizes the detection error probability, dependent on the optical alignment of the entire system, and e_0 is the error rate of the background, which, if we assume it to be random, is $e_0 = \frac{1}{2}$. Since Alice does not know when she might accidentally send a multiphoton state, and Bob also lacks photon-number resolving detectors, experimentally, they can only measure the total information gain and qubit error rate

$$Q_\mathrm{tot} = \sum_i Q_i \tag{2.1-44}$$

$$e_\mathrm{tot} = \frac{1}{Q_\mathrm{tot}} \sum_i e_i Q_i \,. \tag{2.1-45}$$

Hence, their best secure strategy is to combine these measurements with the emission statistics of the SPS $\{p_i\}$ measured by Alice to estimate how much information might have been leaked. Without more detailed knowledge, they must attribute all losses and errors to the single-photon states while assuming that multiphoton states propagate perfectly through the untrusted channel, thus fully leaking the corresponding bit values (see Sec. 2.1.4.3). As a result, they can only bound the single-photon parameters

$$Q_1 \geqslant Q_\mathrm{tot} - p_\mathrm{m} - Y_0 p_0 \tag{2.1-46}$$

$$e_1 \leqslant \frac{e_\mathrm{tot} Q_\mathrm{tot} - \frac{1}{2} Y_0 p_0}{Q_1} \,. \tag{2.1-47}$$

In the limit of an asymptotically long bit string, the GLLP formula then lower bounds how many secure bits can be extracted per state [136]

$$R \geqslant \frac{1}{2} \left[Q_1 \left(1 - H_2(e_1)\right) - f(e_\mathrm{tot}) Q_\mathrm{tot} H_2(e_\mathrm{tot}) \right] . \tag{2.1-48}$$

In this formula, the term $f(e_\mathrm{tot}) Q_\mathrm{tot} H_2(e_\mathrm{tot})$ accounts for the cost of error correction and privacy amplification—$f(e_\mathrm{tot})$ being the error correcting code inefficiency and $H_2(e)$ the binary Shannon entropy—and $Q_1(1 - H_2(e_1))$ states that only the error-free single photon states contribute positively to the secure key generation. Here, common error correction codes include *Cascade* featuring a high correction efficiency of $f = 1.1 - 1.2$ [137], or *low-density parity checks* [138] requiring only one round of classical communication. Privacy amplification is performed using universal hashing (for more detail, see [35, 139]).

Here, Eq. 2.1-46 provides insights into two of the reasons when QKD might fail to generate a secure key: either the multiphoton component exceeds the total gain $p_\mathrm{m} > Q_\mathrm{tot}$, or the gain is saturated by dark count events, $Y_0 p_0 \approx Q_\mathrm{tot}$. Naturally, both effects are cumulative, and the achievable SK rates are especially impaired in the high loss regime where Q_tot is small. Using PDS, the mean photon number (thus, the population distribution $\{p_i\}$) must be optimized for any given set of experimental parameters and the channel loss. Another cause for QKD to fail

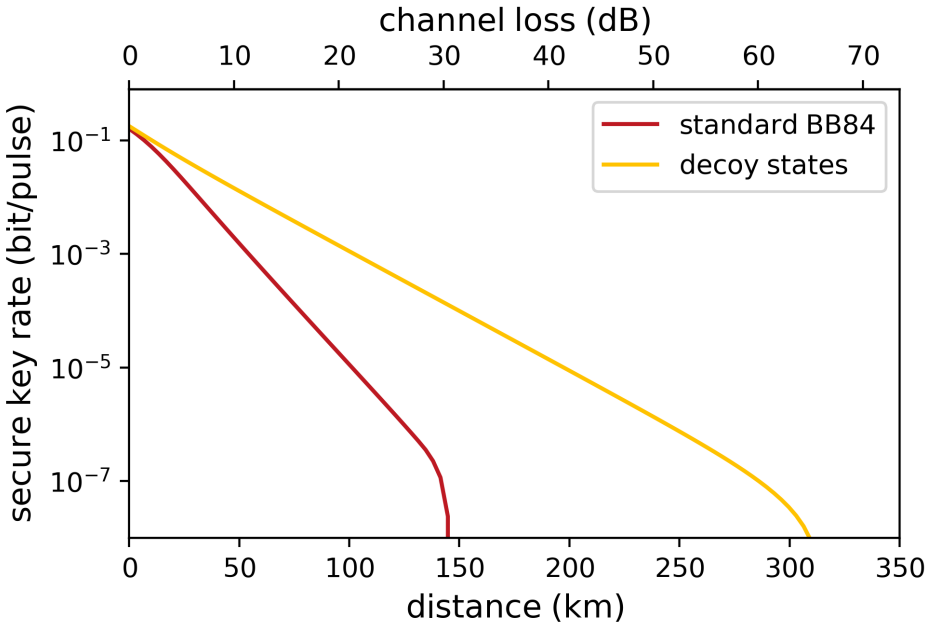


Figure 2.3: **Secure key rate for QKD over channel distance.** Secure key rate calculated as Eq. 2.1-48 for the standard BB84 protocol and using decoy states in the asymptotic limit. For each distance, the mean photon number was optimized to maximize the SK rate. The distance scale assumes a standard telecom single-mode fiber (SMF) with 0.21 dB/km. The other parameters are: alignment error $e_d = 0.02$, detector efficiency $\eta_d = 0.85$, dark count probability $Y_0 = 10^{-8}$ and error correction inefficiency $f = 1.2$

is illustrated by Eq. 2.1-47: if the measured qubit error rate e_{tot} (also known as QBER) is too large R becomes negative. However, it turns out that QKD is remarkably robust against noise; it can tolerate an error probability of up to 11 % [119] and still provide a SK.

Returning to the multiphoton component, it was soon realized that when both legitimate parties cooperate, one can estimate the gain Q_1 directly and avoid the pessimistic assumption that multiphoton states are never lost in the channel. By sometimes sending *decoy states* [40, 41], Alice and Bob may infer the specific loss and error behavior for each of the photon populations individually, even in the presence of an adversary. This improvement of the protocol has significantly boosted performance for PDS as shown in Fig. 2.3, and thus has quickly become the standard version of QKD with WCS.

Over the past decades, even more advanced forms of QKD have been developed. While some aim to provide stronger security guarantees such as measurement device independent QKD [35] or fully device independent QKD [36], others extend the maximum communication distance beyond the PLOB bound [140] using repeater-

based schemes [141] or so-called twin-field QKD [102, 142].

2.1.4.2 Quantum bit commitment

In a nutshell, a bit commitment forces a sender to make a binary choice at one point in time, and allows the receiver at a later point to learn the bit and also verify that the sender did not change their choice after committing. In an often-cited analogy, one party puts a bit $b \in \{0, 1\}$ in a safe and hands it over to a second party but keeps the combination for the safe to herself. This constitutes the *commitment* phase. Only later, in the *reveal* phase, the sender discloses which bit b they had chosen and also provides the combination for the safe, allowing the receiver to verify that the disclosed and committed bits match. The protocol offers two security properties: it is called *binding* if a dishonest sender cannot change the bit b after committing without being detected, and it is *hiding* if a dishonest recipient cannot read the bit b before receiving the combination to decrypt the commitment.

Such a mistrustful primitive may be used to enable secure voting, multi-party computations, or smart contracts [121]. In fact, it was first proposed with classical resources for contract-signing [143, 144]. However, like for key generation, classical implementations can only provide computational security for at least one of the primitive's security properties. Although there were hopes for a perfect, *i.t.*-secure version using quantum resources [128], it was proven impossible not long after [104, 105]. Nevertheless, interest persisted[8] and two alternative approaches were pursued to obtain i.t. security: either to accept an imperfect protocol bias (that is, allow for cheating but with a probability of getting caught) [145, 146], or to limit the adversary's capabilities based on a plausible physical assumption. For the latter, a few options were suggested; the most common include no-signaling constraints derived from flat space-time [107, 108], and the noisy-storage model [106][9]. Both approaches are provably secure and were experimentally demonstrated [32, 149–151].

Outlining the basic bit commitment primitive in accordance with the protocol implemented in Chapter 6, we reverse the roles compared to the original protocol; now, the receiver of the quantum states (to which we will refer as the Prover) commits to a classical bit b, and the sender (from hereon, the Verifier) will check the commitment. In the absence of any experimental imperfections, the protocol works as follows.

The Verifier prepares N random BB84-type quantum states and sends them to the Prover, who measures *all* states in σ_z to commit to $b = 0$ or σ_y to commit to $b = 1$ and records the outcomes. In this way, the bit is encoded in the basis

[8] Parallel works compiling the hierarchy of quantum primitives proved that QBC would also implicate other primitives like quantum oblivious transfer [132].

[9] Early versions of this assumption distinguished between *bounded* [147] and *noisy* quantum storage [148]. Later, as cited above, these notions were unified in one model [106], which nowadays—somewhat confusingly—is simply known as the noisy-storage assumption although the memory might also be bounded in size.

choice, and the commitment is secured by the measurement outcomes. After a pre-agreed time Δt, they proceed with the *reveal* phase. The Prover discloses their committed bit and communicates all measurement outcomes to the Verifier. Whenever the encoding basis matches the measurement basis, the Verifier checks if the communicated outcomes agree with the originally prepared states, and they only accept the commitment if all conform. Since the bases are mismatched in half of the cases ($p = 1/2$), they need N quantum states even in the absence of experimental imperfections to ensure that the commitment can be verified with a probability of $1 - (1 - p)^N$.

Since no information is back-reported to the Verifier before the *reveal* phase, they have no chance to learn anything about the Prover's measurement basis (and thus the secret bit) even if they were to send part of an entangled state instead of a single-photon BB84 state. On the other hand, a Prover with a perfect and unbounded quantum memory could clearly perfectly decommit by simply storing all states until the reveal phase and only then measuring them according to the Prover's delayed choice. However, assuming that the memory either fully decoheres within Δt or is not sufficiently large to store all N states, the Verifier can detect a cheating attempt. In such a setting, a bit commitment that is both binding and hiding with *i.t.*-security may be implemented, which is the underlying idea behind the noisy-storage assumption.

Here, the security of the noisy-storage model is only briefly motivated; a formal proof can be found in [106]. We recast the quantum memory as a channel with capacity $\mathcal{C}_\mathcal{N}$ and completely positive trace-preserving map $\mathcal{F}_t(\rho)$ corresponding to the memory's noise. Using a Markovian noise model, the noise strictly increases with time t, and any dishonest Prover would measure at $t = \Delta t$ to limit decoherence. Thus, we must only consider one channel map $\mathcal{F}$, dropping the index. In the most general approach, the Prover applies an optimal encoding to the input state before transmitting it through the channel and optimizes over all positive operator valued measurements $\{D_x\}_{x\in\{0,1\}^{RN}}$ when reading out the information. Here, R describes the storing rate and N indicates the channel uses (i.e., the number of transmitted qubits). The success probability to correctly send a bitstring $X \in \{0,1\}^{RN}$ is then defined as

$$P^{\mathcal{F}}_{succ}(RN) := \max_{\{D_x\}_x,\{\rho_x\}_x} \frac{1}{2^{RN}} \sum_{x\in\{0,1\}^{RN}} \mathrm{Tr}(D_x \mathcal{F}(\rho_x)), \tag{2.1-49}$$

We continue with an explicit channel model that satisfies the strong converse property [152]. The d-dimensional depolarizing channel with purity parameter r, described by

$$\mathcal{N} = r\rho + (1 - r)\frac{\mathbb{I}}{d} \tag{2.1-50}$$

has a capacity of

$$\mathcal{C}_\mathcal{N} = 1 + \frac{1+r}{2}\log\left(\frac{1+r}{2}\right) + \frac{1-r}{2}\log\left(\frac{1-r}{2}\right). \tag{2.1-51}$$

Naturally, for $R < \mathcal{C}_{\mathcal{N}}$, the channel (i.e., the quantum memory) will not introduce any errors. However, if R exceeds the capacity $\mathcal{C}_{\mathcal{N}}$, the probability of correctly transmitting information decays exponentially, according to the strong converse property

$$P_{succ}^{\mathcal{N}^{\otimes N}}(RN) \leq 2^{-N\gamma^{\mathcal{N}}(R)} \text{ where } \gamma^{\mathcal{N}}(R) > 0 \ \forall \ R > \mathcal{C}_{\mathcal{N}}\,. \tag{2.1-52}$$

Here, the upper bound on the success probability follows from the strong converse exponent

$$\gamma^{\mathcal{N}}(R) = \max_{\alpha \geq 1} \frac{\alpha - 1}{\alpha} \left[R - \log\, 2 + \frac{1}{1-\alpha} \log \left(\left(r + \frac{1-r}{2} \right)^{\alpha} + \left(\frac{1-r}{2} \right)^{\alpha} \right) \right]. \tag{2.1-53}$$

Importantly, at this point, we are still working within the framework of an ideal experiment, and additional precautions must be implemented to provide loss-tolerance and noise robustness. For example, if the Verifier were to accidentally send a multiphoton state, a dishonest Prover could perfectly decommit by simply measuring one photon in each basis[10]. Therefore, the Prover could cheat on all single-photon states they can reliably store in their memory, and additionally, on all multiphoton states the Verifier sends. Furthermore, techniques such as the decoy state method, which drastically increase performance in QKD, cannot always be applied when both parties mutually mistrust each other. Consequently, mistrustful primitives are significantly more sensitive to the multiphoton component p_{m} such that they strongly benefit from sub-Poisson distributed sources [45].

To integrate a loss tolerance, the number of quantum states N that are used to commit to a single bit must be increased to ensure that at least $N_d < N$ photons are detected and that the commitment is checkable with probability of $1-(1-p)^{N_d}$. In such a scenario, the Prover would obviously only need to store N_d instead of N states in their quantum memory to perfectly cheat. Similar arguments can be made for securing the protocol against a finite error probability, and all these considerations must be carefully included in the practical security analysis.

2.1.4.3 Practical security

In the previous two sections, various security concerns arising from experimental imperfections were identified, but strategies to exploit these were not discussed in detail. Generally, attacks either target an imperfection in the photonic state the sender generated (*source attacks*), or are directed at the receiver's setup and attempt to control their measurement outcomes (*detector attacks*).

For trusted primitives, the adversary is a third party (usually referred to as Eve) that intercepts (a part of) the signal and forwards a modified photon stream to Bob. Hence, Eve must typically combine both types of attack to maximize her

[10]When they need to provide the measurement outcome to prove their commitment in the reveal phase, they have both outcomes available to pass the check.

chances. However, in mistrustful settings, where Alice and Bob do not cooperate but work against each other, a dishonest sender (receiver) only needs to implement a successful detector (source) attack to cheat perfectly, making the situation even more precarious. Consequently, many primitives of this class are significantly more sensitive to loss and errors, providing a smaller secure region in parameter space.

Since this thesis focuses on the QDs as SPS for cryptography, we are primarily interested in source attacks. However, at the end of the section, a brief overview of detector attacks and other classical side-channel attacks is also provided. A comprehensive overview of strategies can be found in [35, 153, 154].

In many source attacks, the adversary selectively applies an optimal measurement only to the imperfect states, while discarding all perfect states where their successful cheating probability is lower. To compensate for the additional loss introduced by this approach, the adversary replaces the imperfect quantum channel with a lossless channel. If there are enough imperfect states and the loss budget tolerated by the honest sender is sufficiently large, the attack is successful.

The optimal measurement an adversary can implement is closely related to the field of quantum state discrimination. The concept is briefly introduced below, followed by two explicit strategies for leveraging PNC or multiphoton states.

State discrimination attacks

The Helstrom bound derived in 1969 [155] states the maximal success probability

$$P_{\text{corr}} = \frac{1}{2}(1 + \text{Tr}|p_0\rho_0 - p_1\rho_1|) \tag{2.1-54}$$

for which two quantum states ρ_1 and ρ_2 can be discriminated when a fully conclusive outcome is required. Here, the optimal measurement that one must implement corresponds to the projectors onto the positive and negative eigenspace of $p_0\rho_0 - p_1\rho_1$. Choosing quantum states from conjugate bases (such as the BB84 states) maximizes the overlap between states and thus minimizes the discrimination probability.

This type of measurement provides a fully conclusive result for each state ($P_{\text{conc}} = 1$), but the result is only correct in $P_{\text{corr}} < 1$ of the cases, hence introducing errors in the adversary's string. Alternatively, a measurement can be constructed to produce only correct outcomes if accepting that the projection sometimes remains inconclusive ($P_{\text{corr}} = 1$ but $P_{\text{conc}} < 1$). This type of measurement is known as *unambiguous state discrimination* and poses a significantly more serious threat to realistic implementations with finite channel loss. A broad overview of quantum state discrimination is given in [156].

Recalling that the adversary has access to a lossless channel, Eve's attack strategy in a QKD setting is as follows

- if the USD succeeds, she records the result. As she only knows the bit but not the basis, her best strategy is to prepare another photon in a coherent superposition of both states that represent the recorded logic bit and send it to Bob. For example, if the USD outputs bit 0, Eve prepares $\frac{1}{\sqrt{2}}(|H\rangle + |L\rangle)$.

- if the outcome is inconclusive, Eve does nothing and Bob will report the state as photon loss.

Although the bit discrimination is perfect, in a QKD scenario this intercept & resend method introduces an error probability of $\frac{1}{2} - \frac{1}{2\sqrt{2}}$ since Bob might find the photon in the wrong basis. As long as Eve manages to remain within the loss budget provided by the imperfect channel (that is, introducing less loss caused by inclusive measurements than Alice and Bob expect transmission loss caused by the channel), and and error budget predefined by the honest users, the attack remains undetected and is successful.

To construct a useful USD, the adversary exploits some imperfection in the state that can be classified as phase-discrimination or photon-number splitting attacks.

Phase discrimination

Any realistic source produces a non-zero vacuum component, and the phase relation between the vacuum and single-photon state may be exploited if the photon is in a pure number state instead of a mixed number state (see sec. 2.1.2.3).

Suppose the coherent superposition

$$|\psi\rangle = \sqrt{p_0}\,|vac\rangle + \mathrm{e}^{i\phi}\sqrt{p_1}\,|n=1\rangle = \alpha_0\,|vac\rangle + \alpha_1\,|n=1\rangle \tag{2.1-55}$$

is used to encode a polarization qubit in rectilinear and circular basis (α_0 and α_1 were introduced for convenience). Due to the non-zero vacuum component, this state does not live in the two- but three-dimensional Hilbert space spanned by $\{|vac\rangle, |n=1\rangle_H, |n=1\rangle_V\}$ and the conjugate bases states are

$$\begin{aligned} |\tilde{H}\rangle &= \big(\alpha_0, \quad \alpha_1, \quad 0\big) \\ |\tilde{L}\rangle &= \big(\alpha_0, \quad \alpha_1/\sqrt{2}, \quad i\alpha_1/\sqrt{2}\big) \\ |\tilde{V}\rangle &= \big(\alpha_0, \quad 0, \quad \alpha_1\big) \\ |\tilde{R}\rangle &= \big(\alpha_0, \quad \alpha_1/\sqrt{2}, \quad -i\alpha_1/\sqrt{2}\big)\ . \end{aligned} \tag{2.1-56}$$

By finding the states that are orthogonal to both states that encode the logical 0 or logical 1

$$\begin{aligned} |0_\perp\rangle &= \frac{1}{\sqrt{N_0}}\frac{\alpha_0\alpha_1}{\sqrt{2}}\left(-\tfrac{\alpha_1}{\alpha_0}, \quad i+\sqrt{2}, \quad 1\right) \\ |1_\perp\rangle &= \frac{1}{\sqrt{N_1}}\frac{\alpha_0\alpha_1}{\sqrt{2}}\left(i\tfrac{\alpha_1}{\alpha_0}, \quad -i, \quad 1-\sqrt{2}\right)\ , \end{aligned} \tag{2.1-57}$$

where N_i are normalization factors, one can construct the USD as

$$\begin{aligned} M_0 &= \frac{1}{2}\,|1_\perp\rangle\langle 1_\perp| \\ M_1 &= \frac{1}{2}\,|0_\perp\rangle\langle 0_\perp| \\ M_{\mathrm{inc}} &= I - M_0 - M_1\ . \end{aligned} \tag{2.1-58}$$

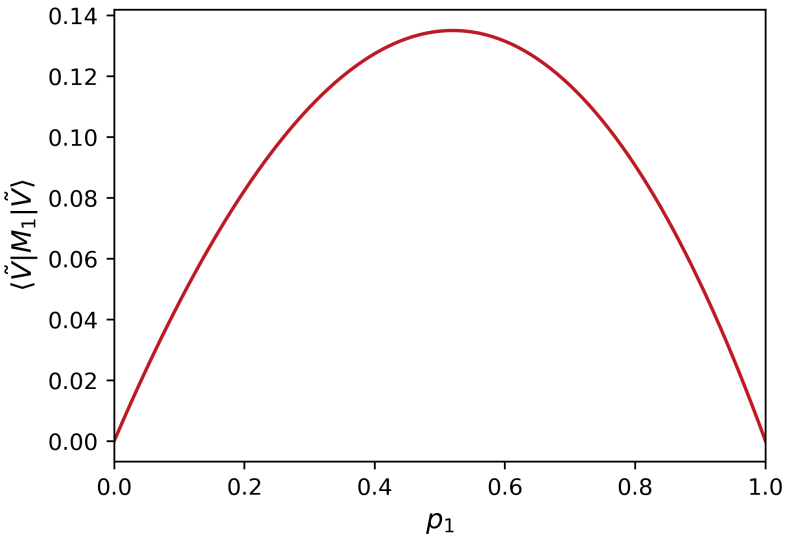


Figure 2.4: **Phase discrimination for a single-photon state.** USD can provide significant knowledge about the encoded state, if the photonic state is not a perfect qubit, but exhibits an imperfection that allows USD in a higher-dimensional Hilbert space. Even in the absence of any multiphotons, phase discrimination is possible between the vacuum and single photon component as long as the state is in a coherent superposition. For equal populations $p_0 = p_1 = 1/2$, the discrimination probability is maximized.

Note that M_1 projects onto $|0_\perp\rangle$, which is orthogonal to both states encoding a logical 0. Hence, if this outcome is found, the adversary knows with certainty that the encoded bit was *not* 0 but must have been 1. Since M_0 and M_1 are fairly parallel (depending on p_0 and p_1), there is a significant amount of inconclusive measurements M_{inc}. The expectation value of M_1

$$\langle\tilde{V}|\, M_1\, |\tilde{V}\rangle = \langle\tilde{R}|\, M_1\, |\tilde{R}\rangle = \frac{(2-\sqrt{2})p_0p_1}{2p_0(2-\sqrt{2})+p_1} \tag{2.1-59}$$

is exemplarily shown in Fig. 2.4 for varying populations. Naturally, the same attack applies to a WCS that has not been phase-randomized [91].

Phase discrimination attacks may be prevented by encoding the polarization onto a mixed state $p_0\,|0\rangle\langle 0| + p_1\,|1\rangle\langle 1|$ instead. If the state generated by the photonic source is pure, phase-randomization effectively transforms the output state into a mixed state. This can be easily verified, since the integration over all

possible phases removes the off-diagonal terms of the density matrix

$$\begin{aligned}\rho &= \frac{1}{2\pi}\int_0^{2\pi} d\phi \ |\psi\rangle\langle\psi| \\ &= \frac{1}{2\pi}\int_0^{2\pi} d\phi \ \sqrt{p_0 p_1}\left(\mathrm{e}^{-i\phi}|0\rangle\langle 1| + \mathrm{e}^{i\phi}|1\rangle\langle 0|\right) \\ &\quad + \frac{1}{2\pi}\int_0^{2\pi} d\phi \ \left(p_0|0\rangle\langle 0| + p_1|1\rangle\langle 1|\right) \\ &= p_0|0\rangle\langle 0| + p_1|1\rangle\langle 1| \ .\end{aligned} \tag{2.1-60}$$

Multiphoton attacks

In the presence of multiphoton states, an adversary uses a quantum non-demolition measurement to infer the photon number without disturbing the encoding degree of freedom and conditions their actions on the measurement outcome. This behavior is known as PNS attack.

Again, in the example of QKD, the best strategy for an eavesdropper is to discard some (or all) single photon states but forward the multiphoton states to Bob after keeping one photon per state for themselves and storing its information in a perfect quantum memory. When Alice and Bob reveal their basis during the key sifting, Eve simply measures all her photons accordingly and obtains full knowledge about these bits of the key. If the state also shows phase coherence, more elaborate methods can extract even more information.

In a bit commitment protocol, where a party commits by measuring the qubit in a particular basis, a dishonest receiver may report all single-photon states as lost and project all multiphoton states in both conjugate bases. In this way, the dishonest receiver can decommit perfectly.

Without appropriate countermeasures, this attack is particularly dangerous in scenarios where the channel loss is high, effectively granting the adversary a large loss budget to selectively "lose" the attack-resistant single-photon states. A common technique to secure a protocol against the PNS attack—other than reducing the multiphoton noise generated by the source—is to include decoy states in the stream of signal states as introduced in Sec. 2.1.4.1. However, this approach applies only to trusted protocols, and for most mistrustful primitives, one can only bound the maximum channel loss for which a secure operation may be guaranteed given a certain p_{m}. This underlines the significance of a highly pure single-photon state in mistrusted scenarios [45].

Side-channel attacks

In addition to methods that directly target the quantum channel, an adversary may attempt to extract information using purely classical means. These types of attacks, referred to as (classical) side-channel attacks, are generally more specific to

a given implementation and are difficult to include directly in the security analysis of a protocol. Hence, each protocol requires that countermeasures for the individual exploits are in place, and failing to do so jeopardizes the security. Two prominent examples are the so-called Trojan horse [157], which targets the sender, and the detector blinding attack [158] aimed at the receiver.

The first applies only to prepare & measure schemes; the strategy is to copy the information classically by exploiting experimental imperfections of Alice's setup in the form of surface reflections. A laser pulse is backpropagated along the quantum channel toward Alice, where it is backreflected from any number of interfaces. If the timing of the Trojan horse pulse is matched so that its reflection passes the encoding stage at the same time as the single-photon state prepared by the legitimate party, the information is imprinted on both photonic states, and the adversary can simply read it out classically from the reflected pulse. For WCS, this attack is easily forestalled by including an optical isolator or attenuator at the end of Alice's trusted setup. The additional losses are simply compensated for by increasing the laser power. Implementations using solid-state SPS must avoid any transmission loss in the preparation stage but can be protected using a timed shutter based on electro-optic modulators (see Ch.6).

The detector-blinding attack, on the other hand, manipulates the outcomes a receiver is recording using strong laser pulses that alter the detectors' operation regime [158]. Once in the linear response regime, the attacker can fully control which of Bob's detectors will click by setting the laser pulse power just above the linear detection threshold and choosing its polarization state. If the pulse ends up in the opposite measurement basis, its power is equally distributed between the two detectors. Since the power per detector is significantly below the click threshold, no detector fires, and the signal appears as lost (introducing 50 % of transmission loss). In combination with any discrimination attack where Eve first learns the encoded bit, she can now ensure that Alice and Bob will not measure an increased error rate during the sifting phase and the attack remains unnoticed. A countermeasure was proposed, e.g., in [159], or can be realized using measurement device independent schemes [35, 160–163].

2.2 Quantum dots as quantum light sources

Throughout the introduction to practical quantum cryptography, we have seen that the photonic state and, in particular, its photon number distribution are vital for the security of quantum cryptographic protocols. While early developments and implementations in the field were powered by WCS or photons generated through SPDC, the intriguing potential of a source that emits light with sub-Poisson statistics was noted from the beginning and motivated researchers to realize such photonic sources. Basing the source on a two-level system would eliminate the intrinsic trade-off between single-photon purity and efficiency present in PDS. Instead, optimizing the photon-number distribution is transferred to a fabrication challenge, but without fundamental limitations.

2.2.1 Quantum dots in a nutshell

Solid-state systems offer exceptional control over the environment and among those QDs benefit especially from advanced epitaxy techniques developed in the III-V semiconductor industry. In the past decade, they have become the most efficient type of SPS by far, thanks to the continuous progress in material quality, fabrication processes, and the co-integration with photonic cavities [55, 56].

Conceptually, a QD represents a structure characterized by a zero-dimensional density of states, leading to a distinct quantization of energy levels [92]. Often described as artificial atoms, these systems exhibit a discrete energy spectrum akin to that of atoms, while benefiting from the advantages of a solid-state environment. These include spatial fixation and the ability to integrate them into additional structures, enabling precise control of their emission properties. Since the first demonstrations [47, 164], the potential of semiconductor QDs as a light-matter interface sparked considerable interest in fundamental studies of quantum optics [62–66] and for a wide range of quantum information technologies [67–73].

Self-assembled QDs are most commonly grown in the Stranski-Krastanov mode [165] (another widespread method is droplet-epitaxy [166]), where the semiconductor material that later forms the QD is deposited on top of a host matrix by metalorganic vapor-phase epitaxy (MOVPE) or molecular-beam epitaxy (MBE)[11]. As a thin layer (wetting layer) builds up, the lattice mismatch between the two different semiconductors induces increasing strain, which is eventually relieved by the development of small islands forming the QDs. Afterwards, they are covered by a capping layer that often consists of the same material as the host matrix, which completes the fabrication process (see Fig. 2.5a). The resulting three-dimensional islands are significantly smaller in the growth direction compared to their in-plane size, overall resembling the shape of a plano-convex lens [167]. Additional processing steps may be necessary to integrate an electronic diode or photonic extraction structure.

[11] For details about epitaxial techniques and their impact on QDgrowth, see the literature, for example [92].

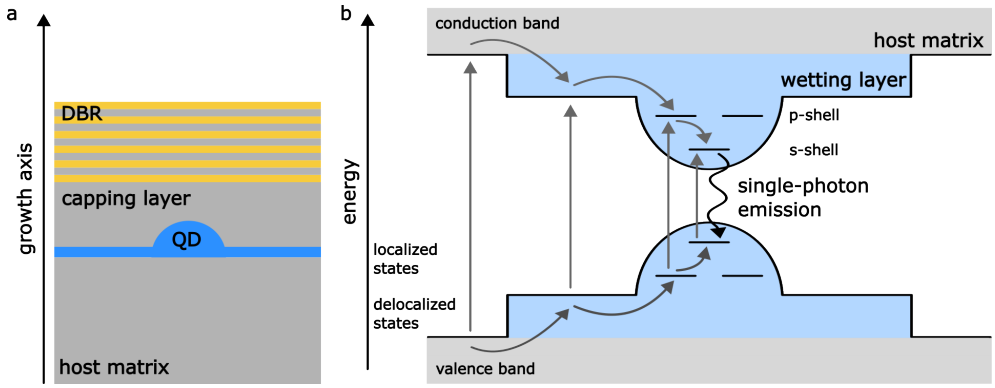


Figure 2.5: **Illustration of a QD and its energy structure.** (**a**) Exemplary layering of semiconductor materials forming a QD with a top distributed Bragg reflector (DBR). (**b**) Schematic energy level of such a QD including host material, the energy band originating from the thin wetting layer and hosting a set of delocalized states, and the QD itself. The set of local QD states features multiple energy levels; here, only the s- and p-shells are shown. Grey straight arrows indicate exciton formation through optical pumping, curved arrows denote relaxation processes into the lowest energy state, and the black wiggly arrow illustrates the single-photon emission from the s-shell recombination.

The QDs studied here are made from In(Ga)As sitting on a GaAs host matrix, where a metamorphic buffer layer (MMB) was used to realize the strain relief required to allow single-photon emission in the telecom C-band around 1550 nm [168].

Energy landscape and optical transitions

The lens shape of the QDs induces a strong confinement along the growth direction but a weaker in-plane confinement which justifies to model the resulting potential as a product of the two components, where only the lowest energy state of quantum well along the growth axis must be considered [169]. In the absence of a magnetic field, the eigenenergies for the single-particle state in a two-dimensional harmonic potential are

$$E_{(n,l)} = (s+1)\,\hbar\omega_0 \tag{2.2-1}$$

where ω_0 is the transition frequency and $s = 2n+|l|$ is in analogy to the atomic shells referred to as the shell index, and n and l are the radial and angular momentum quantum numbers, respectively.

The degeneracy of $2(s+1)$ of each shell follows from Pauli's exclusion principle and the possible spin configurations; that is, the s-shell has two degenerate levels, the p-shell four, and so on. Naturally, a magnetic field will lift the degeneracy in $|l|$. A schematic band structure is shown in Fig. 2.5b.

Symmetry arguments for the electron and hole envelope wave function show that radiative recombinations are only allowed for transitions between levels of

the same shell index $\Delta s = 0$, so for s(p, d, ..)-like electrons and s(p, d, ..)-like holes [92]. Light-induced excitons in a QD are primarily formed from pairs of electrons and heavy holes (heavy refers to the effective mass and the respective dispersion relation). Due to spin-orbit coupling and material strain, the heavy hole represents the highest energy state of the valence band [165, 170] with a total spin of $j_z = \pm 3/2$. When pairing these with an electron of spin $s_z = \pm 1/2$ four exciton configurations with a total angular momentum of $M_J = \pm 1$ and $M_J = \pm 2$ are possible. Due to optical dipole selection rules ($\Delta M_J = \pm 1$), the latter do not couple to light and are therefore dubbed dark excitons. The description in this thesis will be limited to bright excitons with $M_J = \pm 1$, which emit a circularly polarized photon upon recombination.

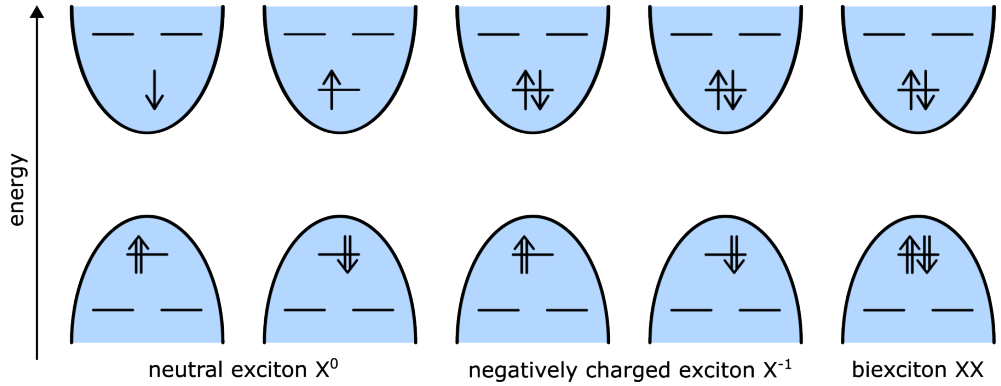


Figure 2.6: **Excitonic complexes in a QD.** A selection of common excitonic complexes possible in a QDare schematically shown in a band structure. The conduction band at the top hosts one or two electrons in the s-shell, where the arrow indicates its spin state, and the valence band at the bottom shows the according spin state of the heavy hole(s). Since the neutral biexciton features only fully populated shells, there is no spin state degeneracy.

Besides this simple excitation complex (neutral exciton $|X^0\rangle$), many more multi-particle configurations are possible within a QD and only a brief overview of the most important types is given here (see also Fig. 2.6). If the QD hosts an additional electron (hole) when the exciton is created through an incident laser field, the resulting state will be a negatively (positively) charged exciton, which is also called a trion and denoted by $|X^{\pm}\rangle$. Starting from the neutral ground state $|G\rangle$, two-photon absorption transfers the system to the biexciton state $|XX\rangle$, from which the QD decays in a radiative cascade through $|X^0\rangle$ back to $|G\rangle$. Note that the transitions $|XX\rangle \rightarrow |X^0\rangle$ and $|X^0\rangle \rightarrow |G\rangle$ are nondegenerate due to the binding energy induced by Coulomb interaction [171]. If the energy splitting of the neutral

exciton states (denoted as the fine structure splitting (FSS)) is zero[12], the emitted photons are in a polarization-entangled $|\phi^+\rangle$ Bell state [54, 173]. However, due to strain, magnetic fields, nondegenerate modes of the photonic extraction structure, or even the excitation laser field [174–176], it is very challenging to remove the FSS entirely, and the output state shows only an imperfect fidelity to a Bell state[13].

More complex situations, including multiply charged or higher-order states, couplings to dark states, and lifting of degeneracy under the influence of a magnetic fields are not relevant for this work and will not be discussed here, but are readily found in the literature.

2.2.2 Figures of merit

Before detailing a theoretical model describing the light-matter interaction of a QD driven by a laser field, it is helpful to define some figures of merit to quantitatively analyze the light emitted by the QD. Two key metrics, system efficiency (Sec. 2.2.2.1) and single-photon purity (Sec. 2.2.2.2), follow directly from the aspiration of developing a source that emits *exactly* one photon in response to a trigger laser pulse. For some applications of SPSs, the photonic coherence of the emission (Sec. 2.2.2.3), manifesting itself in the indistinguishability and PNC, is also vital.

2.2.2.1 System efficiency

Any efficiency is best measured as the ratio of successful trials to the total number of trials. Limiting ourselves to pulsed operation of an SPS, we consider each laser excitation pulse as a trial. Then, the photon generation efficiency is simply

$$\eta_{\text{gen}} = \frac{\#\text{successfull generation events}}{\#\text{trials}} . \tag{2.2-2}$$

However, since some photons are always lost before they can be detected, this efficiency can only be measured indirectly. For example, inevitable loss is caused by the imperfect transmission of the optics rejecting the excitation laser field and background signals η_{fil}, the imperfect collection of the emission from the quantum emitter into the target spatial mode η_{col}, and the non-unity efficiency of the single-photon detector η_{det}. When preparing a stream of single photons for an application, both filtering and collection (typically into a SMF) are strictly necessary, and

[12] Both spin configurations of the neutral exciton ($M_J = \pm 1$) are degenerate, but if the rotational symmetry of the QD is broken (most commonly by strain) they are no longer eigenstates but mix with each other forming new eigenstates. These are the symmetric and antisymmetric superpositions of the angular momentum states and are no longer degenerate [172].

[13] A promising but hard-to-scale approach for removing strain demonstrated a fidelity $> 99.5\,\%$ [177] by adding a multi-leg piezo membrane to compensate the symmetry breaking through external control. Non-resonant two-photon excitation schemes [178–180], on the other hand, can remove (or strongly reduce) the spectral overlap with $|X^0\rangle$ such that the laser-induced level splitting can be avoided.

thus it has become customary to include these inefficiencies when quoting the performance of a SPS using the system efficiency

$$\eta_{\mathrm{sys}} = \eta_{\mathrm{gen}}\eta_{\mathrm{fil}}\eta_{\mathrm{col}}\,. \tag{2.2-3}$$

In the experiment, η_{sys} is determined from the event rate R measured on a click detector with known η_{det} as

$$\eta_{\mathrm{sys}} = \frac{R}{\nu_{\mathrm{rep}}\eta_{\mathrm{setup}}\eta_{\mathrm{det}}}\,, \tag{2.2-4}$$

where ν_{rep} is the excitation repetition rate and η_{setup} accounts for all loss of photons between the collection and the detection stage. In many applications, the stream of single photons does not only need to be in single spatial but also well-defined polarization mode, thus defining a polarized system efficiency is useful

$$\eta_{\mathrm{sys,pol}} = \eta_{\mathrm{sys}}\eta_{\mathrm{pol}}\,. \tag{2.2-5}$$

For a specialized field, even more definitions may be helpful, for which we refer the reader to Refs. [92, 181]. At this point, we must mention an ongoing debate about the metrics and wording used when benchmarking SPSs, where different communities resort to various terminologies. In an effort to avoid ambiguous language, the use of terms such as "brightness" is deprecated [181].

Importantly, η_{sys} as defined in Eq. 2.2-4 does not represent the *single-photon* system efficiency since a click detector cannot distinguish whether one or more photons were emitted; a correction based on the single-photon purity is presented in Sec. 2.2.2.4. The system efficiency η_{sys} may also be expressed in terms of photon number probabilities (at the collection stage)

$$\eta_{\mathrm{sys}} = \sum_{i>0} p_i\,. \tag{2.2-6}$$

Over the two decade, improved fabrication and the integration within photonic structures (see. Sec. 2.2.5) increased the system efficiency of QDs significantly, and values of 57 % [56] and 71 % [55] have recently been achieved in open-cavity systems.

2.2.2.2 Single-photon purity

Every realistic SPS suffers from a non-zero chance to produce more than a single photon in response to the excitation pulse. Throughout this work, we refer to a multiphoton event whenever two or more photons are emitted after the quantum emitter is probed by a single laser pulse. Notably, this includes not only "true" two-photon Fock states but also all events in which two or more photons were found in close succession—more precisely within a time interval $T_{\mathrm{int}} \leqslant T_{\mathrm{rep}}$—instead of in the same temporal mode.

To measure this probability in the experiment, a balanced beam splitter (BS) divides the quantum light equally between two output paths, each connected to a single-photon detector (Hanbury Brown and Twiss experiment (HBT) configuration). Correlating the detection events of both channels in a start-multistop fashion with a time tagger, the second-order correlation function $G^{(2)}(\tau)$ (see Appendix A) is recorded. For a pulsed excitation at ν_{rep}, all coincidences within $[-T_{\mathrm{rep}}/2,\ T_{\mathrm{rep}}/2]$ are counted as multiphoton event

$$G^{(2)}(0) = \int_{-T_{\mathrm{rep}}/2}^{T_{\mathrm{rep}}/2} d\tau\ G^{(2)}(\tau)\,. \tag{2.2-7}$$

For comparison independent of the detection rates R_1 and R_2 and integration time Δt_{int}, the total number of multiphoton events must be normalized

$$g^{(2)}(0) = \frac{G^{(2)}(0)}{CC_{\mathrm{uncor}}} \tag{2.2-8}$$

where

$$CC_{\mathrm{uncor}} = R_1 R_2 \Delta t_{\mathrm{int}} T_{\mathrm{rep}} \tag{2.2-9}$$

is the number of coincidence events an uncorrelated light source with the same detection rates would produce. For a perfect SPS with $\gamma_{\mathrm{rad}} \gg \nu_{\mathrm{rep}}$, coincidences observed at $\tau = \pm T_{\mathrm{rep}}$ represent the so-called Poissonian level, that is, the event statistics are fully uncorrelated as expected from Poisson-distributed light. However, realistic solid-state quantum emitters might experience an effect called blinking [49, 182, 183], where they oscillate between a dark and bright state on time scales $T_{\mathrm{bl}} \gg T_{\mathrm{rep}}$[14]. This introduces correlations between emission events on these longer time scales T_{bl} and the Poissonian level is only reached for $\tau > T_{\mathrm{bl}}$. For example, if the emitter is in a bright state at $\tau = 0$, then it is more likely to detect a coincidence at $\tau = T_{\mathrm{rep}}$ than at $\tau \to \infty$.

Hence, for experimental data, the normalized correlation function is calculated as

$$g^{(2)}(0) = \frac{\int_{-T_{\mathrm{rep}}/2}^{T_{\mathrm{rep}}/2} d\tau\ G^{(2)}(\tau)}{\int_{T_{\mathrm{rep}}(k-1/2)}^{T_{\mathrm{rep}}(k+1/2)} d\tau\ G^{(2)}(\tau)}\,, \tag{2.2-10}$$

where k must be chosen large enough to ensure normalization to Poissonian statistics. Alternatively, one might model and fit the blinking behavior to extract the correctly normalized $g^{(2)}(0)$ from $G^{(2)}(\tau)$ [92].

Finally, we define the single-photon purity

$$\mathcal{P} = 1 - g^{(2)}(0)\,. \tag{2.2-11}$$

[14] Using advanced fabrication, this challenges is successfully addressed by embedding the QD in a electrical diode with tunnel barrier [81, 184] which not only stabilizes the charge environment but allows to prepare the QD in a specific charge state [185, 186]. In this way, blinking-free emission is routinely demonstrated [84].

Much research has been dedicated to eliminating sources of residual multiphoton contributions (for an overview, see Chapter 8 in Ref. [92]) and several causes, including detector background noise [187], emission from the wetting layer [188], other transitions or neighboring QDs, re-filling from trapped charges [189, 190], and re-excitation [96, 191] have been identified. Among those, re-excitation is the only multiphoton process that is not rooted in fabrication challenges, but presents a fundamental limit and, as such, cannot be resolved by technological advancements. All reports of record low $g^{(2)}(0)$ values used appropriate means to mitigate the impact of re-excitation, either by harnessing an effective three-level system (3LS) [97, 187] or, very recently, using specifically tailored spectral filtering [192]. More details on the light-matter interaction leading to re-excitation will be given in Sec. 2.2.4.3.

2.2.2.3 Photon coherence

In optics, coherence describes the ability of two or more light fields to interfere with each other. Extending the classical picture, the single-photon coherence of a quantum emitter is set by the phase stability of the generated photon-number state. In Sec.2.1.4.3, we have seen that even if each individual photon is prepared in a coherent superposition of Fock states, the time-averaged system appears in a mixed state if the phase between the Fock states is not fixed.

For a two-level system (2LS), there are two types of processes that reduce the interference visibility: either population decay through spontaneous emission or scattering processes that alter the phase but do not change the population. The time for which a stable phase relation is maintained—characterized by the coherence time T_2—is therefore affected by the decay time T_1 ($1/T_1 = 1/\tau_{\text{nonrad}} + 1/\tau_{\text{rad}}$) of the upper level and the pure dephasing time[15] T_2^* of the excited state

$$\frac{1}{T_2} = \frac{1}{2\,T_1} + \frac{1}{T_2^*}\,. \tag{2.2-12}$$

For a spontaneous emission that dominates the transition ($T_1 \simeq \tau_{\text{rad}}$) and for negligible pure dephasing ($T_2^* \to \infty$), the coherence reaches an intrinsic limit and the emission is said to be Fourier-transform (FT) limited

$$T_2 = 2\,\tau_{\text{rad}}\,. \tag{2.2-13}$$

Equivalently, the emission is called lifetime limited when the spectrum is only as wide as dictated by time-energy uncertainty, such that $\Delta\nu = \gamma_{\text{rad}}/2\pi$ where $\Delta\nu$ is the emission linewidth measured as full width at half maximum (FWHM) and $\gamma_{\text{rad}} = 1/\tau_{\text{rad}}$ is the radiative decay rate.

[15] This effect is called "pure" dephasing because only the phase is randomized but the population is preserved.

In realistic QDs, the coherence may be affected by a variety of processes that occur in the semiconductor environment and is often inferred as

$$T_2 = \int_{-\infty}^{\infty} d\tau \ |g^{(1)}(\tau)|^2 \tag{2.2-14}$$

from the first-order correlation function $g^{(1)}(\tau)$ (see Appendix A). Here, $g^{(1)}(\tau)$ is measured using Michelson interferometry [193, 194], or alternatively, the emission line shape $g(\omega)$ is recorded, which relates to $g^{(1)}(\tau)$ by the Wiener-Khintchine theorem

$$g(\omega) = \int_{-\infty}^{\infty} d\tau \ e^{i\omega\tau} g^{(1)}(\tau)\,. \tag{2.2-15}$$

Charge noise, spin noise and phonon interactions act on their individual time scales [68, 81, 182, 195] and cause pure dephasing when these are faster than the radiative decay [196]. Charge noise is introduced by nearby defects, charge traps or surface states that may be or may not be occupied and influence the potential of the QD through the confined Stark effect [182, 197–199]. The probabilistic occupation induces Gaussian broadening to the transition line when averaged over many excitation cycles, and is referred to as spectral wandering [198, 200–202]. The Overhauser field created by the randomly oriented nuclear spins around the QD fluctuates typically on faster time scales [182] and couples to the exciton through exchange interactions, contributing to the decoherence [182, 201, 203]. Phonon-induced decoherence can arise from a variety of phonon-exciton interactions as discussed also in Sec. 2.2.3.3 but for QDs at cryogenic temperatures pure dephasing caused by coupling to LA phonons dominates [195, 204–208].

Only dephasing-free single-photon states yield unity indistinguishability when tested in a Hong–Ou–Mandel (HOM) experiment [209] that quantifies the two-photon interference (TPI) visibility. Achieving visibilities close to unity is critical for quantum information processing, as it enables Bell-state measurements and linear optic fusion gates [210–213]. Although spectral filtering can boost the TPI visibility, it is only a transitional solution, as it comes at the cost of reduced efficiency, and achieving dephasing-free emission remains the main pathway toward indistinguishable photons.

Finally, a realistic quantum emitter generates a photon-number state with zero-, one-, and potentially two-photon contributions, which can be either in a statistically mixed state, a coherent superposition, or anywhere in between [214]. For sufficiently slow dephasing, coherent excitation schemes[16] will imprint a fixed phase relation between the number states [214, 215], which is inherited from the laser pulse and thus stable over many excitation cycles. Using an HOM interferometer with a stable and adjustable phase, the PNC can be extracted from the intensity visibility observed during a phase scan. While a low TPI visibility necessarily impairs the

[16]The observation of Rabi rotations in the population probabilities is a clear signature of coherent driving. Coherent excitation methods include mostly, but not exclusively, resonant techniques. For more details, see Sec. 2.2.4.

PNC, the reverse is not true. Highly indistinguishable photons can be obtained using incoherent driving schemes [57, 216–218] that do not feature any measurable PNC [219].

Importantly, for some applications, e.g., twin-field QKD [102, 142] PNC is required, whereas for others, including most cryptographic protocols, it must be avoided [45, 91].

2.2.2.4 Extracting photon number probabilities from the experiment

Deriving the security bounds for cryptographic protocols as introduced in Sec. 2.1.4 assumes knowledge of the photon number probabilities $\{p_i\}$ of the emitted state. However, experimentally, these are not directly accessible but must be inferred from a set of measurements and/or assumptions. While for sources with Poissonian photon statistics a measurement of the mean photon number suffices, the independence of $\{p_i\}$ for SPSs requires, in principle, infinitely many measurements of high-order correlation functions. In practice, assuming $p_{i\geq 4} = 0$ is well justified [220] for QD sources and the remaining multiphoton probabilities are summarized in $p_\mathrm{m} = p_2 + p_3$. Then, making use of the normalization condition $\sum_i p_i = 1$, $\{p_i\}$ can be calculated from the system efficiency η_sys and $g^{(2)}(0)$ value as is shown in the following.

Starting from

$$g^{(2)}(0) = \frac{\langle n(n-1)\rangle}{\bar{n}^2} = \frac{\sum_{i=1} i\,(i-1)\,p_i}{\left(\sum_{i=1} i\,p_i\right)^2} \tag{2.2-16}$$

as given in Appendix A, the steps we derived in our work [1] are retraced for the readers comfort. Using $p_1 > p_\mathrm{m}$, we write

$$\begin{aligned} g^{(2)}(0) &= \frac{2p_2 + 6p_3}{(p_1 + 2p_2 + 3p_3)^2} \\ &= \frac{2p_\mathrm{m} + 4p_3}{(p_1 + 2p_\mathrm{m})^2}\left[1 + \frac{p_3}{p_1 + 2p_\mathrm{m}}\right]^{-2} \\ &\simeq \frac{2p_\mathrm{m}}{(p_1 + 2p_\mathrm{m})^2}\,, \end{aligned} \tag{2.2-17}$$

and truncate it at the zeroth order in p_3. Computing the error $F(p_1, p_2, p_3)$ introduced by our approximations of $g^{(2)}(0)$, we note that $F(p_1, p_2, p_3) > 0$ for all $p_i \in (0, 1)$ where $i = \{1, 2, 3\}$—showing that the right-hand side of Eq. 2.2-17 provides an actual lower bound—and $F(p_1, p_2, 0) = 0$ showing that the bound is tight in the limit of vanishing p_3. After rewriting Eq. 2.2-17 as a lower bound and expressing it in terms of η_sys,

$$g^{(2)}(0) \geq \frac{2p_\mathrm{m}}{(\eta_\mathrm{sys} + p_\mathrm{m})^2}\,, \tag{2.2-18}$$

we expand and rearrange again

$$p_{\mathrm{m}}^2 + \left(2\eta_{\mathrm{sys}} - \frac{2}{g^{(2)}(0)}\right) p_{\mathrm{m}} + \eta_{\mathrm{sys}}^2 \geq 0. \tag{2.2-19}$$

Since $p_{\mathrm{m}} \geq 0$ and $1 \geq \eta_{\mathrm{sys}} \geq 0$, the inequality has only one solution for a SPS (that is $\frac{1}{2} \geq g^{(2)}(0) \geq 0$ [221]) and the set of photon number probabilities follows

$$p_{\mathrm{m}} \leq \frac{1 - \eta_{\mathrm{sys}}\, g^{(2)}(0) - \sqrt{1 - 2\eta_{\mathrm{sys}}\, g^{(2)}(0)}}{g^{(2)}(0)} \tag{2.2-20a}$$

$$p_1 = \eta_{\mathrm{sys}} - p_{\mathrm{m}} \tag{2.2-20b}$$

$$p_0 = 1 - p_1 - p_{\mathrm{m}}\,. \tag{2.2-20c}$$

Eqs. 2.2-20 hold for any source with sub-Poissonian emission, at any distance, and provide an explicit bound that only depends on experimentally accessible parameters. Furthermore, a first-order Taylor expansion of Eq. 2.2-20b provides a simple correction based on a lower bound to estimate the *single-photon* system efficiency (which is equivalent to p_1, see Eq. 2.2-6) from the experimental measurement

$$\begin{aligned} p_1 &\geq \eta_{\mathrm{sys}} - \frac{\eta_{\mathrm{sys}}^2\, g^{(2)}(0)}{2} \\ &\simeq \eta_{\mathrm{sys}} \left(1 - \frac{\eta_{\mathrm{sys}}\, g^{(2)}(0)}{2}\right). \end{aligned} \tag{2.2-21}$$

2.2.3 Light–matter interaction

After introducing the basis of what a quantum dot is and how to benchmark its performance as SPS, we will now describe how the system interacts with a laser field. We will start with the simplest model that only includes the Hamiltonian of the 2LS and a monochromatic laser field. Subsequently, a reservoir allowing for spontaneous emission and a free-dephasing Hamiltonian accounting for phonon interactions are added. In parts, this section follows the derivation given in [222]. The complementing introduction of the phonon coupling is based on a series of publications indicated accordingly.

2.2.3.1 Two-level system in a laser field

We denote the two states of the QD as ground and excited state, $|G\rangle$ and $|X\rangle$, with eigenenergies of 0 and $\hbar\omega_0$ and which are connected via the lowering and raising operators

$$\sigma_- = |G\rangle\langle X| \qquad\qquad \sigma_+ = |X\rangle\langle G|\,. \tag{2.2-22}$$

Using the photonic creation and annihilation operators, $a^\dagger$ and a, and the angular frequency ω_L of the monochromatic laser field, the Hamiltonian reads

$$H_0 = \hbar\omega_0\, |X\rangle\langle X| + \hbar\omega_L (a^\dagger a + 1/2)\,. \tag{2.2-23}$$

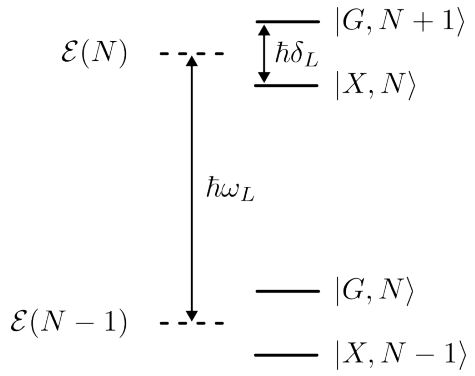


Figure 2.7: **Energy spectrum of QD-laser system without interaction potential.** Each manifold $\mathcal{E}(N)$ hosts N photons and is split into two states depending on the occupation of the QD. If the QD is excited, there is a photon less in the laser field and the state's energy is lowered by $\hbar\delta_L$ since $\omega_L > \omega_0$. Neighboring subspaces are separated by the energy of one laser photon $\hbar\omega_L$.

Since we are interested in near-resonant interactions, we use $\delta_L \ll \omega_0$, where $\delta_L = \omega_L - \omega_0$ is the detuning between the 2LS resonance and the light field and is always positive for our purposes. The spectrum of eigenstates of H is shown in Fig. 2.7 and can be divided in N subspaces, where N denotes the number of photons present in the laser field.

Each subspace, or manifold, consists of a ground state $|G, N+1\rangle$ and an excited state with one laser photon less $|X, N\rangle$ that are $\hbar\delta_L$ apart from each other

$$\mathcal{E}(N) = \{|G, N+1\rangle, |X, N\rangle\}. \tag{2.2-24}$$

Neighboring subspaces $\mathcal{E}(N)$ and $\mathcal{E}(N-1)$ have an energy difference of $\hbar\omega_L$.

Including an interaction potential

$$H_{QD-L} = -\boldsymbol{d} \cdot \boldsymbol{E}(\boldsymbol{r}), \tag{2.2-25}$$

where $\boldsymbol{E}(\boldsymbol{r})$ is the single-mode, quantized electric field of the laser at position $\boldsymbol{r}$ of the QD

$$\boldsymbol{E}(\boldsymbol{r}) = \sqrt{\frac{\hbar\omega_L}{2\varepsilon_0 V}}\boldsymbol{e}_L(a + a^\dagger) \tag{2.2-26}$$

with $\boldsymbol{e}_L$ being the laser polarization axis. The electric dipole operator of the 2LS is

$$\boldsymbol{d} = \boldsymbol{d}_{XG}(\sigma_- + \sigma_+) \tag{2.2-27}$$

assuming only real transition elements, that is $\boldsymbol{d}_{XG} = \langle X|\,\boldsymbol{d}\,|G\rangle = \langle G|\,\boldsymbol{d}\,|X\rangle = \boldsymbol{d}_{GX}$. For simplicity, we assume a laser polarization parallel to the dipole orientation and introduce the light-dipole coupling constant

$$g = -d_{XG}\sqrt{\frac{\omega_L}{2\hbar\varepsilon_o V}} \tag{2.2-28}$$

to rewrite the interaction potential as

$$H_{QD-L} = \hbar g(\sigma_+ a - \sigma_- a^\dagger)\,. \tag{2.2-29}$$

Here, the rotating wave approximation (RWA) was already applied such that the terms with $\sigma_- a$ and $\sigma_+ a^\dagger$ can be safely neglected. Since these terms correspond to creating (destroying) a photon while additionally raising (lowering) the 2LS, one also intuitively understands that these processes should have a vanishing probability. For a coherent laser field with many photons (average photon number $\langle N\rangle \gg 1$), the coupling strength induced by the interaction potential between both states of a manifold can be expressed by

$$\nu_N = \langle G, N+1|\, H_{QD-L}\, |X, N\rangle \simeq \hbar g\sqrt{\langle N\rangle} \tag{2.2-30}$$

and is approximately constant for small variations in $\langle N\rangle$. In other words, for a given driving field, the coupling does not significantly change even if a few photons in the laser field are lost, and we can consider a constant coupling across all manifolds relevant for our description. Introducing the Rabi frequency by analogy from a semi-classical description [223]

$$\hbar\Omega = -d_{XG}E_0 \qquad \text{where} \qquad E_0 = 2\sqrt{\frac{\hbar\omega_L}{2\varepsilon_0 V}\langle N\rangle} \tag{2.2-31}$$

we find the coupling

$$\nu_N = \frac{\hbar\Omega}{2}\,. \tag{2.2-32}$$

The interaction potential alters the energy structure such that the product states of the bare 2LS and photonic states used in Fig. 2.7 are no longer eigenstates. In the rotating frame, the new eigenstates, the so-called laser-dressed states, are

$$|\varphi_G, N\rangle = \cos(\vartheta)\,|G, N+1\rangle - \sin(\vartheta)\,|X, N\rangle \tag{2.2-33a}$$
$$|\varphi_X, N\rangle = \sin(\vartheta)\,|G, N+1\rangle + \cos(\vartheta)\,|X, N\rangle \tag{2.2-33b}$$

using the admixing angle

$$\vartheta = \frac{1}{2}\arctan\left(\frac{\Omega}{\delta_L}\right)\,. \tag{2.2-34}$$

The notation of the dressed states is chosen to indicate that $|\varphi_g, N\rangle$ ($|\varphi_X, N\rangle$) coincides with the ground (excited) state for vanishing laser drive $\Omega \to 0$. Further, $|\varphi_g, N\rangle$ is always higher in energy since $\delta_L > 0$ and using the effective Rabi frequency

$$\Omega_{\text{eff}} = \sqrt{\Omega^2 + \delta_L^2} \tag{2.2-35}$$

we find an energy level splitting of $\Delta E = \hbar\Omega_{\text{eff}}$. The eigenstates and eigenenergies for the three scenarios are pictured side-by-side in Fig. 2.8 using a rotating frame.

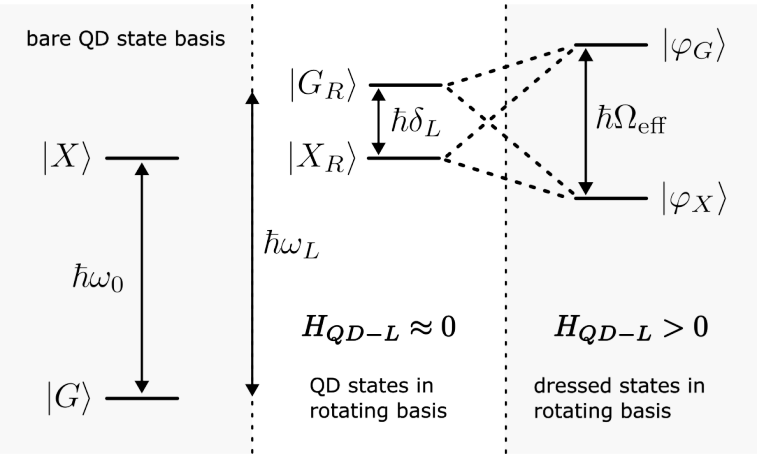


Figure 2.8: **QD eigenstates depending on laser field.** If there is no laser field present, the eigenstates are simply the undisturbed states of a 2LS Hamiltonian (left panel). For a QD–laser system but without interaction (such as Eq. 2.2-23), the basis states are expressed in the frame rotating at the laser frequency (center panel). This description also captures the case of formally including the interaction potential but allowing only a very small interaction strength (small dipole moment or driving field, such that $\Omega \to 0$, see Eqs. 2.2-29 and 2.2-31). For significant coupling (right panel), the rotating basis states mix, and the resulting new eigenstates are referred to as dressed states.

2.2.3.2 Spontaneous emission

Now, we turn toward the spontaneous emission of a photon, which is considered an instantaneous process, sometimes referred to as a quantum jump. Consequently, the dynamics of the driving field are much slower and can be neglected during spontaneous emission. However, the time scale between two quantum jumps given by the radiative lifetime τ_{rad} is much longer than the characteristic time of absorption or stimulated emission, $\tau_{\text{rad}} > 1/\Omega_{\text{eff}}$, and the 2LS can undergo many cycles of radiation coupling between two events of spontaneous emission.

Recall that the dipole operator d_{XG} couples the bare (that is, undressed) states $|X\rangle$ and $|G\rangle$ of the 2LS. Thus, when considering the dressed states in the $\mathcal{E}(N)$ manifold—which have a $|X, N\rangle$ component (see Eq. 2.2-33)—and the dressed states in $\mathcal{E}(N-1)$ that contain $|G, N\rangle$, we find that all four possible transitions between neighboring energy subspaces are allowed by dipole coupling. The complete energy scheme is sketched in Fig. 2.9. Since the photon is emitted into a reservoir of modes that is not captured in the Hamiltonian, the system "loses" energy through spontaneous emission, which is indicated by the radiative decay arrows. Two of the decay paths are degenerate, resulting in three distinct frequency bands at ω_L and $\omega_L \pm \Omega_{\text{eff}}$. The transition probabilities are determined by the mixing angle ϑ as

$$\langle \varphi_G, N|\, e\rangle \langle g|\, \varphi_G, N-1\rangle \propto -\cos(\vartheta)\sin(\vartheta) \tag{2.2-36a}$$
$$\langle \varphi_G, N|\, e\rangle \langle g|\, \varphi_X, N-1\rangle \propto -\sin^2(\vartheta) \tag{2.2-36b}$$
$$\langle \varphi_X, N|\, e\rangle \langle g|\, \varphi_G, N-1\rangle \propto \cos^2(\vartheta) \tag{2.2-36c}$$
$$\langle \varphi_X, N|\, e\rangle \langle g|\, \varphi_X, N-1\rangle \propto \cos(\vartheta)\sin(\vartheta)\,. \tag{2.2-36d}$$

Recalling that $0 \leq \vartheta < \pi/4$, we find that the transition $|\varphi_X, N\rangle \to |\varphi_G, N-1\rangle$ dominates, especially for $\Omega < \delta_L$. Notably, there are no dipole-allowed transitions between dressed states of the same subspace or between any manifolds with $\Delta N > 1$. Deriving the master equations for the dressed atom and projecting it on

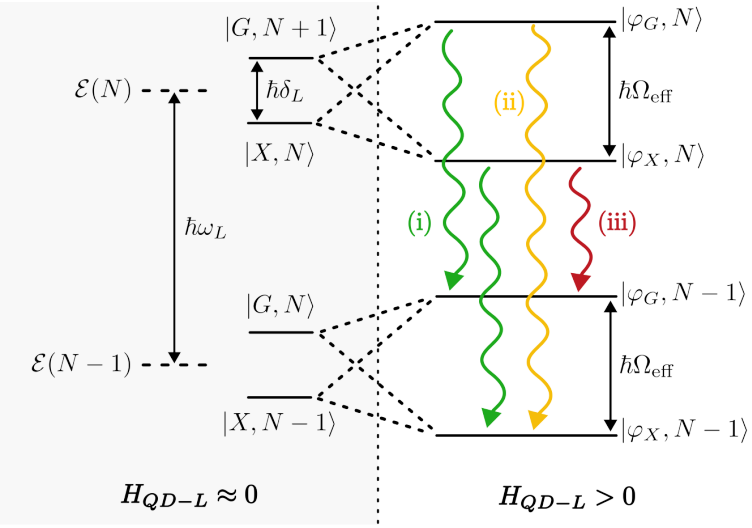


Figure 2.9: **Spontaneous emission from a dressed 2LS.** The energy scheme from Fig. 2.8 is extended by another manifold to indicate the possible dipole transitions of spontaneous emission. In the unpertubed basis, only transitions $|X, N\rangle \to |G, N\rangle$ are dipole-allowed but since both dressed states of the upper (lower) manifold have a $|X, N\rangle$ ($|G, N\rangle$) admixture, all four inter-manifold transitions of the dressed states are possible. There are two degenerate transitions (i) at $\omega_e = \omega_L$ and one at either the high (ii) and low (iii) frequency side $\omega_e = \omega_L \pm \Omega_{\text{eff}}$. Under the right conditions, all three frequency bands of spontaneous emission form, which are known as Mollow-Triplett [224] (for more details, see Sec. 2.2.4.1).

the dressed-state basis in the secular limit (that is $\tau_{\text{rad}} \gg 1/\Omega_{\text{eff}}$), one can show that the steady-state populations of the upper and lower dressed-state types are [222]

$$\pi^{st}_{\varphi_G} = \frac{\cos^4(\vartheta)}{\sin^4(\vartheta) + \cos^4(\vartheta)} \tag{2.2-37a}$$
$$\pi^{st}_{\varphi_X} = \frac{\sin^4(\vartheta)}{\sin^4(\vartheta) + \cos^4(\vartheta)}\,. \tag{2.2-37b}$$

Again, considering the range of ϑ, most of the population is found in $|\varphi_G\rangle$, which is easy to understand from Eqs. 2.2-36 showing that transitioning from $|\varphi_X\rangle$ to $|\varphi_G\rangle$ is more likely than vice versa. Since $|\varphi_G\rangle$ is dominated by the ground state contribution $|G\rangle$ (see Eq. 2.2-33a), the probability of finding the 2LS in the excited state remains low.

2.2.3.3 Interaction with a phonon bath

So far, we have described a 2LS coupled to a laser field and spontaneously emitting into a reservoir of photonic modes. While an atomic system would additionally feature rotational or vibrational degrees of freedom extending the level structure, QDs are embedded in a solid-state environment that eliminates rotational degrees of freedom but interacts through the crystal's vibrational modes with the 2LS. These interactions are diverse, and their nature depends on several parameters, including crystal symmetry and material, as well as the energy range.

We introduce the electron-phonon coupling in our model using a free-phonon and pure-dephasing Hamiltonian [225], sometimes referred to as the independent boson model [226],

$$H_{QD-P} = \sum_{\boldsymbol{q}} \hbar\omega_{\boldsymbol{q}} b_{\boldsymbol{q}}^\dagger b_{\boldsymbol{q}} + \sum_{\boldsymbol{q}\nu} \hbar n_\nu \left(\gamma_{\boldsymbol{q}} b_{\boldsymbol{q}} + \gamma_{\boldsymbol{q}}^* b_{\boldsymbol{q}}^\dagger\right) |\nu\rangle\langle\nu| \,, \tag{2.2-38}$$

where $b_{\boldsymbol{q}}^\dagger$ ($b_{\boldsymbol{q}}$) creates (destroys) a phonon with frequency $\omega_{\boldsymbol{q}}$ and wave vector $\boldsymbol{q}$, n_ν counts the excitons in state $|\nu\rangle$ and γ_q denotes the electron-phonon coupling constants. In the most general description, the coupling $\gamma_{\boldsymbol{q}}$ captures a range of different phonon types, including longitudinal optical (LO), LA, transversal acoustic (TA), and surface phonons. For QDs, we can safely ignore surface effects and confine ourselves to bulk phonons [195], except in cases where surface acoustic waves are actively enhanced [227].

The interest in the phononic coupling of a 2LS is two-fold. Either for studying their impact on the coherence properties of the 2LS, or to investigate how phonon-assisted excitation schemes can provide benefits for the generation of quantum light. With this motivation in mind, optical phonons play only a minor role. On the one hand, it was shown that at low temperatures ($T < 4\,\mathrm{K}$) optical phonon interactions are negligible for dephasing effects and the acoustic phonon branch (being lower in energy) dominates [228, 229]. On the other hand, while early studies demonstrated LO phonon-assisted excitation approach for QDs [230–232] the sharp resonance of LO phonons [230] and a small interaction strength (governed by the Fröhlich-coupling) compared to the acoustic branch [195, 225, 233] shifted the interest to excitation schemes using acoustic phonons [234, 235].

Interactions between a 2LS and acoustic phonons can be mediated by either a deformation or a piezoelectric potential. The latter requires a non-centro symmetric crystal structure, dissimilar wavefunctions for the electron and hole, and, for typical QD sizes, is generally much weaker. The deformation potential describes how

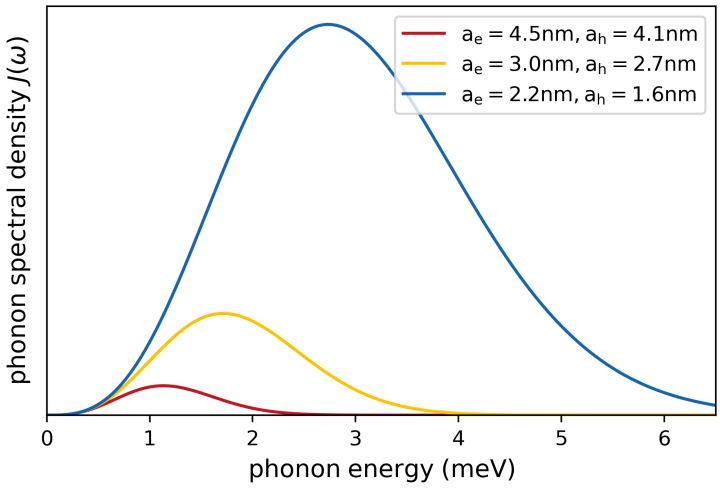


Figure 2.10: **Spectral density of the electron-phonon coupling** $J(\omega)$. Assuming spherically symmetric wave functions for the electron and hole, the exciton-phonon interaction is shown for three sets of radii as indicated by the legend. The spectral density is area normalized and follows from Eq. 2.2-39 using the following parameters: speed of sound $c_s = 4200\,\mathrm{m/s}$, density $\rho = 5680\,\mathrm{kg/m^3}$, electron potential $D_e = 3\,\mathrm{eV}$, hole potential $D_h = -6\,\mathrm{eV}$.

the exciton energy is shifted due to the lattice deformations associated with acoustic phonons, and dominates phonon-induced dephasing in QDs [195, 206, 207]. Primarily LA phonon modes contribute to the deformation potential.

Combining all these considerations, in this work, we will discuss only coupling constants γ_q that capture interactions with LA phonons. The interested reader is referred to Ref. [195] for a complete analysis of the various phononic couplings.

Approximating the phonons using bulk modes (justified when the lattice and dielectric properties of the QD and its host matrix are similar), the coupling can be factorized as a coupling matrix element depending on the material parameters and the electron and hole deformation potentials $D_{e/h}$, and a form factor given by wave functions of the electron and hole. It was shown that the exact shape of the wave function only slightly influences the phonon interaction spectrum [207], and in later experiments good agreement was obtained using spherically symmetric electron and hole wave functions [235]. Hence, for simplicity, we use ground-state solution wave functions of a spherically harmonic potential with a single parameter $a_{e/h}$ representing its radius.

Using the linear dispersion relation of LA phonons $\omega_q = c_s|\boldsymbol{q}|$, where c_s is the speed of sound, we finally express the exciton-phonon coupling as the spectral density

$$J(\omega) = \sum_{\boldsymbol{q}} \gamma_{\boldsymbol{q}} \delta(\omega - \omega_{\boldsymbol{q}}) = \frac{\omega^3}{4\pi^2 \rho \hbar c_s^5} \left(D_e e^{-\omega^2 a_e^2/(4c_s^2)} - D_h e^{-\omega^2 a_h^2/(4c_s^2)} \right) \quad (2.2\text{-}39)$$

with the mass density ρ. Exemplary spectral densities using typical QD parameters are shown in Fig. 2.10.

Returning to the steady-state populations of the laser-dressed states—but now including the LA phonon interaction—the results change significantly. The coupling between the dressed states of the *same* manifold is activated by $J(\omega)$, and it was shown to approximate a thermalization process well [236]. In other words, the exciton-phonon coupling $J(\omega)$ allows efficient population transfer from the high-energy dressed state $|\varphi_G\rangle$ to the low-energy dressed state $|\varphi_X\rangle$. The final exciton occupation reads as [234]

$$p_X^\infty = \frac{1}{2}\left[1 + \frac{\delta_L}{\Omega_{\text{eff}}}\tanh\left(\frac{\hbar\Omega_{\text{eff}}}{2k_BT}\right)\right] \tag{2.2-40}$$

and easily reaches $p_X^\infty > 0.8$. This observation already anticipates a scheme to transfer the QD from its ground to its excited state using the phonon coupling, as will be detailed in the following section.

From Eqs. 2.2-23, 2.2-29 and 2.2-38, we summarize the full Hamiltonian of the system that will be considered for the remainder of this work

$$H = H_0 + H_{QD-L} + H_{QD-P}\,. \tag{2.2-41}$$

2.2.4 Optical excitation

Since the advent of SPSs based on semiconductor QDs, a wide variety of optical excitation techniques have been proposed and demonstrated; for a general overview, see [71, 92–94, 237]. Although all have unique features, they can be broadly classified as coherent or incoherent driving schemes, and the choice between them may have significant implications for applications, as outlined in Sec. 2.1.4.3 and Sec. 2.2.2.3. Notably, not all coherent schemes require a laser pulse that resonantly addresses the target transition[17], and even incoherent schemes may produce highly indistinguishable photons [57, 216–218].

In the beginning of this section, two widely used schemes, resonance fluorescence (RF) in Sec. 2.2.4.1 and LA phonon-assisted driving in Sec. 2.2.4.2, are introduced in more detail, and later a brief comparison with other schemes and their benefits is given (Sec. 2.2.4.4). Sec. 2.2.4.3 discusses the phenomenon of re-excitation that is strongly related to the employed excitation method. It fundamentally limits the single-photon purity of QD sources and plays a central role in Ch.5.

2.2.4.1 Resonance fluorescence

In the simplest approach to generate quantum light from a 2LS—at least conceptually—the system is excited by a resonant laser field ($\omega_L = \omega_0$), and its resonance fluorescence signal is collected. This scheme is widely used in applications and provides a high system efficiency and indistinguishability [55, 58, 241, 242].

[17] For example, see dichroic excitation [238] or the Swing-UP of the quantum EmitteR population (SUPER) scheme [80, 239, 240]. For more details about excitations schemes, see Sec. 2.2.4.4.

Under continuous wave (CW) excitation, two regimes emerge that have been described in the seminal paper by Mollow for a 2LS [224]. For a weak-driving field $\Omega \ll \gamma_{\mathrm{rad}}$ (sometimes referred to as Heitler regime), the laser field is predominantly scattered coherently and the resulting linewidth is inherited from the cw laser such that it can be well below the FT limit [243]. While this scattered light exhibits antibunching ($g^{(2)}(0) < 0.5$), filtering it down to a subnatural linewidth impairs its single-photon nature [79, 244, 245]. Under strong driving $\Omega \gg \gamma_{\mathrm{rad}}$, the incoherent part of the spectrum becomes significant, two side bands appear symmetrically around the central peak at $\omega_L \pm \Omega_{\mathrm{eff}}$ and complete the well-known Mollow-triplett, which can be explained by the dressed-state picture introduced above (Sec. 2.2.3.2). Notably, the emission into the side peaks is time-correlated, indicating a two-photon scattering process from atoms [246] and QDs [247] that may even yield entanglement [248].

Since most applications of quantum light sources require clocked operation, pulsed excitation schemes achieving population inversion within a short laser pulse are generally preferred. In the absence of decoherence mechanisms, the system undergoes coherent oscillations (Rabi rotations) between $|G\rangle$ and $|X\rangle$ as a function of the interaction laser pulse area θ. For a 2LS initially in the ground state, the probability of finding it in the excited state after the laser pulse has fully interacted is

$$p_X(\theta) = \sin^2(\theta/2) \quad \text{where} \quad \theta = \int_{-\infty}^{\infty} dt\, \Omega(t)\,. \tag{2.2-42}$$

More generally, the state reads as

$$|\Psi_f\rangle = \sqrt{1 - p_X(\theta)}\,|G\rangle + e^{-i\phi}\sqrt{p_X(\theta)}\,|X\rangle\,. \tag{2.2-43}$$

Because the state is entirely described by θ (up to a phase ϕ set by the laser), this model is called the pulse area theorem and approximates the state well if the pulse length Δt_L is much shorter than the radiative decay time τ_{rad} of the 2LS. It is common to refer to pulses with $\theta = \pi, 2\pi, ..$ as π-, 2π-pulse and so on. Within the limits of the pulse area theorem, the first achieves perfect inversion, while the latter rotates the system back to its ground state.

Early demonstrations of Rabi rotations using QDs [53, 249] were strongly affected by decoherence that dampens the oscillations. Although various mechanisms contribute, excitation-induced decoherence due to acoustic phonon interactions was identified as the dominant source of noise at larger pulse areas [228, 229, 250, 251]. Later, a theoretical study predicted the reappearance of coherent oscillations if the pulse area is further increased [252], and the effect was recently observed [253]. Since the coupling strength to the phononic environment depends on the instantaneous effective Rabi frequency, $\Gamma_{\mathrm{phonon}} \propto J\,[\Omega_{\mathrm{eff}}(t)]$, exceeding the dressed-state splitting, which causes maximal phonon coupling for a significant duration of the pulse interaction lowers the induced decoherence (see Fig. 2.10 and Fig. 2.11).

For many applications, resonant π-pulse driving of the exciton level is the default scheme to generate high-quality single photons. Using cavity designs

with a polarization mode splitting that is significantly larger than the emission linewidth, π-pulse excitation enabled record system efficiencies >50 % [56] and recently up to 71 % [55] thanks to a state preparation fidelity close to unity. In addition, this driving scheme limits the introduced decoherence resulting in very high indistinguishabilities [56, 254] and phase-coherent generation of photon-number states [214].

On the other hand, most samples must be operated in a confocal microscopy arrangement (laser and quantum light share the same spatial mode), such that rejecting the reflected laser requires a cross-polarization setup that must achieve suppression factors of $10^6 - 10^8$. Generally, this suppression method also causes a single-photon loss of up to 50 %[18] and prohibits the generation of polarization-entangled cluster states using the solid-state spin [57, 90, 256]. Furthermore, this approach is incompatible with rough sample surfaces that can depolarize the scattered laser light [257]. Using an alternative sample geometry with a waveguide, the confocal layout can be avoided, which largely mitigates the need for a cross-polarization setup [60, 258, 259]. However, waveguide-based designs also achieve lower Purcell factors (see Sec. 2.2.5.2).

Finally, resonant population transfer is highly sensitive to fluctuations in laser wavelength and power [92], making the scheme less favorable for long or unsupervised operating periods.

2.2.4.2 LA phonon-assisted excitation

In the end of Sec. 2.2.3.3, an alternative, incoherent excitation scheme was briefly introduced. In contrast to RF, which aims to limit phonon interactions, this method actively incorporates phonons to enable dipole-forbidden intra-transitions within one manifold of the laser-dressed states.

Such an off-resonant scheme addresses two of the main drawbacks present in RF. First, the separation of scattered laser light from single photons can be easily accomplished in the frequency domain through the use of readily available, high-efficiency narrow bandpass or notch filters. Second, the width of the spectral density of the exciton-phonon coupling $J(\omega)$ (see Eq. 2.2-39 and Fig. 2.10) circumvents any strict resonance conditions and makes the process robust against laser fluctuations [260].

So far, we discussed only LA phonon-assisted excitation for CW driving[19], but again, most applications require clocked operation. Hence, we must understand if and how this pumping scheme works in the pulsed regime. Eq. 2.2-40 formulates that a high steady-state population of the excited state is expected, but this might require > 100 ps of laser interaction. In this case, and considering decay

[18] An asymmetric QD–cavity interaction allows to excite the system through the weakly coupled polarization mode while preferentially emitting into the orthogonal polarization mode of the cavity such that the single-photon loss can be reduced drastically [86, 255]. All works cited above that achieved record efficiencies use this special cavity configuration.

[19] Driving in the CW was regime was demonstrated for example in [261].

times $< 1\,\mathrm{ns}$, the simplified steady-state picture without radiative decay breaks down. Nonetheless, first studies focused on dynamical excitation found that exciton preparation fidelities can exceed 95 % for pulses as short as 20 ps [234]; Experimental confirmation followed shortly after [178, 235]. Both theoretical and experimental works have already noted the intricate interplay between pulse length Δt_L, laser–QD detuning δ_L, and pulse area θ, and their effect on the exciton population.

Given the instrumental role of the pulsed LA driving scheme for this thesis, a more detailed description of the excitation process is presented here. Fig. 2.8 depicts the energy levels involved in three different bases, which might facilitate the understanding in different pictures and how they can be linked. These include the bare QD states, the states in the rotating frame where the ground state is higher in energy due to $\delta_L > 0$, and the laser-dressed states with the level splitting $\hbar\Omega_{\mathrm{eff}}$. Naturally, the latter two are only defined for the presence of a laser field, and the rotating frame approximates an eigenbasis only in the limit of a vanishing interaction strength. To capture the dynamics of the process, we will use the dressed states picture exclusively from this point forward. Importantly, the effective Rabi frequency is now time-dependent, and consequently, the level splitting and phonon scattering rate also change dynamically, as shown in Fig. 2.11 for a laser pulse peaking at t_0. The color coding of the dressed-state levels estimates their relative population[20], which can be inferred from the phonon scattering rate obtained using Fermi's Golden rule [262]

$$\Gamma_{\mathrm{phonon}}(t) = \frac{\pi}{4}\frac{\Omega(t)}{\Omega_{\mathrm{eff}}(t)} J\left[\Omega_{\mathrm{eff}}(t)\right] . \tag{2.2-44}$$

To elucidate how the phonon coupling and level populations are influenced by the driving strength, the evolution is displayed for three pulse areas θ. For low field amplitudes, the dressed state splitting is relatively low throughout the pulse, failing to activate sufficiently strong phonon interactions, and thus the population inversion remains incomplete. Increasing the pulse area while maintaining the same pulse length requires a stronger driving field, which in turn increases the phonon coupling and, consequently, also the lower dressed state population. However, once the dressed-state level splitting exceeds the maximum of the spectral density $J(\omega)$ (see Fig. 2.10), the QD is increasingly decoupled from the phonoic environment; the phonon scattering rate is reduced and exhibits a local minimum at the peak of the laser pulse (see Fig. 2.10f). In a first approximation, the occupation of the excitonic dressed state $|\varphi_X\rangle$ follows from the integrated phonon scattering rate. So even if a pulse induces a level splitting that overshoots the optimal phonon coupling, there are still considerable exciton-phonon interactions during its rise and fall time, and only excessive "overdriving" reduces the preparation fidelity significantly. To optimize the exciton population for increasing pulse lengths, the pulse area must also be increased to maintain a similar Rabi frequency.

[20] For now, spontaneous decay is neglected, which is well-justified for $\Delta t_L \ll \tau_{\mathrm{rad}}$.

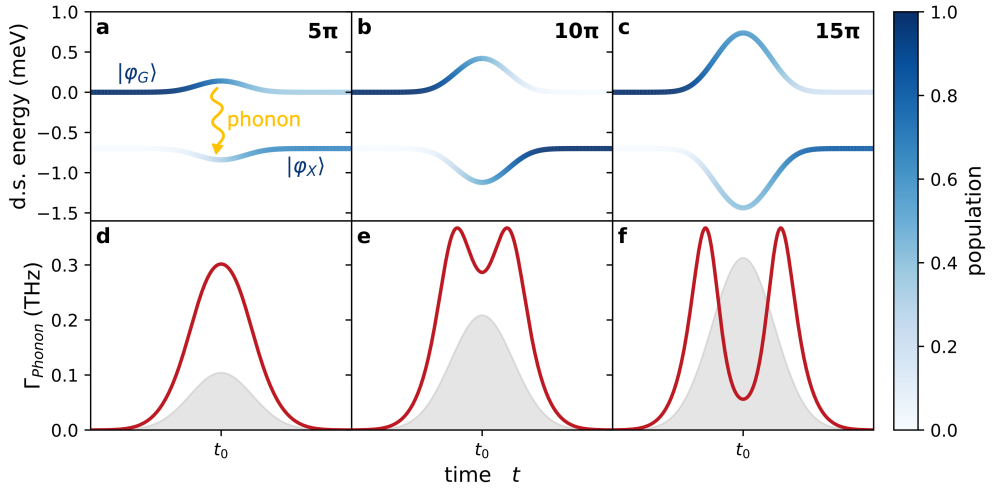


Figure 2.11: **Dynamic LA phonon coupling**. For a laser pulse with an intensity FWHM $\Delta t_L = 10\,\text{ps}$ and detuning $\hbar\delta\omega_L = 0.7\,\text{meV}$, the dressed state energies are shown in the top panels (**a-c**) for increasing pulse area. The color coding indicates the approximate level occupation based on integration over the instantaneous phonon scattering rate $\Gamma_{\text{phonon}}(t)$ shown in the lower panels (**d-f**). This approximation neglects spontaneous emission and Γ_{phonon} is computed using Fermi's Golden Rule (see Eq. 2.2-44). The laser pulse is indicated in gray as a reference for the timing and pulse area. $J(\omega)$ is calculated using Eq. 2.2-39 for the following parameters (red curve in Fig. 2.10): speed of sound $c_s = 4200\,\text{m/s}$, density $\rho = 5680\,\text{kg/m}^3$, electron radius $a_e = 4.5\,\text{nm}$, hole radius $a_e = 4.1\,\text{nm}$, electron potential $D_e = 3\,\text{eV}$, hole potential $D_h = -6\,\text{eV}$.

In this comparison, the detuning $\delta\omega_L$ is fixed; however, in general, it provides an additional tuning knob alongside laser power and pulse length. The joint optimization of all three pumping parameters varies between each QD[21] and is not always straightforward, especially if more than one figure of merit must be maximized, most often system efficiency η_{sys} and single-photon purity $\mathcal{P}$ [1]. These optimization challenges will be explored in depth in Chapter 4, and we show that the spectral width of $J(\omega)$ renders this excitation scheme remarkably insensitive to small deviations from the optimal parameter set. Hence, the photon number distribution varies little even over a notable parameter range [1, 260], making LA driving ideally suited for robust implementations of cryptographic protocols. Finally, Gaussian pulses were shown to approximate the ideal laser pulse shape well [263] since they naturally fulfill the adiabaticity condition[22] required to transfer the population of the excitonic dressed state efficiently to the bare exciton [95, 263].

[21] Most prominently, the phononic coupling $J(\omega)$, radiative decay time τ_{rad}, and constraints on the spectral filtering will affect the ideal operation point.

[22] Adiabaticity is generally given if the envelope of the Ω_{eff} changes slowly compared to $\Omega(t)$ which might be expressed as $d\Omega_{\text{eff}}(t)/dt \ll \Omega(t)$. In experiments, laser pulses $\Delta t_L \gtrsim 20\,\text{ps}$ performed best [264].

A study found that the system efficiency of the LA scheme reaches 85 % of the RF counterpart while providing a higher single-photon purity. The mean wavepacket overlap in this direct comparison was 91 % for both schemes [264]. These results indicate that the dephasing additionally induced by phonon-assisted excitation is small compared to residual noise processes also present for RF; this inference is further supported by Ref. [217], but an experimental comparison using QDs with even lower dephasing noise is still outstanding. Recently, indistinguishabilities exceeding 90 % were also reported for QDs emitting in the telecom C-band using LA phonon-assisted driving [218, 265].

However, although the TPI visibility can be almost as good as for RF, the phase between ground and excited state is different for every pulse due to the interaction with the phonon bath. Hence, the electronic coherence of the state is diminished, resulting in the absence PNC of the photonic state [219]. In cryptographic applications, this phonon-induced phase randomization provides an intrinsic defense against discrimination attacks (see Sec. 2.1.4.3) that would otherwise require active phase scrambling. Finally, compared to other off- or quasi-resonant schemes such as excitation via the wetting layer or higher excited states (p-shell), LA phonon-assisted driving features higher preparation fidelities and significantly shorter relaxation times into the radiative level [217], which reduces emission timing jitter and photon distinguishability.

2.2.4.3 The re-excitation phenomenon

All pulsed schemes use *short* laser pulses to excite the QDs. Here, *short* is defined in relation to the spontaneous decay time, $\Delta t_L \ll \tau_{\mathrm{rad}}$. Then, one assumes the probability of emitting a photon *during* the laser pulse to be negligible, just as in the pulse area theorem for RF stated in Eq. 2.2-42. However, when striving for a high-performance SPS, even minor deviations from the ideal behavior must be considered.

If a photon is spontaneously emitted while the laser pulse still interacts with the QD, there is a chance to excite the system a second time with the same laser pulse, so that a second photon is generated when the QD decays again shortly after. Hence, this process—referred to as re-excitation [266] and schematically depicted in Fig. 2.12—gives rise to multiphoton noise, even for an ideal 2LS[23]. The effect was first studied in detail for RF [96] using a photon-counting technique on quantum trajectories, and an extended analysis was presented shortly after in [267]; the key concept is briefly recapitulated here.

Spontaneous emission can be treated as a stochastic process that occurs at any time with rate $\gamma_{\mathrm{rad}}\, p_X$, where p_X is the occupation probability of the excited state. To understand how p_X evolves in time, it becomes necessary to evaluate the pulse area that interacts with the QD up to a time t (assuming a laser pulse starting at

[23] As defined earlier, whenever two or more photons are generated from a laser pulse, it is considered a multiphoton event, although the photons are not emitted in the same mode.

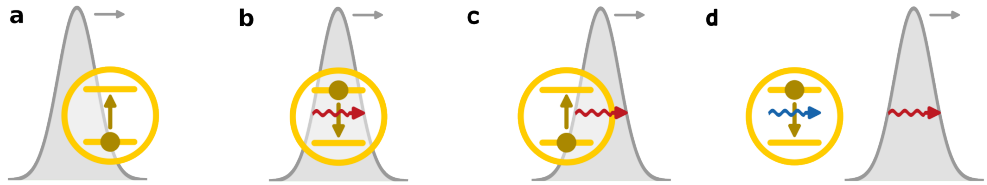


Figure 2.12: **Schematic time series of a re-excitation event.** Modeling the QD as a two-level system, the re-excitation process from a single laser pulse is described using a representative quantum trajectory, broken down into four distinct events. (**a**) Early during the interaction with the laser pulse, the QD is excited and (**b**) decays shortly after, emitting a first single photon (red arrow). (**c**) Towards the end of the interaction with the laser pulse, the QD is excited a second time. (**d**) The laser pulse and the first photon have passed when the QD spontaneously decays again, emitting another photon (blue arrow). Thus, the two photons are emitted consecutively in a strict time-ordered manner. The figure has been adopted from the publication presented in Chapter 5 under CC-BY-4.0 license [2].

$t = 0$ and ending at $t = T$)

$$A(t) = \int_0^t dt' \, \Omega(t') \tag{2.2-45}$$

instead of only considering the total pulse area $A(T) = \theta$. In cases where the first photon is emitted at t_1 during the pulse ($0 < t_1 < T$), the probability follows as

$$f_1(t_1) \approx \gamma_{\mathrm{rad}} \, p_X\left(A(t_1)\right) \, \mathrm{e}^{-\gamma_{\mathrm{rad}} t_1/2} = \gamma_{\mathrm{rad}} \sin^2\left(\frac{A(t_1)}{2}\right) \mathrm{e}^{-\gamma_{\mathrm{rad}} t_1/2}, \tag{2.2-46}$$

where the factor $\mathrm{e}^{-\gamma_{\mathrm{rad}} t_1/2}$ accounts for the probability that no photon was emitted before t_1. If the photon was emitted after the pulse ($t_1 > T$)

$$f_1(t_1) \approx \gamma_{\mathrm{rad}} \sin^2\left(\frac{A(\infty)}{2}\right) \mathrm{e}^{-\gamma_{\mathrm{rad}}(t_1 - T/2)}, \tag{2.2-47}$$

and there will be no re-excitation. However, for the first case, the emission projects the system in the ground state and one can think of it as "restarting" the pulse area theorem at t_1, where the probability of finding the QD in the excited state at time t_2 depends on the pulse area that is "left" after t_1, that is, $p_X\left(A(t_2) - A(t_1)\right)$.

Extending this formalism to $f_n(t_1, ..., t_n)$ (see [267]), any inclusive n photon probability F_n can be computed by integrating over all time arguments

$$F_n = \int_0^\infty dt_1 \ldots \int_{t_{n-1}}^\infty dt_n \; f_n(t_1, ..., t_n) \, . \tag{2.2-48}$$

The inclusive probabilities F_n consider all events with *at least* n photons, whereas p_n reflects only events with *exactly* n photons. Using

$$p_n = F_n - F_{n+1} \tag{2.2-49}$$

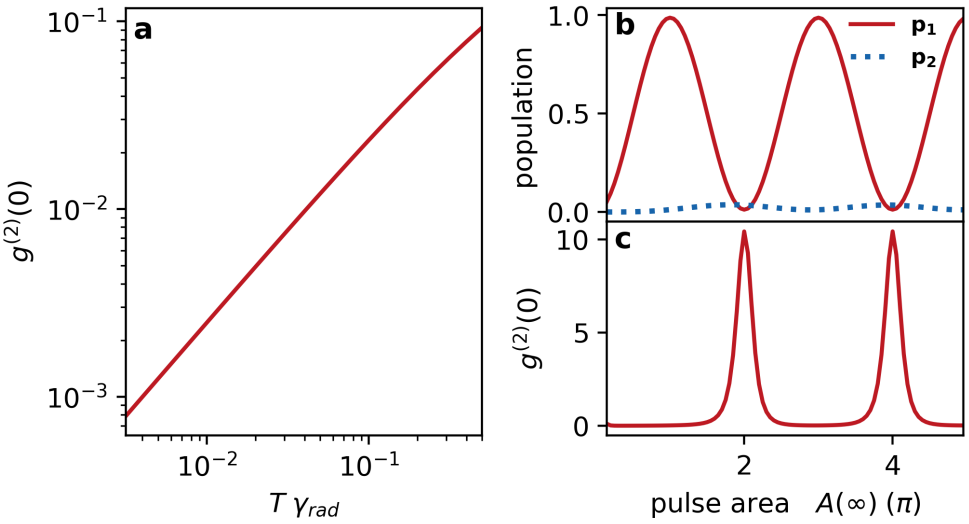


Figure 2.13: **Multiphoton probability of an ideal 2LS under RF.** For a perfect 2LS, the multiphoton contribution is solely caused by re-excitation, which can be approximated well using the photon-counting technique described in the text for laser pulses shorter than the decay time, $T\,\gamma_{\mathrm{rad}} < 1$. Here, T is assumed to include the full pulse area $A(T) \approx A(\infty)$, whereas $A(\Delta t_L) \approx 3/4\,A(\infty)$ for a Gaussian pulse. (**a**) Tuning the ratio $T\,\gamma_{\mathrm{rad}}$ shows a linearly increasing $g^{(2)}(0)$ for $T\,\gamma_{\mathrm{rad}} < 0.1$, while for longer pulses the $g^{(2)}(0)$ slowly deviates from the trend. $g^{(2)}(0)$ is calculated using Eq. 2.2-17 assuming $p_{i>2} = 0$. (**b**) The photon number populations calculated for $T\,\gamma_{\mathrm{rad}} = 0.1$ according to Eq. 2.2-49 oscillate as a function of the pulse area with a period of 2π. Interestingly, p_2 reaches its maximum not for $A(\infty) = 2\pi$ but slightly before. (**c**) The drastic change in $g^{(2)}(0)$ from strongly antibunched at $A(\infty) \approx \pi$ to strongly bunched at $A(\infty) \approx 2\pi$ illustrates how the photon statistics of the quantum light can be controlled by the pulse area of the driving laser.

all photon number probabilities and hence, η_{sys} (Eq. 2.2-6) and $g^{(2)}(0)$ (Eq. 2.2-17) can be calculated for arbitrary pulse areas and ratios of pulse length to decay time.

For a π-pulse, Fig. 2.13a shows the $g^{(2)}(0)$ for the short to medium pulse length regime $T\,\gamma_{\mathrm{rad}} < 1$, for which the approach agrees well with the quantum regression or similarly complete models [267]. However, it fails to capture the saturation behavior $g^{(2)}(0) \approx 1$ for very long pulses that signify Poissonian statistics. In the short pulse limit, the behavior can be further simplified [267]

$$g^{(2)}_{\mathrm{2LS}}(0) \propto T\,\gamma_{\mathrm{rad}}\,. \tag{2.2-50}$$

These results motivated investigations of 3LS, where two cascaded spontaneous emission processes significantly delay the return to the ground state and, consequently, reduce the multiphoton noise considerably. Indeed, for a cascaded 3LS [97],

$$g^{(2)}_{\mathrm{3LS}}(0) \propto T\,\gamma^{X}_{\mathrm{rad}}\;T\,\gamma^{XX}_{\mathrm{rad}} \tag{2.2-51}$$

and the lowest $g^{(2)}(0)$ values to date are measured using this level structure [97, 187][24].

Although it seems clear from these approximations that the shortest possible pulse produces the fewest multiphoton events from a 2LS, the situation is more complicated for a real QD. For one, decreasing the pulse length necessarily increases its spectral width, and the laser might induce other transitions [268, 269] (i.e., a 2LS model is no longer valid), which can contribute to the multiphoton background. For another, these spectrally broad pulses are significantly more difficult to suppress in a cross-polarization setup [257].

Returning to pulse lengths for which the 2LS approximation holds well, the quantum trajectory framework also captures the photon bunching observed for $A(\infty) = 2\pi, 4\pi, ...$ [96, 220] as shown in Fig. 2.13b. For most quantum trajectories, the laser drive rotates the QD back to the ground state, but whenever a first photon was emitted during the pulse, there is still some pulse area available to induce a second photon.

So far, all theoretical and experimental investigations of re-excitation have been concerned with resonant π-pulse driving. While the fundamental mechanism of creating a second exciton from one pulse is expected to persist also for other pumping schemes, there are essential differences that might change the nature of the two-photon response. The first study of re-excitation for an off-resonant scheme—specifically, for LA phonon-assisted driving—is presented in Chapter 5.

2.2.4.4 Spoiled for choice: finding the ideal excitation scheme

The impressive variety of demonstrated excitation schemes can make choosing the *ideal* approach for any given situation a daunting task. Throughout this section, methods other than RF and LA excitation have already been mentioned in passing; however, a more structured overview is still missing and is provided here for the reader's convenience. All possible combinations of requirements, applications, and experimental realities make it nearly impossible to compile a comprehensive guide; instead, the more common schemes are briefly summarized, along with a selection of key elements to be taken into account.

In addition to the previously defined figures of merit of η_{sys}, $\mathcal{P}$, TPI visibility and the presence of PNC (see Sec. 2.2.2.3), experimental complexity, potential entanglement generation, and noise robustness of each scheme affect the evaluation. Depending on the application, these metrics might be weighted differently. For instance, a QKD experiment might benefit from higher single-photon rates even if they come at the price of increased multiphoton noise, whereas a mistrustful primitive such as quantum bit commitment prioritizes a high single-photon purity. To further complicate the optimization, the security analysis depends on channel

[24] Recently, [192] reached similar low values using RF by applying narrow spectral filtering on the order of the emission linewidth γ_{rad}.

loss and might yield different outcomes even for the same primitive when evaluated for increasing communication distances [1]. While these protocols are agnostic to photon distinguishability (as long as their spectral signature is not correlated with the encoding degree of freedom), interference-based applications, including boson sampling [241], or cluster state generation based on fusion gates [58], must factor the TPI visibility constraints in. Yet others, including twin-field QKD [102, 142] or photon-number state engineering [214] require all of the above but also PNC.

Above-band pumping is typically only used for characterization purposes because of its simple implementation but poor quality of generated quantum light. Both CW and pulsed implementations are possible. The delocalized exciton state is created in the wetting layer and may eventually relax to the lowest energy state of the quantum dot. The laser is far blue detuned and easily filtered, but the probabilistic nature of the relaxation reduces the state preparation fidelity and causes a large temporal uncertainty when the excited state is occupied, strongly limiting photon indistinguishability.

Similarly, excitation through higher-order states (mostly **p-shell**) introduces significant jitter and impairs the indistinguishability to various degrees [217, 270]. Otherwise, it generally outperforms above-band pumping in η_{sys} and $\mathcal{P}$ while still enabling spectral laser suppression. On the other hand, the p-shell state is resonantly driven, which makes the scheme more sensitive to laser fluctuations.

Adiabatic rapid passage (ARP) [94, 271, 272] is another scheme (in addition to the incoherent phonon-assisted approach) that offers robustness while preserving high performance; it has been derived from stimulated Raman adiabatic passage (STIRAP) techniques [273]. By chirping the laser pulse, the frequency is effectively swept over the QD resonance to exploit the avoided crossing of dressed states while maintaining an intensity high enough to ensure that the system follows the instantaneous eigenstate of the Hamiltonian. Usually, negatively chirped pulses are used so that the system follows $|\varphi_X\rangle$ (sometimes referred to as lower dressed state), which preserves a high population better than $|\varphi_G\rangle$ since phonon absorption is much less likely than emission at low temperatures. The low excitation-induced decoherence expected for ARP [272] was confirmed by achieving a record high indistinguishability of 99.5 % [216]. Due to the broad frequency sweep, the scheme is not overly sensitive to the exact resonance of a QD and a route to collectively exciting multiple QDs was proposed [274]. However, the spectral broadness also makes suppressing the laser more challenging.

The recently introduced **SUPER scheme** [239], employs two differently red-detuned laser pulses to create a beating and promises to combine coherent state preparation with high fidelity and the advantages of spectral filtering. To date, however, experiments show varied results, where the single-photon count rate was sometimes close to the rate obtained using resonant two-photon excitation (TPE)

[240, 275] but in other implementations, the count rate, single-photon purity, and TPI visibility (significantly) remained lower than for RF [276] or for LA excitation [80]. A study suggests that the phonon coupling could significantly affect the performance [277], offering a possible explanation. From an experimental perspective, the scheme adds considerable complexity by requiring two phase-locked lasers where their relative delay, their intensities, and detuning with respect to the QD must be finely adjusted, which also makes it susceptible to variations in any of those parameters.

Another set of methods revolves around preparation of the QD in the biexciton state using **TPE**, first demonstrated using laser pulses resonant to half the biexciton energy [278, 279]. An important prerequisite for this technique is the presence of a neutral exciton; this often implies that the QD must be embedded in a diode with tunneling barrier to allow charge tuning [81, 82, 184]. The biexciton decays through a radiative cascade back to the ground state by emitting two nondegenerate photons which may be entangled in polarization [54, 173] (also see Sec. 2.2.1 2.2.1). The scheme exhibits the same pump parameter sensitivity as RF, but due to the nonzero binding energy, the laser wavelength falls between the two single-photon wavelengths and can be filtered spectrally. While using this scheme, η_{sys} may be as high as for RF, and $\mathcal{P}$ even exceeds the values achievable using RF due to the reduced re-excitation probability (see Sec. 2.2.4.3), the (unfiltered) photon indistinguishability is fundamentally limited by the ratio of the radiative lifetimes of X^0 and XX [280]. Without a specially tailored cavity featuring frequency-selective Purcell enhancement, the TPI visibility is typically $< 70\,\%$ [280, 281]. Nonetheless, this scheme recently enabled the first teleportation experiment between two remote quantum dots, where frequency matching was achieved with quantum frequency conversion [76] or magnetic and strain tuning [176].

A variation of TPE called **stimulated two-photon excitation (sTPE)** [215, 282] addresses the drawback of limited TPI visibility by adding a second laser pulse to stimulate the $|XX\rangle \rightarrow |X^0\rangle$ transition. At first glance, this might seem to only overcomplicate the generation of a single photon from the $|X^0\rangle \rightarrow |G\rangle$ transition unnecessarily; however, the approach restores high indistinguishabilities by eliminating any temporal correlations with the XX photon and, at the same time, enhances the emission properties compared to RF in several ways. First, due to optical selection rules, the single photon and stimulating pulse are co-polarized, thus allowing one to channel the stream of single photons in a preselected, linearly polarized mode and circumvent the $50\,\%$ loss associated with randomly polarized output. Second, laser suppression can be performed in the frequency domain, simplifying the experimental setup, and third, sTPE inherits a reduced re-excitation probability common to cascaded emission processes [97]. These benefits come at the cost of increased complexity, as two laser pulses with carefully adjusted temporal delays are necessary, and an uncharged ground state must be available. Finally, quantum light produced via sTPE was shown to feature PNC if both pulses are

phase-locked [215].

Lastly, most of the above schemes that prepare the QD in $|X^0\rangle$ can also be modified to prepare $|XX\rangle$ instead. The purpose is to benefit simultaneously from the features of the respective excitation method and the (entangled) two-photon output state provided by the XX–X cascade. Specifically, there are demonstrations of ARP TPE [94, 283], SUPER TPE [180, 240] and LA phonon-assisted TPE [178, 179].

This overview of benefits and drawbacks of the various excitation schemes did not expand on the specifics that arise when the QD is coupled to a photonic cavity. In particular, narrow-band cavities can significantly impact the efficiency with which the laser couples to the QD, while more elaborate cavity designs might mitigate some of the challenges associated with a particular scheme. For example, purposely induced polarization mode splitting as briefly discussed in Sec. 2.2.4.1 can funnel single-photon emission almost exclusively into one polarization mode, thereby increasing efficiency.

The following section details the most common approaches to enhance the performance of a SPS using photonic structures, and indicates aspects that might influence the choice of pumping scheme were appropriate.

2.2.5 Photonic extraction structures

Close to unit system efficiencies require not only high single-photon generation efficiency but also that the collection of those photons is optimized (see Eq. 2.2 -4). However, semiconductor QDs are made from high-index materials and for a bare sample only $\approx 1\,\%$ of the emission escapes through the top surface and can be collected with a high NA objective. To improve the situation, photonic structures that enhance outcoupling and ideally funnel all photons into a single Gaussian mode in the far field are integrated around the QD. Single-photons propagating in a well-defined Gaussian mode can then be easily collected by a lens or directly coupled to SMF. This challenge has been addressed by a number of approaches that may be divided into three general strategies:

I The emission characteristic of the QD is unchanged, but the semiconductor to vacuum interface is modified to reduce total internal reflection and the divergence angle. Common examples are solid-immersion lenses or microlenses [284–286].

II Modifying the QD environment to inhibit spontaneous emission into any *unwanted* modes. This can be achieved using strong mode confinement provided by waveguides [258, 287–291].

III Instead of suppressing the decay into unwanted modes, the emission into the target mode is accelerated by Purcell enhancement [87] created by photonic resonators or cavities [51, 57, 58, 80, 173, 259, 292–299].

The first strategy is the simplest to implement but also the least efficient. Since the emission behavior of the QD is not altered, at best, all photons that are radiated towards the upper half-space can be collected. The limitation may be addressed by combining the lens with a simple planar cavity, which has almost no Purcell enhancement but helps to direct all photons towards the sample's surface [300].

The other two strategies directly change the emission rates into various modes. For the purpose of this overview, it is sufficient to consider only two rates, namely emission into the target mode and into all other modes. The β-factor measures the fraction of photons that can be found in the target mode

$$\beta = \frac{\gamma_{\text{target}}}{\gamma_{\text{target}} + \gamma_{\text{other}}}, \tag{2.2-52}$$

and illustrates the two pathways to increase β: either γ_{other} is suppressed (strategy II) or γ_{target} is enhanced (strategy III).

2.2.5.1 Photonic waveguides

The second strategy is often implemented using waveguides that create strong single-mode confinement, thereby reducing the density of states (DOS) of all modes that are otherwise available in a bulk crystal. Following Fermi's Golden Rule, a lower DOS causes a lower transition rate for these modes.

When the QD sits within a waveguide that is directed out of the sample plane, the structures are also called nanowires [287–290] and often feature a vertical taper that transforms the waveguide mode to a single-mode Gaussian in the far field, thus allowing high outcoupling efficiencies. Both top-down [288] and bottom-up [289, 301] approaches have been demonstrated, and a bottom reflector helps to redirect photons toward the top, further enhancing the collection [288, 301]. The broadband operation of the waveguide easily accommodates the spread of emission wavelengths common to QD growth, and no post-fabrication tuning is necessary. In bottom-up growth, the spatial alignment of the QD in the center of the waveguide is also intrinsically ensured [301].

Alternatively, the waveguide may be oriented in-plane, which is often realized as photonic crystal waveguide (PhC WG) [291, 298, 299]. In this context, the array of periodically alternating refractive indices consists of the high-index semiconductor and a grid of etched air cylinders with a spacing and diameter that are carefully tuned to the target wavelength. Bragg scattering between these interfaces controls light propagation, which can be described with optical Bloch modes. For pronounced Bragg interference, and by embedding the PhC WG in a membrane to reduce the system's dimensionality, a photonic band-gap develops that strongly suppresses the local DOS. Waveguides can be created by omitting a row of nanoholes in the air-cylinder etching pattern and using this design, β-factors close to unity are viable [259]. Typically, grating couplers are used to efficiently extract photons

from the structure[60, 259, 302], and the waveguides can be designed as single-sided structures to ensure emission in the direction of the coupler. Notably, most PhC WGs also locally engineer the DOS to achieve some Purcell enhancement as detailed in the next section. For more details on PhC WGs, see Ref. [303] and for an extensive review of these nanostructures for quantum light applications, see Ref. [258].

2.2.5.2 Photonic cavities

The third strategy aims to maximize β by increasing the decay rate of the target mode γ_{target} based on the Purcell effect [87]. The rate enhancement is quantified by the Purcell factor

$$F_P = \frac{\gamma_{\text{target}}}{\gamma_{\text{rad}}}\,, \tag{2.2-53}$$

where γ_{rad} is the radiative decay rate in bulk. For most cavity designs, where no special measures are taken to reduce γ_{other}, one can use $\gamma_{\text{other}} \approx \gamma_{\text{rad}}$ to approximate

$$\beta \approx \frac{F_P}{1+F_P}\,, \tag{2.2-54}$$

which indicates that $\beta \to 1$ for $F_P \gg 1$.

Before turning toward the calculation of the Purcell enhancement, we briefly recall the interaction between a 2LS and a single-mode optical resonator. It can be characterized using three rates:

i The 2LS decay rate in bulk γ_{rad}.

ii The rate at which photons leave the cavity, often called the cavity loss rate κ.

iii The coupling between the single-mode field of the resonator and the 2LS g_{cav}.

In analogy to the coupling between a 2LS and a monochromatic laser field (see Sec. 2.2.3), the interaction strength g_{cav} can be inferred using the dipole interaction with the vacuum field of the cavity

$$g_{\text{cav}} = \frac{|\boldsymbol{d}_{XG}\cdot\boldsymbol{E}_{\text{vac}}|}{\hbar} = \sqrt{\frac{d_{XG}^2\,\omega_0}{2\varepsilon_0\,\hbar\,V}}\,, \tag{2.2-55}$$

where ω_0 is again the frequency of the 2LS and V is the mode volume of the cavity. For a cavity within a semiconductor, the calculation of V is beyond the scope of this introduction, but details for a self-consistent evaluation of the Purcell enhancement can be found in Ref. [304]. When the excitation is lost fast compared to the strength of the light-matter coupling enhanced by the cavity $g_{cav} \ll (\kappa, \gamma_{\text{rad}})$, the system is in the weak-coupling regime where dissipative processes dominate. The photon may be lost by leaving the cavity or by emission from the 2LS into a mode that is not resonant with the cavity. On the other hand, when the exchange interaction is strong

compared to the loss channels (strong-coupling regime), the eigenstates hybridize in a superposition of photonic and electronic excitation, forming a quasi-particle called polariton. This regime has been achieved in various QD–cavity systems [294, 297, 305–307] and causes vacuum Rabi splitting, single-photon nonlinearities, or cavity-mediated entanglement, but is less suited for the efficient outcoupling of single photons. Hence, the following focuses on the weak-coupling regime, where the influence of the cavity vacuum field on the 2LS can be treated as a perturbation. Then, Fermi's Golden Rule allows one to calculate the emission rate in the target mode. What follows partially retraces the derivation found in Ref. [112].

Assume a cavity with only a single mode at ω_c that is close to the 2LS transition frequency ω_0, and with a spectral DOS $h(\omega)$ with FWHM $\Delta\omega_c$. Using a Lorentzian cavity mode

$$h(\omega) = \frac{2}{\pi\Delta\omega_c}\frac{{\Delta\omega_c}^2}{4(\omega-\omega_c)^2+{\Delta\omega_c}^2} \tag{2.2-56}$$

that obeys single-mode normalization $\int_0^\infty d\omega\, h(\omega) = 1$, we can evaluate the DOS at the emitter frequency $\omega = \omega_0$. For a resonant cavity $\omega_c = \omega_0$, this expression reduces to

$$h(\omega_0) = \frac{2}{\pi\Delta\omega_c} = \frac{2\,Q}{\pi\omega_0}\,, \tag{2.2-57}$$

where the cavity quality factor

$$Q = \frac{\omega_c}{\Delta\omega_c} \tag{2.2-58}$$

was introduced. In addition to the DOS, the transition matrix element must be calculated to evaluate Fermi's Golden Rule

$$\Gamma = \frac{2\pi}{\hbar^2}|M_{XG}|^2 h(\omega)\,. \tag{2.2-59}$$

For dipole interaction, and assuming parallel orientation of dipole and cavity field, the matrix element is

$$|M_{XG}|^2 = \langle \boldsymbol{d}_{XG}\cdot\boldsymbol{E}_{\mathrm{vac}}\rangle^2 = d_{XG}^2\,E_{\mathrm{vac}}^2 = \hbar^2 g_{\mathrm{cav}}^2 \tag{2.2-60}$$

such that the emission rate into the target mode follows from Eqs. 2.2-56, 2.2-59, and 2.2-60 as

$$\gamma_{\mathrm{target}} = 2\pi g_{\mathrm{cav}}^2\,\frac{2}{\pi\Delta\omega_c}\frac{{\Delta\omega_c}^2}{4(\omega-\omega_c)^2+{\Delta\omega_c}^2}\,. \tag{2.2-61}$$

On resonance, the rate simplifies to

$$\gamma_{\mathrm{target}} = 4g_{\mathrm{cav}}^2\frac{Q}{\omega_0} = \frac{2d_{XG}^2}{\varepsilon_0\hbar}\,\frac{Q}{V}\,. \tag{2.2-62}$$

Similarly, the transition matrix element for a decay into a free-space mode is calculated using a large volume V_0 and accounting for a random orientation of dipole and field,

$$|M_{XG}^{\mathrm{bulk}}|^2 = \frac{1}{3}\frac{d_{XG}^2\omega_0}{2\varepsilon_0\hbar V_0} \tag{2.2-63}$$

and using the photonic DOS in free space $h(\omega) = \frac{\omega^2 V_0}{\pi^2 c^3}$, the bulk decay rate follows again from Eq. 2.2-59

$$\gamma_{\text{rad}} = \frac{d_{XG}^2}{3\varepsilon_0 \hbar \pi} \frac{\omega_0^3}{c^3} . \tag{2.2-64}$$

Finally, the Purcell factor is calculated from Eqs. 2.2-53, 2.2-62 and 2.2-64

$$F_P = \frac{3}{4\pi^2} \frac{Q}{V} \left(\frac{\lambda_0}{n} \right)^3 , \tag{2.2-65}$$

where λ_0 is the emitter wavelength in vacuum and n is the refractive index of the cavity medium. Assuming perfect reflectivity for the cavity mirrors, the cavity loss rate approaches the FWHM of the cavity mode $\kappa \approx \Delta\omega_c$ and the Purcell factor can be re-expressed using the three parameters describing the cavity-QD interaction[25]

$$F_P = \frac{4 g_{\text{cav}}^2}{\kappa \gamma_{\text{rad}}} . \tag{2.2-66}$$

Eq. 2.2-65 illustrates that for a given wavelength, the enhancement is strongest for large quality factors and small cavity volumes, while Eq. 2.2-66 indicates that for a given bulk decay rate, the cavity loss rate must be reduced or the coupling strength increased to boost F_p.

In the experiment, the Purcell factor is inferred from time-correlated photon counting (TCPC) histograms, measuring the bulk and cavity-accelerated $1/e$-decay times separately. The bulk decay time is usually averaged from an ensemble of similar emitters or, if possible, by detuning the emitter far from the cavity. Assuming sufficient temporal resolution of the detection system and a predominantly radiative decay, the corresponding rate is simply $\gamma_{\text{rad}} = \tilde{\tau}_{\text{rad}}^{-1}$ (the tilde indicates the measured value). Although in principle γ_{target} cannot be measured in the same fashion[26], the distortion caused by the cavity can often be neglected since $\kappa \gg \gamma_{\text{target}}$. Only when the measured decay trace clearly deviates from a monoexponential curve does a more careful analysis become necessary.

Over the past two decades, various cavity designs have been realized with QDs, and the most common ones will be briefly introduced here. Spatial and spectral overlap of the QD dipole and cavity mode is key to any design and fabrication process [85]. Typically, structures with large Q-factors are more sensitive to spectral mismatch, and those with low mode volume V are more susceptible to spatial misalignment.

[25] In the strong-driving regime, this expression is often also called the cooperativity.

[26] The photons are first emitted into the cavity mode at rate γ_{target} but leave the cavity on average only $1/\kappa$ later. Hence, the decay trace measured outside the cavity is the convolution of two exponentials instead of a monoexponential (the product of two Lorentzians is no longer a Lorentzian). However, for most systems $\kappa \gg \gamma_{\text{target}}$, such that the approximation $\gamma_{\text{target}} \approx \tilde{\tau}_{\text{cav}}^{-1}$ holds reasonably well, where $\tilde{\tau}_{\text{cav}}^{-1}$ is the measured decay time in the presence of a cavity.

Micropillar

Along the growth direction, the mode is confined using a bottom DBR and a top DBR with slightly lower reflectivity. For lateral confinement, the design leverages the high refractive index contrast between the semiconductor and vacuum. The sample is etched several µm deep and only a cylindrical pillar encompassing the QD with a diameter of 2 µm to 4 µm is left standing. Even smaller diameters would produce smaller mode volumes, but surface effects, such as sidewall roughness, start to strongly affect cavity performance [308]. Micropillar cavities require pre-characterization to select suitable emitters and deterministic fabrication based on position markers.

Early demonstrations with an ensemble of QDs [292] were soon followed by cavities hosting a single emitter [51, 293], and the strong-coupling regime was approached shortly after [294]. Compared to other designs, the mode volume is relatively large $V \approx 5 - 10\,\lambda^3$, but quality factors can exceed 10^4 and $F_P \approx 10$ is routinely achieved. Although high Q-factor cavities provide in-built spectral filtering that suppresses wide-band background, such as wetting layer or phonon-sideband emission, they also require precise control over the QD–cavity detuning (using electric or magnetic fields or piezo-controlled strain) to ensure maximal enhancement. Furthermore, they cannot simultaneously support the transition lines of the XX–X cascade, which inhibits applications utilizing (entangled) two-photon states. Finally, depending on the excitation scheme, the laser may be far detuned from the cavity resonance, requiring high excitation powers and, accordingly, causing high scattered laser intensities. While most micropillar designs aim for a circular cross-section, there were early observations of polarization mode splitting in elliptical pillars [309, 310]. Later, those inspired a resonant π-pulse excitation scheme that exploits the induced polarization mode splitting to generate highly polarized single-photon emission, thus minimizing the loss penalty of cross-polarization laser suppression [86, 255].

Photonic crystal cavity

Photonic crystal cavities are formed by a defect in a photonic crystal, which can be a single missing nanohole [298], a group of three missing holes with slightly off-placed holes at either end (so-called $L3$ cavity) [305], or a whole missing row that creates a waveguide (referred to as $W1$ type) [259, 299]. Due to the limited outcoupling efficiencies of high-Q designs such as $L3$, the waveguide-based approach has been established as the most promising solution, offering broadband operation, albeit with a lower Purcell enhancement. Here, the guided mode is highly dispersive and can be tailored to approach group velocities v_g close to 0, which is why they are also called *slow-light* waveguides. The associated group index $n_g(\omega) = c/v_g(\omega)$

increases sharply which also causes the Purcell factor to increase[27]. Furthermore, the typical mode volumes of $V \approx 0.2 - 1\,\lambda^3$ are significantly smaller than those of micropillars. Although theoretical investigations suggest that $F_p > 50$ is possible [258], fabrication imperfections significantly limit performance due to the small hole diameters required, and only $F_p \approx 3 - 5$ is routinely achieved in $W1$-type cavities [60, 302] with remarkable exceptions such as [296].

Circular Bragg grating

Another in-plane cavity design for QDs known as CBG employs similar principles and creates second-order Bragg-type interference between alternating rings of low- and high-index materials (alternating air-trenches and semiconductor ridges) [295]. The circular arrangement provides lateral mode confinement, and the distinct shape inspired its nickname as Bull's eye resonator. Early implementations realized the vertical confinement with a suspended membrane [295], but recent designs use a flip-chip process [311] to transfer the membrane on a metallic bottom mirror with a spacer layer [80, 173, 257, 312], thus avoiding fragile free-standing membranes.

The directionality of the emission is significantly enhanced, and the far-field mode consists, to very good approximation, of a single Gaussian [295, 313]. With a central disk diameter of $\approx \lambda$ and thickness of $\approx \lambda/4$, the mode volume ranges between $0.3\,\lambda^3$ and $1.5\,\lambda^3$. Hence, a Purcell enhancement of $F_P > 25$ could be demonstrated [257]. At the same time, the relatively low Q-factors of $\approx 100 - 400$ enable broadband operation.

On the other hand, the design is very sensitive to spatial alignment and thus requires precise ($\sim$ 30 nm) emitter localization and fabrication techniques. If the emitter is offset by 50 nm from the center of the CBG, F_p can already drop to 1/3 of the maximum value [257]. Surprisingly, for an even larger offset, the QD–cavity coupling increases again significantly if the emitter position coincides with the first side maximum of the cavity field distribution [313]. Importantly, such an off-centered emitter interacts almost exclusively with one of the cavity's polarization modes, and thus produces highly polarized single-photons [80, 313]. However, the Purcell factor remains lower than for a centered QD, and a more efficient approach of harnessing selective mode coupling to achieve highly polarized emission introduces some ellipticity to break the polarization mode degeneracy (as described above for micropillar designs) [86, 255].

The very broadband geometric enhancement combined with modest Q-factors makes this design ideally suited to support both photons of the XX–X cascade equally, which may be used to generate polarization-entangled pairs [173] if the pho-

[27] For PhC WGs the calculation of the Purcell factor must be adapted to account for the lower dimensionality. See section VI A. in Ref. [258] and in particular Eq. 49

$$F_P = \frac{3}{4\pi n} \frac{\lambda^2/n^2}{V/a}\, n_g(\omega)\,,$$

where a is the optical lattice constant.

tons are collected for emission angles $\lesssim 15°$ [314]. Finally, the in-plane architecture can be co-integrated with piezo-strain tuning [312], adding an additional degree of freedom to tune the emission properties [177]. In this way, a recent experiment demonstrated for the first time the teleportation of a quantum state using photon pairs generated from remote QDs [176].

Open cavity

All of the above structures are fully integrated within the semiconductor platform. While this provides a mechanically stable cavity, the spatial alignment of QD and cavity cannot be adjusted after fabrication, and the spectral alignment is limited to the Stark tuning range. An open cavity addresses these shortfalls [315, 316]. This design is of Fabry-Perot type with two highly reflective mirrors in hemispherical configuration, yielding a Gaussian output mode. The bottom mirror is planar and sits below the QD in the semiconductor sample, and the miniature top mirror is curved, made of a specially carved fiber tip—as also known from fiber-based cavities [315, 317]—and mounted externally. As a result, the top and bottom mirrors can be freely tuned relative to each other. This flexibility forms the basis for record-high system efficiencies [55, 56] by enabling spatial matching of QD and cavity through in-plane positioning, and spectral adjustment of the cavity mode using the out-of-plane axis.

The design features reasonably low mode volumes[28] $V \approx 1\,\lambda^3$ and simultaneously high Q-factors matching those of micropillars and exceeding 10^4. However, the high Q-factors make the cavity extremely sensitive to vibrations, and a misalignment of only 1 nm can shift the cavity mode by 50 GHz, dramatically reducing Purcell enhancement [55]. Although the cavity is geometrically symmetric, a residual (but often wanted) birefringence in the host matrix induces a polarization mode splitting, which can be exploited to generate highly polarized emission [56, 255].

2.2.5.3 Fiber-pigtailed QD sources

Lastly, efforts to fully integrate the optical alignment of the excitation and collection path, and to provide a "plug & play" QD SPS are briefly mentioned [318]. These techniques aim to directly couple single photons to an optical fiber, and most recent experiments also aim to permanently fix the fiber to the semiconductor chip to ensure highly efficient, long-lasting outcoupling.

There are three approaches, which can be divided into near-field [319–321] or far-field vertical coupling [322, 323], or horizontal evanescent coupling using appropriate tapers [324–327]. The most recent and promising pigtailing demonstrations combined near-field coupling with an on-chip cavity structure, where significant advances have been shown for both CBGs [270, 328] and micropillars [329, 330].

[28] The mode is partially in free space and partially within the semiconductor, which complicates the estimation of the mode volume V. For exemplary calculations, see [316].

For an in-depth review on fiber-pigtailing of quantum emitters, see [318] and for a recent overview of the best performing QD demonstrations, see Tab. 1 in [330].

3 Experimental methods

3.1 Quantum dot excitation and collection

All experimental results presented in the following chapters prepare the quantum light from a QD using several components that are shared between all the experiments. Before these key elements of the optical setup are described (Sec. 3.1.2), the QD samples themselves are briefly introduced (Sec. 3.1.1).

3.1.1 Quantum dot samples

The QDs were fabricated by our collaborators from the group of Prof. Peter Michler at the University of Stuttgart. The epitaxially grown QDs consist of In(Ga)As sitting on a GaAs host matrix, where a MMB was used to realize the strain relief required to allow single-photon emission in the telecom C-band[168]. These sources are therefore ideally suited for applications in quantum communication leveraging existing low-loss infrastructure.

The sample employed for the study detailed Chapter 4 features a bottom distributed Bragg reflector where the distance between the 23 pairs of AlAs/GaAs constituting the reflector and the semiconductor/vacuum interface corresponds to a nominal, weak λ cavity, and the QD layer is situated in its anti-node [168]. The attribution of the QD transition to a positive trion is based on power- and polarization-resolved μ-PL measurements, as well as previous experimental and theoretical investigations [168, 331–333] on similar samples.

The results shown in Chapter 5 and Chapter 6 were obtained using a modified sample design, coupling the QD to a circular Bragg grating cavity with a FWHM of $\approx$1250 GHz (see Fig. 3.1a). Starting from a GaAs substrate, a non-linear metamorphic buffer layer [168] is deposited. Reducing the lattice mismatch to the subsequently grown InAs QD layer enables emission in the telecom C-band. In the next step, circular Bragg grating cavities are etched into the sample in a non-deterministic process. More details on the fabrication and the exact geometry of the final structure can be found in [78] and the corresponding supplementary material, following the same design principles.

The studied QD transition is attributed to a charged exciton that features highly polarized emission ($> 95\%$) [80] at $\lambda = 1552.25(3)$ nm. Under optimized

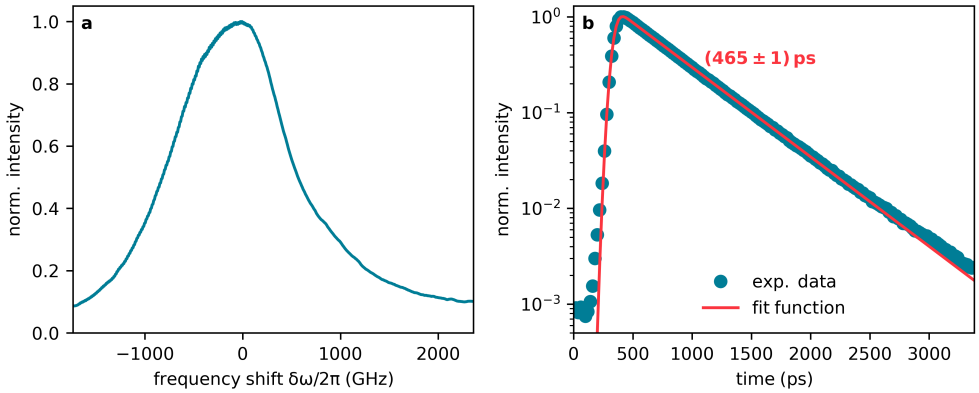


Figure 3.1: **Characteristics of the CBG QD sample.** (**a**) The cavity mode is measured using a strong above-band laser and collecting the emission of the sample on a spectrometer. The frequency is given as detuning from the spontaneous emission line, that is $\delta\omega = \omega - \omega_0$. The spectrometer resolution is 10 GHz per pixel. (**b**) The measurements are performed with pulsed LA excitation at 75.95 MHz, a blue-shifted detuning of 175 GHz, a laser spectral bandwidth of 26(1) GHz and excitation power of 250(15) nW. The mono-exponential fit function (solid black line) yields a decay time of 465(1) ps. The figure is reprinted from the Supplementary Material of [2] under a CC-BY-4.0 license.

LA phonon excitation conditions, a polarized and fiber-coupled single-photon generation efficiency of $\eta_{\mathrm{sys,pol}} = 5.7(2)\,\%$ at $g^{(2)}(0) = 1.3(1)\,\%$ is achieved. The spontaneous emission line is subject to spectral diffusion and not FT limited. The inhomogeneous broadening results in a predominantly Gaussian line shape.

3.1.2 Optical setup

The core optical setup, as shown in Fig. 3.2 may be divided into three sections: laser pulse preparation, laser suppression & collection of the QD emission, and signal processing & detection. The first two are modified only marginally between the experiments and are explained here, whereas the last section will be introduced in the respective chapters.

All experiments use an Er-doped fiber mode-locked laser (*PriTel UOC*) that creates pulses at a repetition rate of $\nu_{\mathrm{rep}} = 75.95\,\mathrm{MHz}$. A filter with 1 nm bandwidth placed within the laser cavity stretches the pulses to a temporal pulse width of 9(3) ps (spectral width FWHM ≈450 pm, the pulses are not fully Fourier-limited). The laser provides a wavelength tunability between 1530 nm and 1555 nm with an average output power of 200 mW. All remaining broadband amplifier noise is removed by a free-space 4-f pulse shaper based on a reflective grating (*Spectrogon*, 1200 lines/mm, blaze at 1550 nm) with an efficiency of ≈ 90 %, C-coated cylindrical lens with a focal length of 500 mm and a variable filtering slit. Depending on the desired excitation parameters, the pulses are further stretched in this shaper setup.

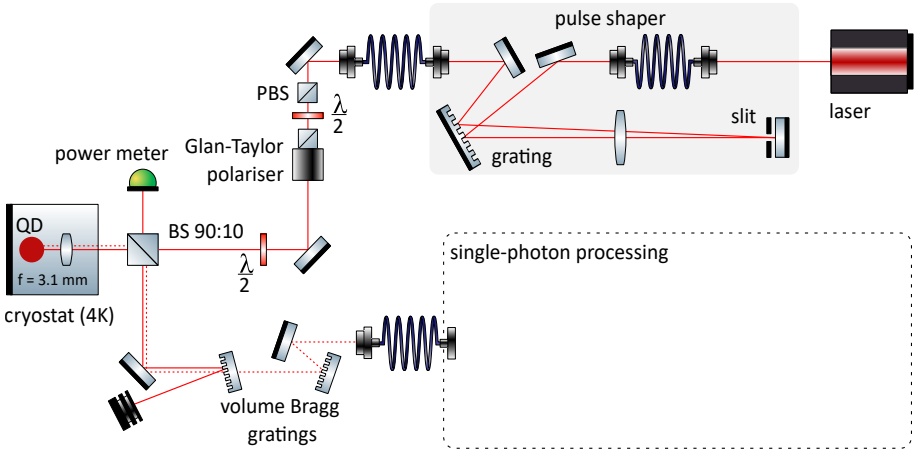


Figure 3.2: **Optical setup for QD excitation and collection.** A schematic representation of the main building blocks of the experimental setup, such as the pulsed excitation laser, pulse shaper, optional tuneable bandwidth filter, excitation of the QD sample, filtering of the QD transition line and the analysis stage. Solid (dashed) red lines indicate the path of the laser beam (QD emission). PBS: polarizing beam splitter, BS: beam splitter, SNSPD: superconducting nanowire single-photon detector.

The excitation beam is collimated into free space using an 8 mm lens collimator (*Schäfter+Kirchhoff 60FC-SF-0-M08-08*). A Glan-Taylor polarizer (*Leysop GTB10-M*, polarization extinction ratio $> 60\,\mathrm{dB}$) sets the excitation beam polarization to linear; the polarization axis can be adjusted by a subsequent half-wave plate (*Bernard Halle*). Just before the cryostat chamber, a 90:10 BS cube (*Thorlabs BS078*) is placed to separate the incoming excitation beam from the single photons emitted by the QD. Approximately 90 % of the excitation beam (depending on its polarization state) is reflected by the BS to a power meter (used to control the QD excitation power) and only $\approx 10\,\%$ of the laser power is guided to the cryostat chamber (*attocube attoDry800*) in which the QD sample is placed.

The QD emission and the reflected laser light are collected by a lens ($f = 3.1\,\mathrm{mm}$) with an NA of 0.68. Leaving through the reflection port of the 90:10 BS, QD emission and laser light are separated using up to two volume Bragg grating (VBG) notch filters (*OptiGrate*) that block the laser wavelength with a spectral bandwidth (FWHM) of 120(1) GHz and an individual suppression of OD6. Finally, the QD emission is filtered from the remaining spectrum by reflecting off another VBG (various options) and coupled by a $f = 8\,\mathrm{mm}$ lens collimator into a SMF. The laser-free, background-reduced emission is then further processed according to the specific research interests of each experiment.

3.2 Analyzing temporal information in quantum optics

The arrival time of a single photon already carries valuable information, but even more can be learned when detecting multiple photons. Consequently, analyzing the temporal information of detection events is essential in quantum optics and has become even more powerful with the advent of fast, highly efficient single-photon detectors and modern time taggers.

For example, for an exciton transition, the radiative decay rate is inferred from the distribution of the photon emission times. In combination with additional measurements, this rate then provides information about the emitter's coupling to a radiation field (Purcell effect, see. Sec. 2.2.5.2), the presence of decoherence mechanisms (*lifetime-limited* emission, see Sec. 2.2.2.3), or cooperative phenomena such as super- and subradiance [334]. When processes involve two or more photons, events can be sorted by their detection patterns, enabling selective discarding of certain events or the application of different analysis methods. This may be as simple as announcing the presence of a single photon in the target mode through the detection of an ancilla photon (heralded single-photon sources), or more complex when postselecting the successful generation of a cluster state using linear optics [58], or boosting the success probability of a Bell state measurement using many ancillas [335].

Although these multiphoton examples rely on coincidence clicks, the coincidence window is large–typically on the order of the experiment's clock cycle–and when exactly any of those ancillas arrives might not be crucial to the experiment. However, there are also applications where a more precise time analysis is required; for instance, to resolve the time ordering of different frequency components [336, 337], to confirm stimulation emission at a few-photon level [338], or to analyze the temporal correlations found in a single-emitter response in detail [63, 96, 220].

In modern setups, the optical signal is converted to digital timing information by two main components: the single-photon detector (often a superconducting nanowire single-photon detector (SNSPD) for highest performance) that transforms the optical pulse into an electrical analog pulse, and a time tagger that detects (typically rising) edges in the input signal and compares them against a high-precision local oscillator to assign an arrival time to the event. The digitized data can be further processed on dedicated field-programmable gate array (FPGA) hardware, or transferred to a computer for live feedback or later in-depth analysis. Both SNSPD and time tagger contribute to the final timing resolution with which an optical signal is recorded, and high-precision detection systems achieve temporal jitter with $\mathrm{FWHM} < 25\,\mathrm{ps}$ per channel. The timing accuracy of a given configuration of the measurement apparatus can be derived from the system response function (SRF). For example, sending a strongly attenuated laser pulse, which is significantly shorter than the expected timing resolution, to an SNSPD and correlating its output with a reference seeded by the laser produces a temporal distribution whose width reflects the overall timing precision. In other cases, one is interested in the SRF

for correlating two SNSPD channels (see Chapter 5).

3.2.1 Basic measurements

Histogram

While a time tagger can be used as a simple event counter, its primary function is to temporally correlate different input signals. At a fundamental level, most measurements[1] simply determine the time differences between two or more input signals and count how often each time difference occurred[2]. For example, a two-channel histogram is a multi-start, multi-stop measurement that records all time differences $\tau_1 = t_1 - t_0$, where t_0 denotes all detection times of the start channel and t_1 those of the stop channel. Due to the construction of the measurement, $\tau_1 > 0$. For convenience, the offset is adjusted such that $t_0 = 0$ and therefore $\tau_1 = t_1$. The time differences are discretized using a bin width (bw) to form the histogram showing how many events were counted for each time difference. Often, an electrical reference signal is used as the start channel–in pulsed experiments, it is derived from the laser clock–and the optical signal(s) serve as stop channels. This type of simple two-channel histogram (see Fig. 3.3a) is used to infer the radiative lifetime of an emission process and is referred to as TCPC or sometimes as first-order correlation $G^{(1)}(t_1)$.

Multi-channel histogram

The measurement can be easily extended to higher dimensions by adding more stop channels. Then, each event represents a multi-channel coincidence $G^{(n)}(\tau_1, ..., \tau_n)$ with τ_i delays, where $i \in [1, ..., k]$ with k being the number of stop channels. Here, the arrival time differences of all input axes are measured analogously as $\tau_i = t_i - t_0$, which is again simplified to t_i instead of τ_i by defining $t_0 = 0$. While the measurement style is often changed to single-start, single-stop to reduce computational complexity, a multi-start, multi-stop measurement can be used when uncorrelated peaks are required for normalization or other purposes. The complete graphical representation of such a data set becomes quickly impractical, and often only certain cross-sections are used. However, for our purposes, we are only interested in at most two stop channels (stemming from two optical input signals), and a two-dimensional (2d) false color map, as shown in Fig. 3.3b, can fully display the 2d histogram of the coincidence data $G^{(2)}(t_1, t_2)$. For all data shown in this chapter, the optical input of both stop channels is identical, and therefore, the same bin width is used along both axes. The significantly lower coincidence counts along the diagonal of Fig. 3.3b where $t_1 = t_2$ indicate that the

[1] Modern time tagger allow to record and save the raw stream of time-tagged events. Hence, in this context, a *measurement* refers rather to how the data is analyzed than how it is acquired. Naturally, to apply a different measurement to the same raw data (i.e., reprocess the raw data in a different way), all required input signals must have been recorded during the acquisition.

[2] An extensive overview of measurement options, including visual representations, can be found, for example, in the documentation of Swabian Instruments.

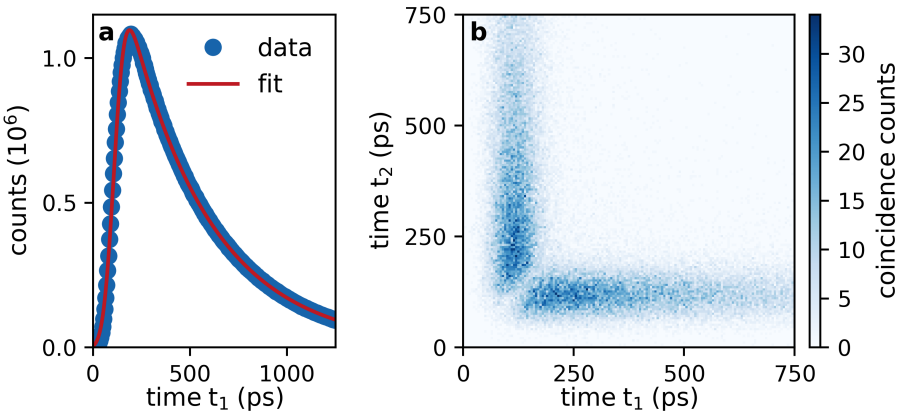


Figure 3.3: **Examples of a one- and two-dimensional histogram.** Exemplary TCPC data obtained from the telecom QDused in Chapter 5 and Chapter 6 under pulsed LA excitation with a repetition rate $\nu_{\mathrm{rep}} = 75.95\,\mathrm{MHz}$ and a pulse width $\Delta t_L \approx 85\,\mathrm{ps}$. The single photons were detected using SNSPDs from Single Quantum featuring a FWHM timing jitter of $\approx 25\,\mathrm{ps}$ and correlated using the TimeTaggerX from Swabian Instruments (FWHM timing jitter of $\approx 4\,\mathrm{ps}$). The electronic reference signal was provided directly by the pulsed laser, and the temporal offsets were adjusted such that $t_0 = 0$ and the peak of the TCPC of each stop channel sits at $t_i = 200\,\mathrm{ps}$. All bins are 5 ps wide and the data acquisition time is 6 min. The optical signal is split with a balanced fiber BS (HBT configuration), and both plots are generated from the same time tags. (**a**) The one-channel TCPC is generally used to infer the radiative lifetime of a selected transition. The rising edge of the signal is affected by the laser pulse width, possibly a relaxation time to the exciton state, and the timing jitter of the detection system. Here, the fit yields $\tau_{\mathrm{rad}} = 426(2)\,\mathrm{ps}$ and half-width at full maximum for the rise time of 56(1) ps. (**b**) Each event on the color map represents a two-photon coincidence between both optical channels at the respective delays t_1 and t_2.

QD rarely emits two photons at the "same" time (more accurately, within the resolution of the detection system).

Correlation

The temporal structure of these multiphoton events[3] can also be analyzed independently of a reference clock signal by directly correlating the two optical channels in

[3]As explained in Sec. 2.2.2.2, we refer to a multiphoton event whenever two photons are emitted in response to one laser pulse.

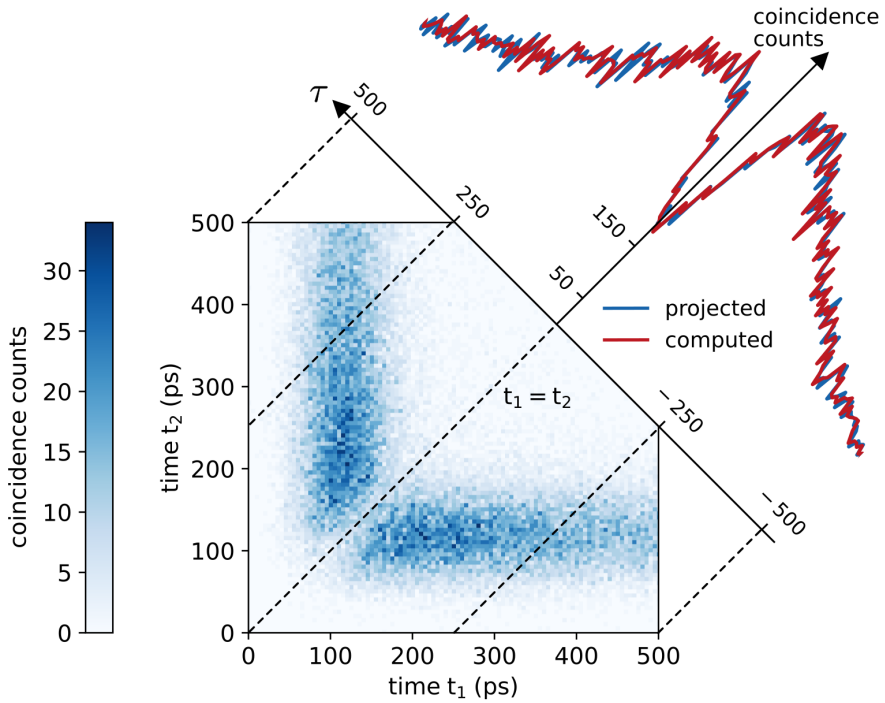


Figure 3.4: **Reduction from $\mathbf{G^{(2)}(t_1, t_2)}$ to $\mathbf{G^{(2)}(\tau)}$.** The line plots are obtained either by directly correlating the optical channel (red) or reconstructed from the $G^{(2)}(t_1, t_2)$ as illustrated. If the "absolute" arrival time is irrelevant to the object of interest, the integration over all bins where the relative delay between the optical channels is the same provides significantly better statistics. The same experimental data as shown in Fig. 3.3 is used.

multi-start, multi-stop fashion[4]. For a symmetrically split input signal, this yields the autocorrelation $G^{(2)}(\tau)$ as defined in Appendix A, where we use τ without any index to denote the time difference between the two optical channels. Since $G^{(2)}(\tau)$ is independent of the arrival times with respect to the clock, it can be reconstructed from $G^{(2)}(t_1, t_2)$ by summing over all events that occur the same delay $\tau = t_2 - t_1$ irrespective of their "absolute" detection times, as illustrated in Fig. 3.4 (effectively, integrating the 2d data along the bisectrix $t = t_1 = t_2$). Naturally, the reverse operation is not possible since the time reference to the clock is missing. While both line plots in Fig. 3.4 quite similar, they show deviations from bin to bin, which

[4]Before modern time taggers enabled multi-start, multi-stop measurements, the combination of a time to amplitude converter and multi-channel analyzer could only implement single-start, single-stop measurements. When the start signal is determistic (i.e. reference clock in a histogram), the results are identical but for a correlation with two optical signals, such an arrangement can cause distortions if the signal rates approach the repetition rate of the experiment. Effectively, the probability of a coincidences drops exponentially to 0 for large τ and the decay rate is related to the click probability. See [339] for more details.

is caused by the discretization and subsequent integration along the diagonal.

3.2.2 Time-gated measurements

In pulsed experiments, the target signal typically arrives in a time window much smaller than the repetition cycle. To improve the SNR, one sometimes chooses to analyze a smaller time window, $T_{\mathrm{int}} < T_{\mathrm{rep}}$. For which cases such a postselection provides valid results and for which it can significantly misrepresent the quantity of interest will not be discussed here[5]. Rather, various options for implementing time-gating are introduced.

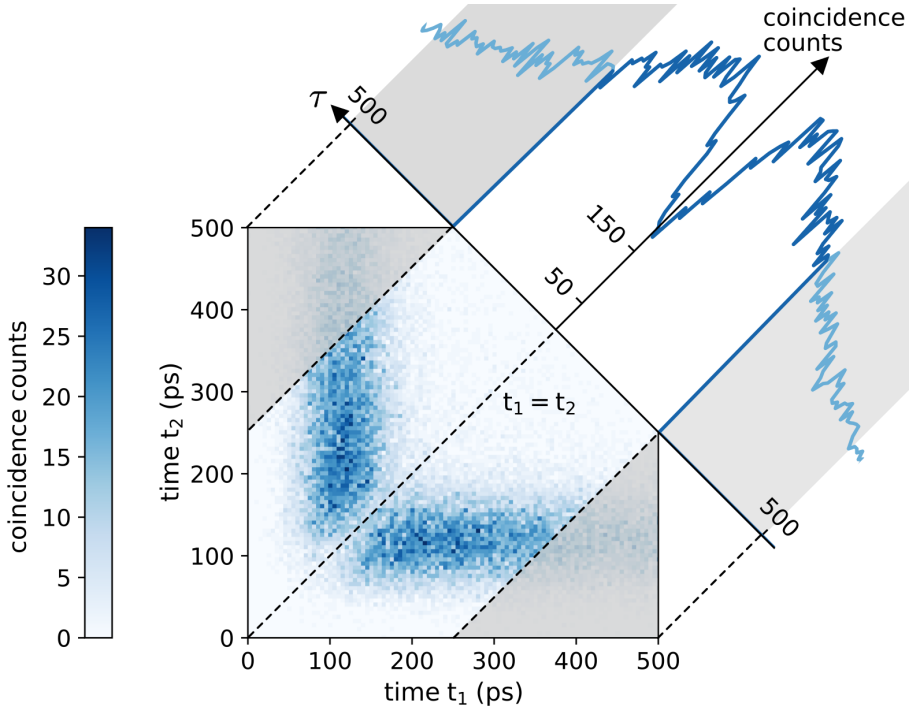


Figure 3.5: **Illustration of a reduced integration window for $\mathbf{G^{(2)}}(\tau)$.** Data as in Fig. 3.3. The shaded areas in the 2d histogram and the line plot indicate which coincidence events are neglected for $T_{\mathrm{int}} = 500\,\mathrm{ps}$. Only the line plot for the direction correlation is shown here for clarity.

Time-gated correlation

A common approach simply uses $T_{\text{int}} < T_{\text{rep}}$ when evaluating the single-photon purity via $g^{(2)}(0)$ or when estimating the TPI visibility using an HOM setup. The reasoning is typically that the "true" emitter properties are degraded by external noise introduced by the preparation and detection setup, and that postselection reduces the impact of those background contributions. Such an evaluation is presented in Sec. 4.3 as a function of T_{int}. The effect is visualized in Fig. 3.5 on both a 2d histogram and a direct correlation. Here, $T_{\text{int}} = 500\,\text{ps}$ is unrealistically narrow and purely chosen for illustration purposes. More appropriate windows respect $T_{\text{int}}/2 > 5\tau_{\text{rad}}$ to cover at least 99 % of the spontaneously emitted photons. Obviously, the same integration window T_{int} must be applied to the reference signal (i.e., the uncorrelated side peaks) to maintain proper normalization.

Time-gated input channels

Alternatively, one can apply the gating window to one (or multiple) channel as an initial step in the postprocessing workflow. This approach mimics a physically gated optical input channel while achieving higher temporal resolution, and crucially, the gating can be adjusted and optimized after the raw data has been recorded. The difference between time-gating before or after correlating two optical channels can be drastic, as displayed in Fig. 3.6. Two gating windows with the same width $T_{\text{gate}} = 500\,\text{ps}$ but differing starting points $t_{GS,1} = 50\,\text{ps}$ or $t_{GS,2} = 200\,\text{ps}$ are illustrated on the TCPC signal in Fig. 3.6a. Applying those windows symmetrically on both channels, effectively ignores all coincidence events outside the boxes indicated on the 2d histogram in Fig. 3.6b, and the $G^{(2)}(\tau)$ is clearly expected to differ significantly depending on the start gate time. The effect that the gating has on the uncorrelated peak is shown in Fig. 3.6c, where the blue line represents the ungated case. Since T_{gate} is so narrow that a significant part of the spontaneous emission is disregarded, the number of uncorrelated coincidences is reduced. Notably, the temporal shape changes depending on whether the rising edge is included (green and red curves).

Finally, Fig. 3.6d highlights the stark difference observed in the correlated peak for the same three gating configurations. While the $G^{(2)}(\tau)$ is almost unchanged when the gating window includes the rising edge, it is drastically reduced and no longer features the pronounced dip at τ if the gate start time coincidences with the TCPC peak; a clear sign that at least one of the photons must be emitted early on. The underlying re-excitation mechanism was introduced in Sec.2.2.4.3 and will be discussed in detail for LA phonon-assisted driving in Chapter 5. Here, we simply

[5]For example, for the honest parameter estimation required in most cryptographic protocols, such an analysis is only permissible if the stream of encoded single photons is also physically time-gated and the postselection solely suppresses background noise from detector dark counts. On the other hand, if a fundamental property of a SPS is studied for low signal counts or in the presence of excessive external noise that can be identified but not isolated (e.g., backreflections or detector afterpulsing), time-gating can be a valuable tool.

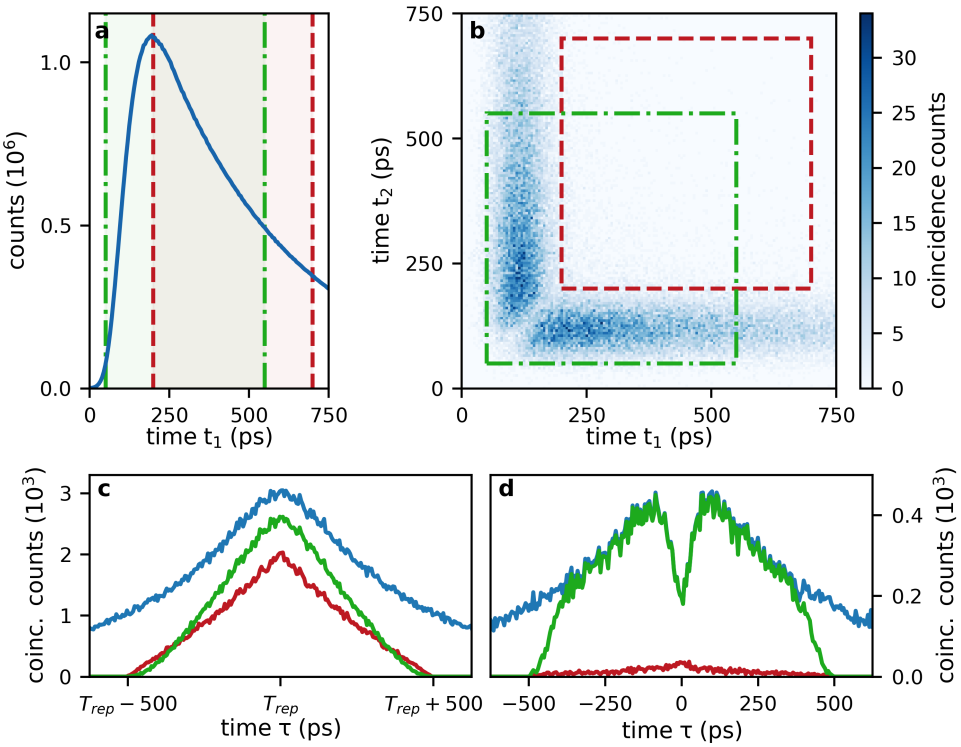


Figure 3.6: **Comparison of gating strategies.** For the same data set as used in Fig. 3.3, four different evaluations of the multiphoton component are compared: i) raw correlations, ii) raw correlations but $T_{\mathrm{int}} = 1000\,\mathrm{ps}$, time-gating the input channels with $T_{\mathrm{gate}} = 500\,\mathrm{ps}$ starting at $t_{GS,1} = 50\,\mathrm{ps}$ (iii, green, dash-dotted) or $t_{GS,2} = 200\,\mathrm{ps}$ (iv, red, dashed). The gating windows are indicated on the TCPC histogram (**a**) and on the 2d histogram (**b**). The resulting $G^{(2)}(\tau)$ is shown for the uncorrelated peak at $\tau = T_{\mathrm{rep}}$ (**c**) and the correlated peak at $\tau = 0$ (**d**). The $g^{(2)}(0)$ values are calculated from the ratio of the areas, yielding $\{0.166, 0.153\}$ for {ungated (i), correlation-gated (ii)} and $\{0.228, 0.014\}$ using the input channel gating with {early window (iii), late window (iv)}.

register that appropriate time-gating of the input channel can change the measured $g^{(2)}(0)$ considerably, whereas reducing the integration window T_{int} mostly reduces the impact of background noise.

This type of analysis provides explicit experimental clues that for high-quality SPSs, the early portion of the signal is essential for the multiphoton contribution. Motivated by these findings, we slightly adapted the gating method in Refs. [338] and [220]. Instead of using a fixed window size, the gate finish time was maintained at $t_{GF} = 2000\,\mathrm{ps}$ while varying the start gate time t_{GS}. The time-gated multiphoton

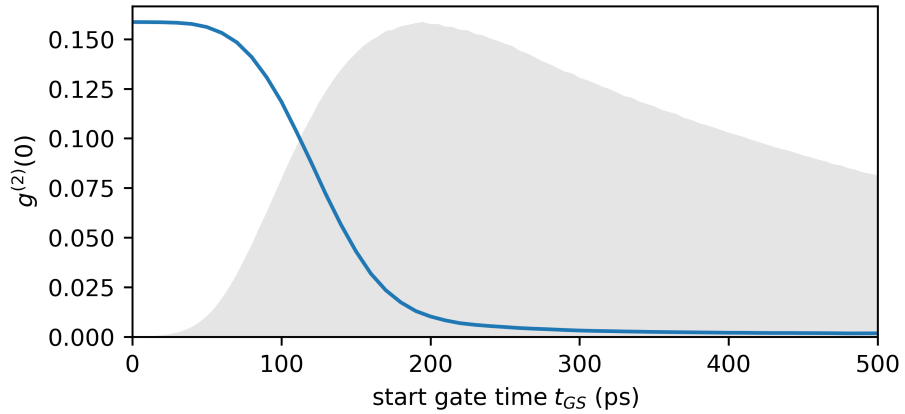


Figure 3.7: **Temporal distribution of the multiphoton component.** The optical input channels are gated such that only events with $t_{GS} < t < t_{GF}$ are accepted for the correlation (see Fig. 3.6). For a fixed $t_{GF} = 2000\,\text{ps}$, the start gate time is varied (in steps of 20 ps), the gated channels are correlated, and the multiphoton contribution is estimated using Eq. 3.2-1. The ungated TCPC signal is shaded in the background as a temporal reference.

contribution then reads as

$$g^{(2)}(t_{GS}) = \frac{\int_{t_{GS}}^{t_{GF}} \int_{t_{GS}}^{t_{GF}} G^{(2)}(t_1, t_2)\, dt_1\, dt_2}{\int_{t_{GS}+T_{\text{rep}}}^{t_{GF}+T_{\text{rep}}} \int_{t_{GS}+T_{\text{rep}}}^{t_{GF}+T_{\text{rep}}} G^{(2)}(t_1, t_2)\, dt_1\, dt_2} \tag{3.2-1}$$

and is shown in Fig. 3.7 for the same data set as before. The $g^{(2)}(0)$ drops smoothly from 0.159 when effectively no gating is applied to $g^{(2)}(0) = 0.010$ when the gating window starts at the peak of the TCPC signal ($t_{GS} = 200\,\text{ps}$) and even further to $g^{(2)}(0) = 0.003$ for $t_{GS} = 300\,\text{ps}$.

Additionally conditioned measurements

Although these methods provide a strong indication that, in a multiphoton emission process, one photon must be emitted early, they do not directly capture its temporal profile. Isolating the individual shapes of each photon can provide valuable insights into identifying the cause of multiphoton emissions.

For realistic QD SPSs, there are several mechanisms that can lead to the emission of two or more photons (see Sec. 2.2.2.2). Some are expected to exhibit a specific temporal structure (including re-filling, re-excitation, cascaded emission) while others should be mostly uncorrelated (for example, emission from neighboring QDs or the wetting layer). Analyzing the time tags in a way that assumes a temporal structure allows us to distinguish between them.

Conceptually, such a measurement searches for coincidences between the optical channels that occur within some window $T_{\text{coinc}} < T_{\text{rep}}$ and records the time stamps

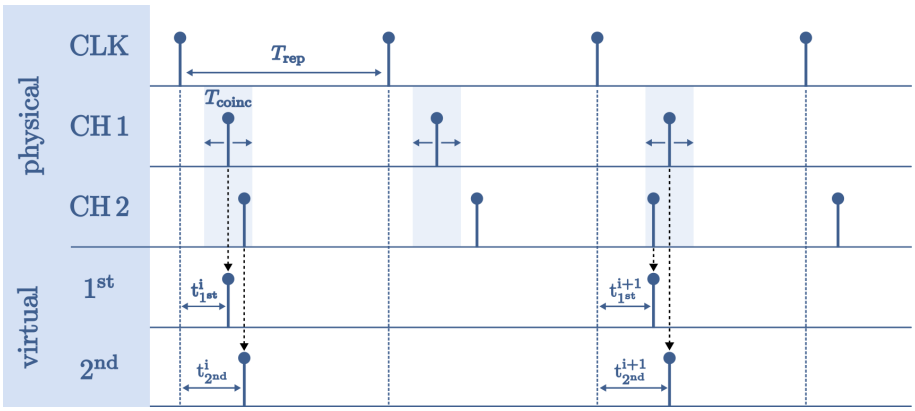


Figure 3.8: **Correlation logic using virtual channels.** The time tagger records three physical channels: the reference clock (CLK) provided by the pulse laser, and the electronic output of two SNSPDs (CH 1 and CH 2. Whenever two events on the optical channels are registered within a coincidence window $\Delta\tau_c = 2\,\mathrm{ns}$, they are accepted as coincidence, and the individual events are "copied" onto the virtual channels where the early click of each pair is assign to one channel (1^{st}) and the late click to the other (2^{nd}). The histograms between the reference clock and the virtual channels then represent the temporal shapes of the individual photons as displayed in Fig. 3.10.

of both events as shown in Fig. 3.8. All of the early events from each coincidence are accumulated in one histogram (against the reference clock), and all late events are assorted in a second histogram. These histograms represent the temporal profiles of the 1^{st} and the 2^{nd} photons. Here, the notation of 1^{st} and 2^{nd} implicitly captures that this method postselects multiphoton events. Due to the timing jitter of the measurement apparatus, some events are falsely assigned and can skew this type of measurement in extreme cases. For more details on the estimation of such an assignment error, see Supplementary Note 4 of our work [2].

Such a measurement can also be visualized using the 2d histograms introduced before. In the area below the bisectrix, only coincidences with $t_1 > t_2$ are considered (remember that t_1 and t_2 are the arrival times at detectors #1 and #2; they do not denote the 1^{st} and 2^{nd} photon). Hence, the 1^{st} photon is always registered on detector #2 and the 2^{nd} photon on detector #1. Consequently, to learn the temporal shape of the 1^{st} or 2^{nd} photon, we simply integrate over all detection times of the other detector, as indicated in Fig. 3.9. For the present data set, there is a distinct difference between the shapes of the early and the late photon, suggesting that the multiphoton emission is indeed temporally structured.

To identify if the unwanted photon stems from re-filling or re-excitation, we can make use of a few additional cues. If caused by re-excitation, we expect that the 1^{st} photon must be emitted while the laser pulse is still present (otherwise the QD could not have been excited a second time), whereas for re-filling the tunneling from the trap states dictates the relevant timescales [190]. If this analysis remains

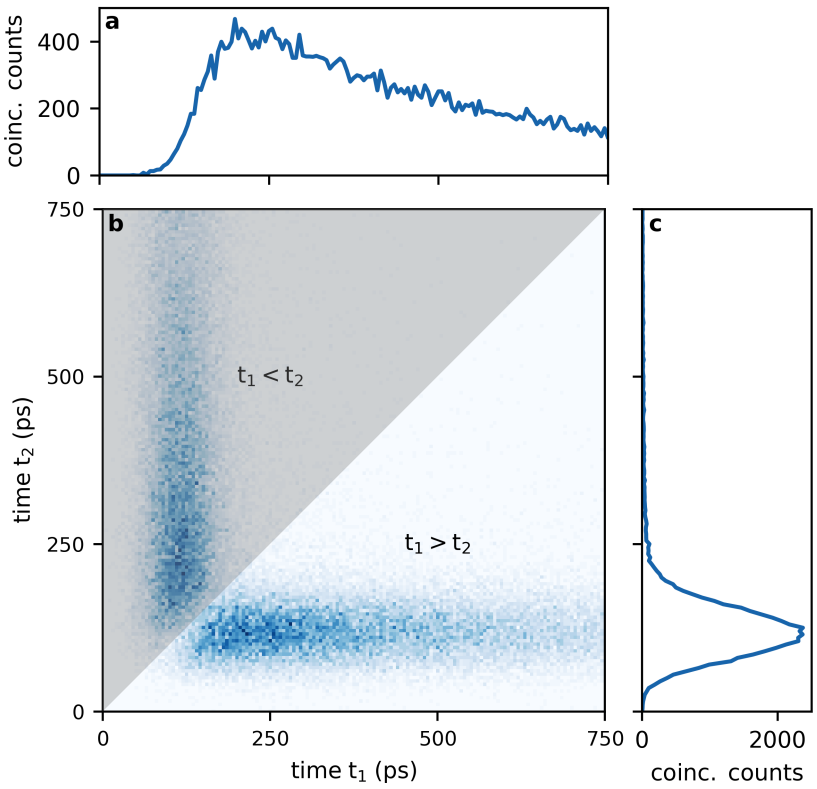


Figure 3.9: **Individual photon time profiles from a 2d histogram.** For the same data set as used in Fig. 3.3, the temporal shape of the 1st and 2nd photon can be differentiated from one another by only considering events satisfying $t_1 > t_2$. Integrating over t_1 yields the distribution of arrival times on detector #2, which can be attributed solely to 1st photons since $t_1 > t_2$. Analog, integration along t_2 provides the temporal shape of the 2nd photon. To improve statistics, all events with $t_1 < t_2$ can also be included (naturally, the channel-to-photon assignment must be reversed).

inconclusive, one can vary the ratio between laser pulse width and radiative decay time, since re-excitation is expected to scale approximately linearly (see Sec. 2.2.4.3) while re-filling rather shows power dependence. In Fig. 3.10, the individual photon shapes are shown next to one another and are overlaid with the laser excitation pulse. The result strongly suggests re-excitation as the cause of the multiphoton emission, but more data and contextualization will be presented in Chapter 5.

Finally, if the multiphoton contribution had been caused instead by emission from a neighboring QD, the analysis would yield identical photon shapes along both time axes (assuming no interaction between the QDs that would induce temporal correlations). In such a case, fitting the temporal shape with a bi-exponential could reveal differences in their individual radiative decay times τ_{rad}.

Yet another analysis strategy resorts to 2d histograms similar to the ones shown

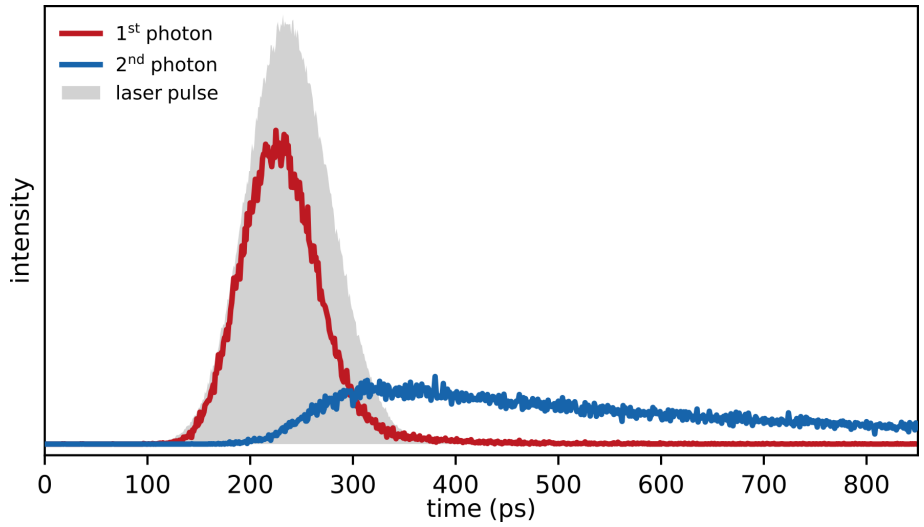


Figure 3.10: **Photon profiles caused by re-excitation of a 2LS.** Displaying the excitation laser pulse alongside the individual photon shapes extracted as in Fig. 3.9 provides insight into the cause of the multiphoton emission. Here, the 1st photon is emitted exclusively during the laser envelope, strongly pointing to re-excitation. To rule out re-filling that might coincidentally act on the same time scale and provide a conclusive result, the measurement could be repeated for various laser pulse widths (as presented in Chapter 2.2.4.3). The figure is reproduced in a modified version from [2] under a CC-BY-4.0 license.

here, but instead of using the "absolute" detection time of the 2nd photon, one can transform to a relative basis using $\tau = t_2 - t_1$ to highlight the delay between the photon emissions [96, 340].

This section presented various methods to evaluate the same set of time tags acquired using a common HBT arrangement (including a reference clock). By modifying the physical setup, additional properties of the quantum light can be analyzed. An example where frequency filtering is applied to only one of the optical channels is presented in Chapter 5.

4

Robust excitation of a C-band quantum dot for quantum cryptography

This chapter presents the publication "Robust excitation of C-band quantum dots for quantum communication" [1] and contextualizes the work in light of this thesis. The project was motivated by recent findings from our research group showing under which circumstances QDs enhance the performance of various cryptographic protocols [45]. Key messages include that there is no single pumping scheme that fits all purposes, and that PNC, shown to exist in resonant schemes [214, 215], poses a significant threat to most protocols (see Sec. 2.1.4.3). For these and other reasons, we were interested in exploring LA phonon-assisted excitation for quantum cryptography using the QD samples from our collaborators in Prof. Peter Michler's group at the University of Stuttgart, which directly emit single photons in the telecom C-band, making them ideally suited for long-distance communication.

In contrast to resonant schemes that induce Rabi rotations using π-pulses, the phonon-mediated driving offers a much larger parameter space that achieves efficient excitation. On the one hand, this makes the scheme inherently more robust [260], but on the other hand, theoretical results suggested that single-photon generation efficiency and multiphoton noise cannot be optimized simultaneously [95]. The goal of this project was then to experimentally determine how to optimize the performance of the QD as a SPS for a given cryptographic protocol.

***My contribution:** With support from Michal Vyvlečka, I built and designed the pulse shaper to consistently tune the excitation wavelength, and the analysis setup. Together with Michal Vyvlečka, we characterized the sample in order to choose a QD for the experiment. The project was developed by Michal Vyvlečka, Mathieu Bozzio, and myself. Together with Michal Vyvlečka, I performed measurements and, with his input, processed and visualized the experimental data. I implemented the theoretical extrapolation of our results for QKD under the guidance of Mathieu Bozzio.*

Large parts of the work presented in this chapter have been published in [1] or its Supplementary Material and are licensed under a Creative Commons Attribution (CC BY 4.0) license.

The paper entitled "Robust excitation of C-band quantum dots for quantum communication" explores how the photon-number probabilities of a QD under LA phonon-assisted driving are affected by the laser power and its detuning from the emitter's wavelength. The scheme is known to provide an exceptionally robust excitation pathway due to the spectral width of the exciton-phonon coupling [95, 260]. However, previous experimental studies have been limited to investigating the impact of parameter variations on the single-photon generation efficiency [235], even though the multiphoton contribution is at least equally important for most applications.

This project implements an extensive parameter scan on a charged exciton emitting at 1549 nm, sweeping the laser power beyond saturation and increasing the laser–QD detuning from 0.7 nm to 3.7 nm. The count rate and $g^{(2)}(0)$ value are recorded for each of the more than 250 parameter pairs, and the respective photon-number probabilities are inferred. Using this experimental data, the joint optimization of laser settings for a specific application is demonstrated, using the secure key rate (SKR) of QKD as a figure of merit. The extrapolation of the SKR is performed for various combinations of external parameters (e.g., error rates, channel losses, dark count rates), and both for standard BB84–QKD and for decoy-state BB84–QKD. The analysis reveals that using a SPS standard BB84–QKD performs better across almost all regimes and that variations of the external parameters call for different ideal laser settings.

4.1 Methods

In addition to the introductions to QKD (Sec. 2.1.4.1), light–matter interaction of QDs (Sec. 2.2.3) and specifically phonon-assisted excitation (Sec. 2.2.4.2), this section briefly explains the methods and experimental tools required to follow the manuscript's reprint in Sec. 4.2 and its accompanying measurements (Sec. 4.3).

Sec. 4.1.1 provides a schematic layout of the setup, while Sec. 4.1.2 presents how the experimental data is processed to extract the performance parameters of the SPS. Sec. 4.1.3 details how these are then used to estimate the photon-number probabilities and extrapolate the SKR for each laser setting. Finally, a simple approximation for the maximally achievable QKD-distance is derived for SPSs (Sec. 4.1.4).

Since this manuscript was published, the use of "brightness" in the context of single-photon sources has been deprecated due to its ambiguous terminology; see [181]. For consistency, the wording was not changed in the reprint. Therefore, the following methods will also use the term brightness to describe the quantity that is physically equivalent to the system efficiency η_{sys} introduced in Sec. 2.2.2.1.

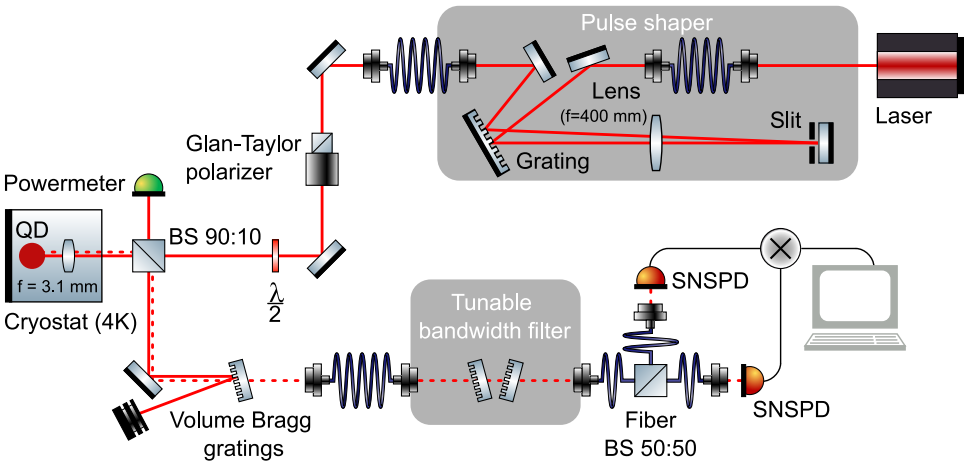


Figure 4.1: **Scheme of the experimental setup.** A schematic representation of the main building blocks of the experimental setup, such as excitation pulse laser, pulse shaper, excitation of the QD sample, filtering of the QD transition spectral line, and detection. The figure is reprinted from the Supplementary Material of [1] under a CC-BY-4.0 license.

4.1.1 Experimental setup

The experimental setup displayed in Fig. 4.1 adopts the structure and instruments explained in Sec. 3.1.2 to a large extent. The full excitation and collection setup exhibits an efficiency of 0.13. The analysis setup consists of a bandpass filter (FWHM= 150 pm) to suppress any spectral background contributions and a fiber-based HBT configuration to measure the photon correlations. The SNSPD channels have an efficiency of 0.86 each, and show dark count rates of 130 Hz and 180 Hz. The temporal resolution of the detection apparatus is 34 ps (FWHM of the SRF).

The QD sample features a bottom distributed Bragg reflector where the distance between the 23 pairs of AlAs/GaAs constituting the reflector and the semiconductor/vacuum interface corresponds to a nominal, weak λ cavity, and the QD layer is situated in its anti-node [168]. The attribution of the QD transition to a positive trion is based on power- and polarization-resolved μ-PL measurements, as well as previous experimental and theoretical investigations [168, 331–333] on similar samples. The TCPC measurement reveals the decay trace as displayed in Fig. 4.2, and a mono-exponential fit function yields a decay time of 1.07(2) ns.

4.1.2 Data processing

For each laser detuning and pump power, a single measurement run is performed from which both the brightness B and purity P are inferred. To obtain enough statistics for the autocorrelation $g^{(2)}(\tau)$, each measurement acquires data until the

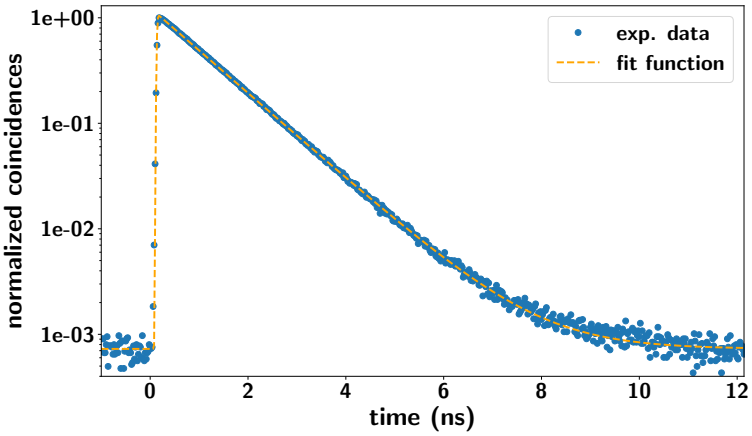


Figure 4.2: **Time-resolved fluorescence measurement.** For this measurement, the excitation laser was spectrally detuned by 1.5 nm and a pump power of 0.89 a.u. was used. The mono-exponential fit function (dashed line) yields a decay time of 1.07(2) ns. The figure is reprinted from the Supplementary Material of [1] under a CC-BY-4.0 license.

coincidence counts of the uncorrelated side peaks exceed a threshold value (here, 700 counts at 100 ps bin width).

Simply summing the count rates of both detection channels results in probabilistic double-counting of multiphoton states

$$\tilde{B} = B + \sum_{i \geqslant 2} p_i \left(1 - \frac{1}{2^{i-1}}\right) = B + \frac{1}{2}p_2 + \frac{3}{4}p_3\,. \tag{4.1-1}$$

Thus, to avoid over-counting, we subtract all coincidence events occurring within one repetition period $T_{\mathrm{rep}} = \frac{1}{\nu_{\mathrm{rep}}}$ such that the measured brightness reads as

$$B = \frac{T_{\mathrm{rep}}}{\eta_{\mathrm{setup}}\eta_{\mathrm{d}}} \times \Big(R_1 + R_2 - CC_{1,2}(t_{\mathrm{coinc}} = T_{\mathrm{rep}})\Big) \tag{4.1-2}$$

where R_i is the raw count rate of the $i-$th detector and CC denotes the coincidence count rate.

Before inferring the $g^{(2)}(0)$ value numerically from the correlation measurement, only the background originating from coincidences arising from a dark count in at least one of the channels is subtracted. It is computed as

$$CC^{\mathrm{d}}_{1,2} = R_1 R^{\mathrm{d}}_2 + R_2 R^{\mathrm{d}}_1 + R^{\mathrm{d}}_1 R^{\mathrm{d}}_2 \tag{4.1-3}$$

where the superscript "d" denotes the dark count rates.

Integrating the area over the central repetition period and dividing by the averaged and blinking-corrected area of the outer peaks then yields the $g^{(2)}(0)$

value. For applications in quantum communication, it is crucial to consider the full repetition period when calculating the source's purity since an adversary has access to all the information leaving the sender's lab.

4.1.3 Extrapolation of the SKR

Since the security analyses of cryptographic protocols are based on photon-number probabilities, these probabilities must be inferred from experimentally accessible data to ensure the protocol's security. Such an estimation with minimal assumptions is derived for SPSs in Sec. 2.2.2.4, requiring only the brightness B and the $g^{(2)}(0)$ value (recall, B is equivalent to η_{sys}).

To extrapolate the performance of the QD for a standard BB84–QKD scenario, the photon-number probabilities are used to derive worst-case bounds by attributing all losses and errors that occur in the untrusted channel to single-photon states, yielding Eqs. 2.1-46 and 2.1-47.

However, these bounds are not tight and degrade the protocol's performance rapidly when a significant multiphoton component is present, as in the case of WCS. As a countermeasure, the sender can randomly modulate the intensities of the states, chosen from the set $\{\rho, \nu_1, \nu_2, ...\}$. In this way, sender and receiver can cooperate to evaluate the total gain and error rate for each state $\{Q_{\mathrm{tot}}^{(i)}, e_{\mathrm{tot}}^{(i)}\}$ and solve the system of equations for $\{Y_i, e_i\}$ (see Eqs. 2.1-41 and 2.1-43) without resorting to bounding. These so-called *decoy* states have been shown to increase the achievable SKR drastically for implementations based on WCS [41]. Even though decoy states decrease the sifting efficiency, as they cannot contribute to the raw key, a tighter estimation of the single-photon contribution is often preferable, especially for large distances. When working with sub-Poissonian sources, the multiphoton component $p_{n\geq 4}$ can be safely neglected, and two decoy states are sufficient to compute the yields and error rates $\{Y_i, e_i\}$ exactly.

In both cases, the achievable SKR is simulated according to Eq. 2.1-48.

4.1.4 Maximum QKD-distance approximation from experimental measures

Estimating the maximum attainable communication distance d_{max} for QKD with a given source is highly relevant for practical implementations. However, due to the complexity of Eq. 2.1-48, one cannot solve it analytically for the channel transmission η_{ch} but has to resort to numerical methods when a tight approximation is required. On the other hand, even the numerical evaluation of the SKR for a given photon-number statistics always involves estimating other protocol parameters such as single-photon detection error, detection efficiency, dark-count probability, and error-correction code inefficiency. Therefore, an analytic approximation for d_{max} can be advantageous—especially if the required quantities are easily accessible.

We start by upper bounding Eq. 2.1-48 by

$$\mathrm{SKR} \leq \frac{1}{2} Q_1 (1 - H_2(e_1)) \tag{4.1-4}$$

as the cost of error correction, $f(e_{\mathrm{tot}}) Q_{\mathrm{tot}} H_2(e_{\mathrm{tot}})$, is strictly positive. The maximum attainable distance $d_{\max}$ is formalized as the minimal channel efficiency $\eta_{\mathrm{ch}}^{\min}$ for which SKR$> \delta$ where a threshold of $\delta = 10^{-8}$ is used here. From Eq. 4.1-4 follows SKR$\to 0$ if $\frac{1}{2} Q_1 (1 - H_2(e_1)) \to 0$, or further simplified $Q_1 \to 0$. Note that $e_1 \to \frac{1}{2}$ results in a vanishing SKR but since $Q_1 \to 0$ also implies $e_1 \to \frac{1}{2}$ (see Eqs. 2.1-44 to 2.1-47) we focus only on $Q_1 \to 0$. Inserting Eq. 2.1-44 into Eq. 2.1 -46 and rearranging yields

$$Q_1 = p_{\mathrm{m}} \left(\frac{p_1}{p_{\mathrm{m}}} Y_1 + Y_2 - 1 \right). \tag{4.1-5}$$

Since $p_{\mathrm{m}} \geqslant 0$ for a realistic source, the expression inside the bracket has to tend to zero to cause $Q_1 \to 0$. Assuming $Y_0 \ll 1$ we can simplify Eq. 2.1-41 to $Y_1 = \eta_{\mathrm{ch}} - Y_0$ and $Y_2 = 2\eta_{\mathrm{ch}} - \eta_{\mathrm{ch}}^2 - Y_0$ such that $Q_1 \geqslant 0$ entails

$$1 \leq \eta_{\mathrm{ch}} \left(\frac{p_1}{p_{\mathrm{m}}} + 2 - \eta_{\mathrm{ch}} \right) - Y_0 \left(\frac{p_1}{p_{\mathrm{m}}} + 1 \right). \tag{4.1-6}$$

As we are looking for long-distance communication, we use $\eta_{\mathrm{ch}} \ll 1$ and rearrange

$$\eta_{\mathrm{ch}} \geqslant \frac{1}{\frac{p_1}{p_{\mathrm{m}}} + 2} + Y_0 \frac{\frac{p_1}{p_{\mathrm{m}}} + 1}{\frac{p_1}{p_{\mathrm{m}}} + 2} \geqslant \frac{1}{\frac{p_1}{p_{\mathrm{m}}} + 2} + Y_0 \,. \tag{4.1-7}$$

The maximum distance at which a secure key can still be generated now corresponds to the minimal channel transmission that satisfies Eq. 4.1-7. Therefore, we write

$$\eta_{\mathrm{ch}}^{\min} \approx \frac{1}{\frac{p_1}{p_{\mathrm{m}}} + 2} + Y_0 \tag{4.1-8}$$

and finally, use Eq. 2.2-20a to re-express the result only in terms of B and $g^{(2)}(0)$ as

$$\eta_{\mathrm{ch}}^{\min} \approx \frac{1}{\frac{B g^{(2)}(0)}{1 - B g^{(2)}(0) - \sqrt{1 - 2B g^{(2)}(0)}} + 1} + Y_0 \,. \tag{4.1-9}$$

Since $B g^{(2)}(0) \ll 1$, we can expand $B g^{(2)}(0)$ in a Taylor series and truncate after the first order such that

$$\eta_{\mathrm{ch}}^{\min} \approx \frac{B g^{(2)}(0)}{2} + Y_0. \tag{4.1-10}$$

For state-of-the-art technology, we have $Y_0 \ll Bg^{(2)}(0)$ such that one can also dismiss Y_0 in the above equation.

Considering the complexity of Eq. 2.1-48, the approximation is strikingly simple and follows directly from the fundamental source parameters. Yet, for a broad parameter range, the results compare well to the numeric solution where Eq. 4.1-10 always overestimates the maximum distance. This systematic error stems from the approximation's derivation as an upper bound and is primarily due to the omission of the error-correction term in Eq. 2.1-48. Typically, the overestimation amounts to $25 - 35$km.

4.2 Manuscript reprint

Building a quantum internet requires efficient and reliable quantum hardware, from photonic sources to quantum repeaters and detectors, ideally operating at telecommunication wavelengths. Thanks to their high brightness and single-photon purity, quantum dot (QD) sources hold the promise to achieve high communication rates for quantum-secured network applications. Furthermore, it was recently shown that excitation schemes such as longitudinal acoustic phonon-assisted (LA) pumping provide security benefits by scrambling the coherence between the emitted photon-number states. In this work, we investigate further advantages of LA-pumped quantum dots with emission in the telecom C-band as a core hardware component of the quantum internet. We experimentally demonstrate how varying the pump power and spectral detuning with respect to the excitonic transition can improve quantum-secured communication rates and provide stable emission statistics regardless of network-environment fluctuations. These findings have significant implications for general implementations of QD single-photon sources in practical quantum communication networks.

4.2.1 Introduction

The emergence of practical quantum technology paves the way to a quantum internet—a network of connected quantum computers capable of reaching computational speed-ups in various tasks such as prime factoring [341], machine learning [342], and the verification of NP-complete problems with limited information [343]. Although such schemes are appealing, most are technologically challenging, while the security advantages provided by quantum cryptography are more tangible [35, 344, 345]. A broad range of quantum-cryptographic primitives including QKD [35, 344, 345], quantum coin flipping [346–348], unforgeable quantum tokens [349–351], and quantum bit commitment [32, 352, 353] have been developed to demonstrate some security advantage over their classical counterparts. The success of a future quantum internet then relies on the development of fundamental quantum hardware (sources, repeaters and detectors) which should adhere to these primitives' security standards, provide high communication rates, and operate reliably in a real-world environment [354].

Non-classical light sources such as spontaneous parametric down-conversion [355, 356], nitrogen-vacancy centers [357] and trapped atoms [358], have been used as hardware for the first quantum networks. In recent years, semiconductor QDs have materialized as highly versatile and quality single-photon sources [56, 86, 242, 359, 360], with outstanding end-to-end efficiencies overcoming 57 % and the potential to reach repetition rates of tens of GHz [56]. Such emission properties of QDs have led to the implementation of complex network building blocks relying on quantum teleportation [361, 362] and quantum entanglement swapping [69, 359, 360]. Regarding the emission wavelength, the spectral regime of the telecom C-band (1530 nm to 1565 nm) is highly appealing, due to its global absorption minimum in standard silica fibers, the possibility to implement daylight satellite communication

[363] and the compatibility with mature silicon photonic platforms [364]. QDs with emission wavelengths in and around the C-band are available on indium phosphide (InP) [365–367] and gallium arsenide (GaAs) material system [78, 168, 368], and circumvent the technical overhead and losses of quantum frequency conversion [369]. Embedded in circular Bragg cavities, QDs based on the well-established GaAs platform have simultaneously demonstrated high brightness and high purity values recently [78].

Previous works have investigated the advantages and drawbacks of various optical pumping schemes (resonant, phonon-assisted and two-photon excitation) in terms of efficiency, single-photon purity and indistinguishability [217, 260, 264]. On the other hand, it was recently shown that such schemes must be carefully tuned to satisfy the security assumptions of each quantum-cryptographic application [45]. Crucially, quantum coherences between the emitted photon-number components must be scrambled for optimal performance, which is inherently provided by some optical pumping schemes such as LA phonon-assisted excitation and TPE [45]. On top of their intrinsic security benefits, LA schemes are fairly insensitive to pump instabilities like power or polarization fluctuations, making them suitable for real-life communication networks [260, 264]. These excitation schemes are also beneficial for QDs with a complex charge environment, while other pumping schemes such as TPE can typically only address charge-neutral transitions. Moreover, unlike for neutral transitions, charged excitons can enhance polarized emission in polarized cavities, an important feature for most applications[56]. Finally, LA schemes do not require challenging single-photon polarization filtering (contrary to the resonant counterpart), and thus promise an experimentally straightforward way to obtain simultaneously high brightness and purity with high reproducibility in quantum dot fabrication and experimental setups [264].

In this work, we combine all aforementioned advantages of LA excitation in the C-band, and exploit its tunable parameters to investigate the complex dependence of brightness and purity on pump power and spectral detuning. We illustrate how this non-trivial behavior affects the security of quantum-cryptographic primitives with the example of single-photon QKD, and how the optimal operation conditions depend on the communication distance. In agreement with theoretical findings [95], our results show that the characteristics of LA excitation can be tuned to achieve the ideal photon-number statistics. This optimization is reminiscent of the mean-photon adjustment required in WCS implementations [370].

4.2.2 Results

To start investigating and optimizing our excitation parameters, it is important to note that most quantum-secured applications rely on few trusted parameters that are typically not all experimentally accessible. Here, we infer the photon-number probabilities $\{p_k\}$ from two measurements, the brightness $B = \sum_{k=1}^{\infty} p_k$ and the single-photon purity $P = 1 - g^{(2)}(0)$, where $g^{(2)}(0)$ is the second-order auto-correlation measurement evaluated at zero time delay. We use an InAs QD based

on an InGaAs metamorphic buffer layer enabling emission in the telecommunication C-band [168]. The tunable excitation is provided by a mode-locked fiber laser with a pulse length of 17(1) ps and a FWHM spectral width of 210(20) pm. The QD transition line is filtered by a set of VBG filters (FWHM = 0.2 nm). The total setup efficiency is 13 %. For more experimental details, see Sec. 4.1.1. The highest single-photon purity under pulsed LA excitation was measured as $P = 0.982$, the corresponding second-order auto-correlation measurement is shown in Fig. 4.3.

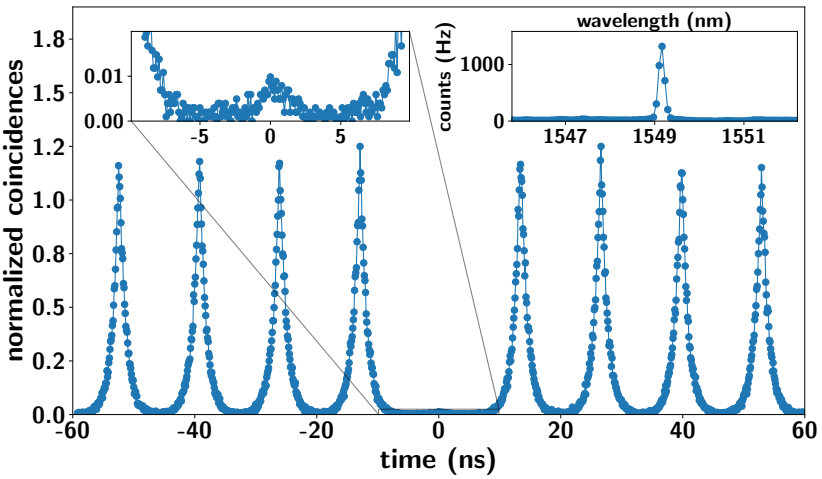


Figure 4.3: **Characterization of the positively-charged exciton transition under pulsed LA excitation.** Second-order auto-correlation measurement $g^{(2)}(\tau)$ for an excitation field strength of 0.46 a.u. and a detuning of 1.5 nm. The well-suppressed peak at zero time delay confirms the high single-photon purity ($g^{(2)}(0) = 0.018\,(1)$). Further details on the analysis of $g^{(2)}(0)$ can be found in Sec. 4.1.2. The inset shows a micro photoluminescence (μ-PL) spectrum of the studied transition, including spectral suppression of the laser with an excitation field strength of 0.89 a.u. and a detuning of 1.5 nm from the QD resonance. The figure is reprinted from [1] under a CC-BY-4.0 license.

While scanning both the power and wavelength of the pump laser, we simultaneously measure the brightness B (experimentally evaluated according to Eq. S2 and corresponding to first-lens brightness) and single-photon purity P. We then compile the results in 2D maps as shown in Fig. 4.4 (a) and (b), respectively. Due to the low phonon density at a sample temperature of ~ 4 K we excite the QD only with positive detunings $\Delta = \hbar(\omega_{\text{laser}} - \omega_{\text{dot}}) > 0$. The brightness map features a single, broad maximum around $\Delta \approx 0.8$ meV($\Delta\lambda \approx 1.5$ nm) agreeing with similar experimental findings [179, 235] and theoretical studies [234, 262]. LA excitation with sufficiently smooth pulses [263] achieves a population inversion of the QD ground and excited state if the effective Rabi splitting of the laser-dressed states, $\hbar\Omega_{\text{eff}} = \sqrt{(\hbar\Omega)^2 + \Delta^2}$, ensures an efficient exciton-phonon coupling that is characterized by the spectral phonon density $J(\omega)$ [234, 235]. The robustness

of this scheme against power and wavelength fluctuations of the excitation laser is demonstrated by the broad maximum of the brightness in Fig. 4.4 (a) and stems from the spectral width of $J(\omega)$. Thus, the large bandwidth of the phonon interaction directly benefits a stable operation of the QD source. Only for large detunings and weak fields, the phonon-induced relaxation to the exciton level fails and the brightness drops significantly. Similarly, for high powers, the effective Rabi splitting is no longer in resonance with the phonon interaction resulting in reduced brightness.

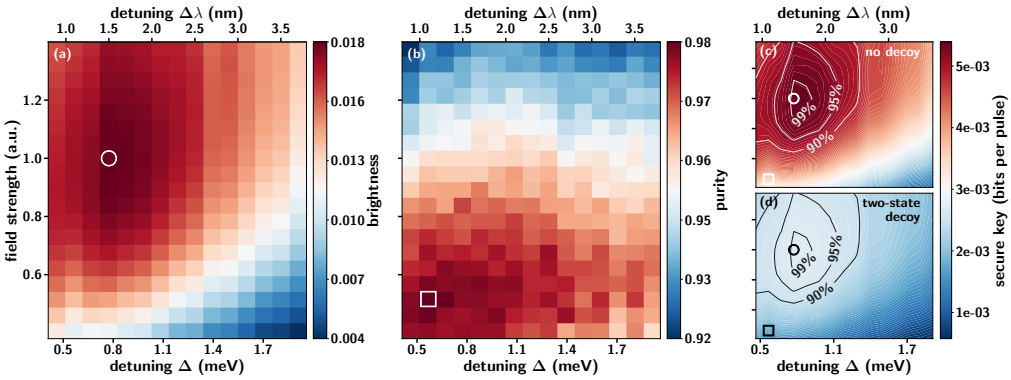


Figure 4.4: **Measured photon-number statistics and extrapolated QKD secure key bits per pulse for LA excitation.** Scanning the excitation parameters while simultaneously measuring (a) brightness $B = \sum_{k=1}^{\infty} p_k$ and (b) single-photon purity $P = 1 - g^{(2)}(0)$ of the QD emission. The white circle (square) marks the set of excitation parameters achieving the optimal brightness (purity). From the photon-number populations $\{p_k\}$, the secure key bits per pulse (SK) are calculated for zero distance based on the BB84–QKD protocol without (c) and with two decoy states (d). For more details on the parameter estimation, see Sec. 4.1.3. The equipotential lines indicate where the SK has dropped to $\{99\,\%, 95\,\%, 90\,\%\}$ of their individual SK maxima. The SK was estimated in the asymptotic limit [136], $SK = \eta_{\mathrm{sif}}[Q_1(1 - H_2(e_1)) - f(E)Q_{\mathrm{tot}}H_2(e_{\mathrm{tot}})]$, where H_2 is the binary Shannon entropy and $\eta_{\mathrm{sif}} = 1/2$. Extrapolation for two-state decoy includes an intensity modulator loss of 3 dB. Parameters for all plots are: single-photon detection error $e_{\mathrm{d}} = 0.02$, detection efficiency $\eta_{\mathrm{d}} = 0.86$, dark-count probability $Y_0 = 1.6 \cdot 10^{-6}$, error-correction code inefficiency $f = 1.2$. The figure is reprinted from [1] under a CC-BY-4.0 license.

Besides emission efficiency, the single-photon purity of the quantum-light source is crucial to the performance of cryptographic protocols [45]. Therefore, we analyze the purity P, depicted in Fig. 4.4 (b), for the same parameter range as the brightness. We identify a broad region of high purity at similar detunings but shifted towards lower powers. At large detunings, the purity degrades because the exciton state preparation via LA phonons becomes less efficient (evident by the low brightness in the same area of Fig. 4.4 (a)) and spurious contributions to the emission, including

neighboring QDs or a quasi-continuum of transitions, are no longer negligible. Considering a perfect two-level system, Ref. [95] predicts an enhanced purity for increasing excitation field strength because the phonon-induced level inversion is delayed until the end of the pulse. As a consequence, the chance of a reexcitation event during the same pulse, as it is known for resonant pumping [97, 264], would be reduced. In our experiment, however, this process competes with, and is eventually out-weighed by, the aforementioned unintended emission decreasing the purity at high powers significantly.

Interestingly, our experimental findings imply the absence of a trivial set of optimal parameters (simultaneously maximizing brightness and single-photon purity), which confirms some of the complex behaviours predicted in previous theory works [45, 95]. Instead, a careful tuning of the excitation parameters is required for each quantum-cryptographic application. Depending on the desired security of merit, the correct weighting of the photon-number populations $\{p_k\}$ used for the optimization [45] must be defined. At the same time, fluctuations in the excitation parameters produce only small changes in photon-number populations. Furthermore, optimal brightness and near-optimal purity are achieved for a pump pulse detuned by $\approx$1.5 nm from the QD transition that can be readily separated from the single-photon emission using efficient, off-the-shelf spectral filters. This simplifies source operation and optimizes brightness by removing the need for a cross-polarization setup, further underlining the practicality of LA excitation for network applications [32, 35, 344–348, 350–353, 371].

We now experimentally show how to perform the excitation parameter optimization for the example of QKD, arguably the best-known primitive in quantum communication. QKD allows two parties to establish a secret key over an eavesdropped channel [30, 35]. In that sense, the most natural figure of merit is the number of secure bits communicated per round of the protocol. This quantity can be computed from two experimental parameters: the total gain Q_{tot}, corresponding to the probability of detecting at least one photon from a given pulse sent by Alice, and the total error rate e_{tot}, indicating the fraction of states for which the wrong (polarization) detector clicks. Naturally, only the error-free single photon states contribute positively to the secure key, while the multiphoton contribution p_{m} leaks significant amounts of information. Starting from experimental data, one therefore needs to estimate the values of the single-photon gain Q_1 and the single-photon error rate e_1, which are not directly accessible. In Sec. 4.1.3, we infer these quantities in two ways: first by deriving an upper bound on the multiphoton emission probability

$$p_{\mathrm{m}} \leq \frac{1 - Bg^{(2)}(0) - \sqrt{1 - 2Bg^{(2)}(0)}}{g^{(2)}(0)} \tag{4.2-1}$$

and second by employing the two-state decoy approach [41, 372]. Compared to previous work [42], Eq. 4.2-1 gives an explicit expression for p_{m} relying only on the experimentally accessible B and $g^{(2)}(0)$ and provides additional intuition to Ref.

[221].

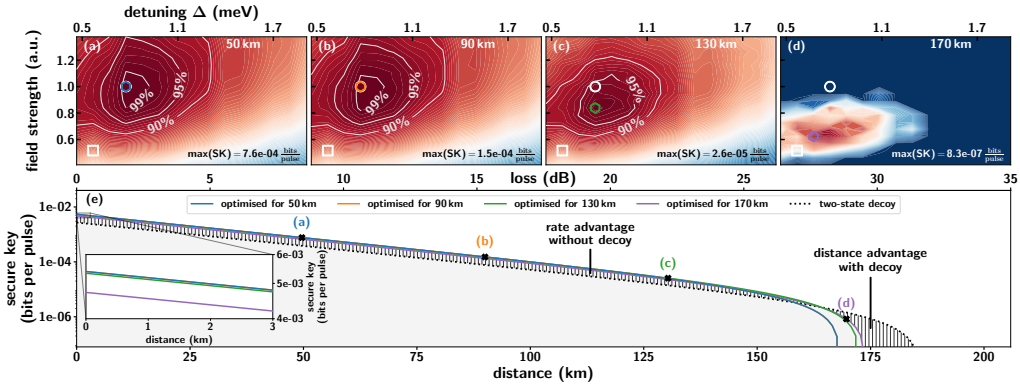


Figure 4.5: **Secure key rates of BB84–QKD for varying communication distances.** Secure key bits per pulse for the LA excitation parameter space in a standard BB84–QKD scenario for increasing channel length {50 km, 90 km, 130 km, 170 km} (a)-(d), where we assumed a fiber attenuation of $\alpha = 0.17\,\mathrm{dB/km}$, typical for the telecom C-band. The white circle (square) marks the set of excitation parameters achieving the optimal brightness (purity) as shown in Fig. 4.4, whereas the colored hexagon marks the trade-off between the two, optimizing the *SK* at the given distance. The color scale of each map is normalized to its maximum *SK* that is noted in the bottom-right corner of each map. (e) Calculating the *SK* for each highlighted parameter set from (a)-(d) as a function of transmission loss demonstrates how the tunability of LA excitation helps to adapt the emission statistics to the channel. The two-state decoy protocol reduces the *SK* by a factor of ~ 3 at short and medium distances but performs better in the high-loss regime. The parameters used to calculate the *SK* are the same as for Fig. 4.4. The figure is reprinted from [1] under a CC-BY-4.0 license.

Following the parameter estimation, we calculate the attainable secure key bits per pulse (*SK*) in the asymptotic regime for standard and decoy-state BB84–QKD for each set of excitation parameters as shown for zero communication distance in Fig. 4.4 (c) and (d), respectively. For the decoy protocol, we include a typical 3 dB loss for a high-bandwidth intensity modulator required to produce the decoys. While the qualitative dependence of the *SK* is very similar for both protocols, the performance gap is evident in the absolute values. Decoy states have been introduced to handle the risk of multiphoton contributions p_m, but since these are inherently small for QDs, introducing the constant loss of the intensity modulator outweighs the effect of an exact bounding of p_m. Furthermore, recalling that quantum cryptography with off-resonantly or two-photon excited QDs does not require any modulator for phase scrambling, adding an intensity modulator for decoy would increase the setup complexity. Comparing Fig. 4.4 finally shows that for zero distance the brightness (more accurately, p_1) dominates the *SK* map making a tight bounding of p_m even less relevant.

However, the impact of the multiphoton events on the *SK* comes into play for non-zero communication distances making the ideal set of $\{p_k\}$ no longer trivial but dependent on the channel loss. Computing *SK* maps at four distances, as depicted in Fig. 4.5 (a)-(d), visualizes the shift in source requirements. Short-distance transmissions benefit most from a bright source, whereas high-loss scenarios such as long-distance communication call for sources with high purity. Fig. 4.5 (e) then shows how these four ideal parameter sets behave over distances. The difference in performance underlines the potential of individually adjusting the excitation conditions with respect to the channel loss. Note that the joint optimization of $\{p_k\}$ by tuning the pump conditions is possible with resonant or TPE but less performant.

The maximum distance for which the generation of a secure key is still possible is of great interest for applications. Since there is no analytical expression, we state the maximal communication distance as minimal channel transmission $\eta_{\mathrm{ch}}^{\mathrm{min}}$ and find that the approximation

$$\eta_{\mathrm{ch}}^{\mathrm{min}} \approx \frac{B g^{(2)}(0)}{2} + Y_0\,, \tag{4.2-2}$$

where Y_0 is the dark-count probability, captures the breakdown of secure key generation due to multiphoton contributions well under realistic assumptions. Due to its construction (effectively lower bounding $\eta_{\mathrm{ch}}^{\mathrm{min}}$), Eq. 4.2-2 always overestimates the distance by $\sim 30\,\mathrm{km}$ (see Sec. 4.3.1). Eq. 4.2-2 also implies that, within the limits of the approximation, a brighter source reduces the maximum communication distance. Counter-intuitive at first sight, this is readily explained as the multiphoton probability p_{m} increases with the source brightness if $g^{(2)}(0)$ is unchanged (see Eq. 4.2-1). While brighter QDs further improve the *SK* at short to medium distances, one must reduce the multiphoton component when communicating over large distances. For this purpose, simply attenuating the signal before launching it into the untrusted channel is sufficient [42]. Note how this approach resembles the mean-photon number optimization used for QKD with WCS [370]. Considering a detector dark-count probability $Y_0 = 10^{-7}$, single-photon detection error $e_{\mathrm{d}} = 0.02$ and a highly pure source ($g^{(2)}(0) = 0.02$), our numerical analysis (see Sec. 4.3.1) identifies the ideal brightness for maximum distance as $B \approx 0.9\,\%$. This is well within range of today's telecom C-band QD -technology.

Finally, we remark that implementing decoy states could be advantageous in the long-distance regime even for highly pure single-photon sources such as QDs, as reflected in Fig. 4.5 (e). However, for low to moderate loss, standard BB84 outperforms the decoy-state protocol.

4.2.3 Discussion

In conclusion, we have investigated the benefits that phonon-assisted excitation of a telecom C-band QD provides for quantum-secured applications. Besides the convenient wavelength for communication applications, the InAs QDs provide a

deep confinement potential, typically spanning over several hundred millielectron volts. As a consequence, their photon-number statistics are relatively insensitive to temperature fluctuations [332]. Moreover, the source can be operated at 25 K, which is feasible for a low-cost Stirling cryostat [78].

In addition to the previously simulated low photon-number coherence [45], the robustness to environmental fluctuations [260] and the efficient single-photon filtering, we have shown that LA excitation allows to effectively optimize the photon-number statistics with respect to the desired application. This feature originates from interaction with the phonon environment and is therefore not common to resonant excitation schemes but can be exploited by tailoring the LA pumping conditions. The complex implications of phonon interactions for brightness and single-photon purity have also been theoretically predicted for idealized systems [95]. Therefore, our observations can be generalized to other QD-based sources.

As a means of improving the emission statistics independently of the excitation mechanism, temporal filtering of the signal was proposed [373, 374] but requires special hardware and comes at the price of additional loss. Moreover, we show in Sec. 4.3.2.2 that optimizing $\{p_k\}$ using only the LA pump power is as efficient as temporal filtering with a fast and lossless modulator. Only at very large distances, temporal filtering performs better since it also reduces the source brightness and detector dark counts.

However, we note that TPE can simultaneously yield a higher brightness and purity than achievable for any parameter set using the LA scheme [97]. Nevertheless, TPE—being a resonant process—is sensitive to environmental fluctuations and thus less suitable for real-world implementations.

Furthermore, we found that even for quantum light sources with inherently low multiphoton contribution, decoy states can push the maximum attainable distance in QKD. Although, in consideration of the low SK at these distances and the experimental overhead involved, we believe that decoy states are not beneficial for QD implementations.

Finally, we would like to stress that we optimized the photon-number statistics in LA excitation with respect to QKD as an example but the process is transferable to other quantum-secured applications [32, 346–353] and prone to improve their performance.

4.3 Complementary material

4.3.1 Photon-number optimization via variable attenuation

In 4.1.3 we inferred all experimental quantities required to calculate the SKR directly from the estimated photon-number populations $\{p_k\}$. To include the variable attenuation, we now model the photon loss first. To this end, we apply

$$\begin{aligned} p_0(\eta_{\mathrm{att}}) &= p_0 + p_1(1-\eta_{\mathrm{att}}) + p_{\mathrm{m}}(1-\eta_{\mathrm{att}})^2 \\ p_1(\eta_{\mathrm{att}}) &= p_1\eta_{\mathrm{att}} + p_{\mathrm{m}}(1-\eta_{\mathrm{att}}^2 - (1-\eta_{\mathrm{att}})^2) \\ p_{\mathrm{m}}(\eta_{\mathrm{att}}) &= 1 - p_0(\eta_{\mathrm{att}}) - p_1(\eta_{\mathrm{att}}) \end{aligned} \tag{4.3-1}$$

where we used $p_2 \gg p_3$ and η_{att} is the probability to transmit a photon. Note that this model is equivalent to a beam splitter with tunable reflectivity. The approach is similar to the one used in Ref. [42].

With the modified set of $\{p_k\}(\eta_{\mathrm{att}})$, we proceed as before estimating Q_1, e_1, Q_{tot}, e_{tot} and calculating the SKR. For each communication distance, we optimize the SKR over η_{att} to assess the full potential of a given source. This process resembles the optimization used to identify the ideal mean-photon number in QKD with WCS [370]. The results are depicted in Fig. 4.6, where we compare our source to idealized sources and detection.

For the experimental source parameters, we find that the maximum brightness is—by coincidence—very close to the point-wise optimized curve. Hence, reducing the effective brightness ($B_{\mathrm{eff}}(\eta_{\mathrm{att}}) = p_0(\eta_{\mathrm{att}}) + p_1(\eta_{\mathrm{att}}) + p_{\mathrm{m}}(\eta_{\mathrm{att}})$) will reduce the maximum communication distance and the SKR at short distances. However, assuming a brighter source, as in Fig. 4.6 (b), the results change drastically and the benefit of adjusting the attenuation according to the channel loss becomes clear. The enveloping curve (i.e. point-wise optimized) now features three regions with successively faster exponential decrease in the SKR. Up to ≈75 km, the best SKR is achieved without any attenuation, since the SKR is predominantly set by p_1. The next region is shaped by the continuous balancing of $p_1(\eta_{\mathrm{att}})$ and $p_{\mathrm{m}}(\eta_{\mathrm{att}})$ to optimize the SKR. Finally, at around 170 km, further attenuation cannot extend the maximum distance, as the impact of dark counts dominates.

We indicate in Fig. 4.6 (a)-(b) the maximum distance approximation obtained from Eq. 4.1-10, where the brightness corresponds to the ideal long-distance brightness. We see that Eq. 4.1-10 overestimates the distance by ≈30 km as discussed above.

To analyze the impact of distinct experimental parameters on the SKR, we simulate results for different purities and dark-count probabilities, Y_0, while the brightness is always set to $B = 1$. Only displaying the attenuation-optimized SKR for each parameter pair, we see how the position of the first inflection point is influenced by p_{m} (see Eq. 2.2-20a for fixed B and changing $g^{(2)}(0)$) whereas the distance at which the SKR curve drops for the second time is determined by both, $p_{\mathrm{m}}(\eta_{\mathrm{att}})$ and Y_0. The two inflection points can readily be associated with different causes for the SKR to break down. In the first case, the information leakage due

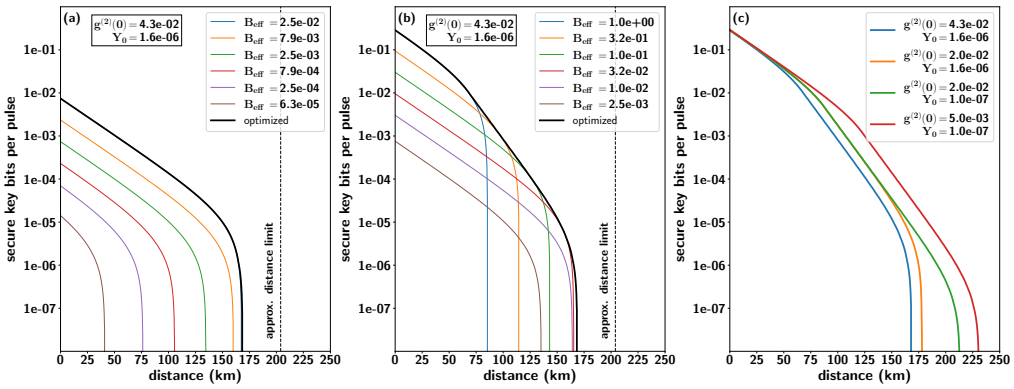


Figure 4.6: **Secure key bits per pulse over distance for variable signal attenuation.** Including a variable attenuator into the sender's setup can improve the secure key bits per pulse (SKR) at large distances. (a)-(b) The SKR as a function of distance is displayed for six effective brightness values, B_{eff}. The thick, black line represents the attainable SKR for a point-wise optimization of B_{eff}. The dotted vertical line indicates the approximated maximum distance according to Eq. 4.1-10. The dark-count probability $Y_0 = 1.6 \cdot 10^{-6}$ and purity of 95.7 % correspond to the experimental data at ideal brightness excitation conditions. For (a) the maximum brightness $B = 2.5\,\%$ corresponds to the experimental value whereas for (b) an ideal brightness, $B = 100\,\%$, is assumed. (c) Point-wise optimized curves for a set of $\{g^{(2)}(0), Y_0\}$ highlighting the different influences on the SKR. The parameters for all plots are: single-photon detection error $e_{\text{d}} = 0.02$, detection efficiency $\eta_{\text{d}} = 0.86$, error-correction code inefficiency $f = 1.2$, fiber attenuation $\alpha = 0.17\,\text{dB/km}$. The figure is reprinted from the Supplementary Material of [1] under a CC-BY-4.0 license.

to multiphoton events is too large as to permit the sifting of a secure key from the raw key, while in the second case, signal clicks are similarly probable as dark-count clicks, resulting in a high error probability e_1.

From the above findings, we conclude that a brighter but similarly pure source, as already available in the C-band [78], increases the SKR at short and medium distances but will not allow for reaching greater distances. The purity, on the other hand, has little effect on the SKR for a short channel but improves the SKR for medium distances and—in combination with a low dark-count probability—boosts the maximum attainable communication distance d_{max}.

4.3.2 Time filtering

In the Sec. 4.2, we showed how tuning the excitation conditions of the LA scheme changes the photon-number populations $\{p_k\}$ of the QD source. In a similar way, temporal filtering can be used to manipulate the photon-number statistics and improve SKRs in QKD [373, 374]. This technique can be implemented irrespective of the excitation scheme but requires a fast amplitude modulator with sufficient suppression at the sender's site. In the following, we compare the two methods

based on our experimental data.

The idea of a temporal filter is to enhance the single-photon purity of the source, and at the same time, reduce the dark-count probability at the receiver. In an experimental realization, the fast amplitude modulator is phase-locked to the driving laser and transmits only during a gating window defined by its widths τ_A and its delay t_0 to the reference input. By adjusting t_0 such that the gating window starts just before the probability of a passing photons peaks (i.e. the peak in the TCSPC measurement), one can vary τ_A to decide how much of the exponential decay trace should be transmitted. While reducing the effective brightness, this technique usually improves the purity by excluding two-photon events caused by refilling, or by reducing the single-photon contribution from neighbouring QDs.

4.3.2.1 Time filtering via post-selection

In this section, we resort to time filtering by post-processing. We, however, emphasize that a secure implementation of time filtering necessarily requires a physical gating of the signal before it is sent through the untrusted channel. Nonetheless, only investigating the potential advantage of time filtering, post-processing yields the same results as a physically gated signal stream.

To mimic this gating using post-selection, we consider only events in the auto-correlation measurement that occur within a post-selection window, $g^{(2)}(0)[\tau_\mathrm{A}, t_0 = 0]$, where τ_A is the tuning parameter. When calculating the resulting $g^{(2)}(0)$ value, we also apply the post-selection window to the n uncorrelated peaks where $t_0 = \pm\, nT_\mathrm{rep}$ (see Sec. 4.1.2). Furthermore, we compute the corrected brightness as

$$B_\mathrm{corrected}(\tau_\mathrm{A}) = B\,\frac{A_\mathrm{uncorr}(\tau_\mathrm{A})}{A_\mathrm{uncorr}(T_\mathrm{rep})} \tag{4.3-2}$$

where A_uncorr denotes the average, blinking-corrected area of the uncorrelated peak. Finally, the receiver could choose to disregard signals occurring outside an acceptance window τ_B either by gating the single-photon detectors or by post-processing. To satisfy the security requirements of the parameter estimation, $\tau_\mathrm{B} \geqslant \tau_\mathrm{A}$ must hold. For the following analysis, we set $\tau_\mathrm{B} = \tau_\mathrm{A}$ and assume the dark counts to be constant in time ($Y_0(t) = Y_0$) such that the time-filtered dark-count probability reads as

$$Y_0(\tau_\mathrm{A}) = Y_0 \frac{\tau_\mathrm{A}}{T_\mathrm{rep}}\,. \tag{4.3-3}$$

4.3.2.2 Time filtering for QKD

4.7 (a)-(b) shows the resulting brightness and purity as a function of the window widths τ_A and—for comparison—as a function of excitation power. To assess the impact on a QKD implementation, we compute the SKR over distance and optimize for the SKR by tuning either τ_A or the pump power at each step (see Fig. 4.7 (c)). For simplicity, we restrict the analysis to a fixed detuning $\Delta\lambda = 1.5\,\mathrm{nm}$.

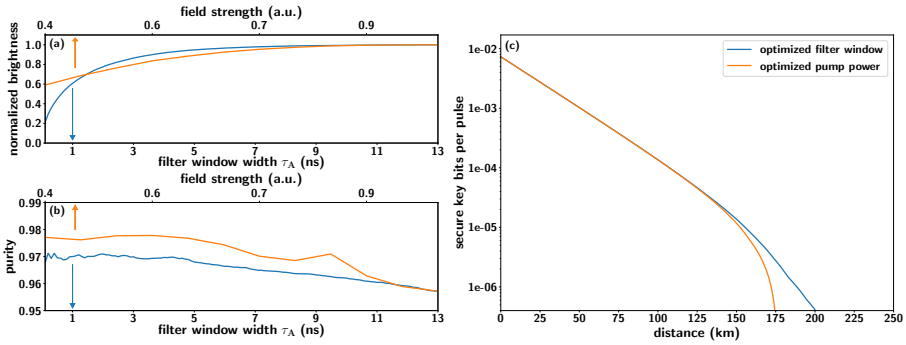


Figure 4.7: **Comparing time filtering and pump power tuning for QKD.** Both changing the excitation power or the width of a post-selection window τ_A alter the effective photon-number statistics of the QD source. For a laser detuning of $\Delta\lambda = 1.5\,\text{nm}$, we analyze their effects in terms of brightness (a), single-photon purity (b) and secure key bits per pulse in a BB84–QKD protocol (c). When the power is chosen as tuning parameter, no time filtering is applied (i.e. $\tau_A = 13.16\,\text{ns}$), whereas the field strength is set to 1 a.u. when varying τ_A. The parameters for (c) are: single-photon detection error $e_d = 0.02$, detection efficiency $\eta_d = 0.86$, dark-count probability $Y_0 = 1.6 \cdot 10^{-6}$, error-correction code inefficiency $f = 1.2$, fiber attenuation $\alpha = 0.17\,\text{dB/km}$. The figure is reprinted from the Supplementary Material of [1] under a CC-BY-4.0 license.

While both methods perform similarly over short to medium distances, we find that time filtering improves the rate at large distances. Interestingly, the enhanced SK at high loss is attributed not to the improvement of purity but to a reduction of brightness and, even more important, a reduction of dark-count probability. For the optimal filter window at 200 km, $\tau_A = 0.2\,\text{ns}$, the purity is in fact lower than for power-tuning but the simultaneous decrease of dark counts to $Y_0(\tau_A) = 2.4 \cdot 10^{-8}$—almost two order of magnitudes lower than unfiltered—preponderates. As discussed in 4.3.1, the SKR at long distances is ultimately given by the multiphoton population p_m and dark-count probability Y_0 where p_m is not just affected by the purity but also the brightness (see Eq. 2.2-20a).

However, we remark that such a short filter window requires a modulator with $> 10\,\text{GHz}$ bandwidth and would introduce significant loss at all distances. Alternative routes to decrease Y_0 include the technological advancement of SNSPDs in the long run and—already feasible today—the optimization of their biasing. Reducing the bias current (or voltage, depending on the model) of the SNSPDs affects the detection efficiency but will also lower the dark-count probability. This approach is especially appealing for long-distance communication and can be employed in combination with the power and detuning optimization of photon-number statistics in LA excitation.

5

Asymmetric two-photon response of an incoherently driven quantum emitter

This chapter presents the paper "Asymmetric two-photon response of an incoherently excited quantum emitter" [2] and contextualizes the work in light of this thesis. After diving deeply into phonon-assisted excitation and discovering the subtleties the scheme encompasses, some questions remained unanswered, which eventually sparked this project. At that time, I was working simultaneously with Lena Hansen on unveiling the temporal distribution of the multiphoton component of a resonantly excited QD—mostly as a tool to distinguish consecutive from genuine two-photon emission [338]. These efforts motivated me to consider more complex temporal analysis methods and how they might yield additional insights into the underlying physical processes. Influenced by the seminal works on re-excitation using RF [96, 97, 267] and my previous study of the dressed-state picture for [1], I expected qualitative differences when using phonon-assisted driving. Out of curiosity and anticipating benefits in single-photon applications, I decided to explore these effects experimentally.

My contribution: *The research plan and project goals were conceived and developed by myself. While the original QD excitation and collection setup used in Chapter 4 required only recalibration for the new QD sample and some minor modifications for the laser pulse preparation, I had to design and build the analysis setup for the temporal and spectral correlation measurements. I collected the experimental data, with participation of Thomas W. Sandø, and processed it with input from Lena M. Hansen and Patrik I. Sund.*

Large parts of the work presented in this chapter are accessible online [2] under a Creative Commons Attribution (CC BY 4.0) license. The manuscript is currently under peer review.

In the paper entitled "Asymmetric two-photon response of an incoherently excited quantum emitter", the re-excitation phenomenon (see Sec. 2.2.4.3) is explored for the first time using an off-resonant driving scheme. Almost a decade after the early works on re-excitation [96, 97, 267], the idea behind this project

is twofold: first, to experimentally demonstrate the strict temporal ordering of the two-photon emission and the spectro-temporal correlations manifesting in the spectral profiles and widths for each of the photons. These signatures are expected to be common to all excitation methods. On the other hand, some properties of re-excitation, including its probability or additional spectral shifts, should be affected by the optical driving scheme. Hence, the second goal of the project is to reveal qualitative differences in the two-photon response when using off-resonant pumping rather than RF.

Choosing LA phonon-assisted excitation is both practical, given our experience with the scheme, and scientifically motivated. From an application point of view, the method offers compelling advantages, especially for quantum cryptography as extensively discussed in the previous chapter. From a more fundamental perspective, a dynamic treatment of the dressed-state picture suggests that if an incoherent photon is emitted during the interaction with the laser pulse, its wavelength is redshifted. Since in a re-excitation event the early photon must be generated during the emitter's interaction with the laser pulse, it will be spectrally offset from the unperturbed QD emission line. This shift greatly facilitates suppression of the multiphoton component associated with the re-excitation signal via spectral filtering. Importantly, such a shift is not expected for RF, so LA excitation might offer another significant advantage for applications.

Since cavity designs have advanced considerably in recent years—achieving high Purcell factors and decay times as low as 25 ps [257, 296, 312]—the challenge posed by multiphoton emission through re-excitation is increasingly pressing for applications of QDs as SPSs since the spectral bandwidth of ever shorter laser pulses causes practical limitations for the laser suppression.

The following section (Sec. 5.1) provides information on the experimental setup and techniques to facilitate the reading of the original manuscript's reprint in Sec. 5.2. Additional measurements presented in Sec. 5.3 support these results and grant further insights into the process's dynamics.

5.1 Methods

The basic concepts necessary to follow the findings presented here are introduced in the Preliminaries: the dressed-state model in Sec. 2.2.3 and on its basis, pulsed LA phonon-assisted excitation in Sec. 2.2.4.2 in the limit of short pulses[1]. Sec. 2.2.4.3 elucidates the re-excitation phenomenon.

In addition, Sec. 3.2 provides a comprehensive introduction to analyzing temporal correlations. In the following, the experimental setup is detailed (Sec. 5.1.1) and the specific temporal and spectral measurements are explained (Sec. 5.1.2).

[1]Please note the slightly different notation for the dressed states. The upper (lower) dressed denoted as $|\varphi_G, N\rangle$ ($|\varphi_X, N\rangle$) in Sec. 2.2.3 is here referred to as $|\alpha, N\rangle$ ($|\beta, N\rangle$).

5.1.1 Experimental setup

For this study, the core setup shown in Fig. 3.2 is extended by a tunable bandpass filter after the pulse shaper to further reduce the laser bandwidth and consequently elongate the pulses. No other modification to the free-space excitation and collection setup is necessary. The VBG separating the QD emission from broadband background contributions has a FWHM of 120(1) GHz, and its central frequency is redshifted by ≈ 40 GHz with respect to the undisturbed QD line at ω_0 to accommodate the full spectral response expected from phonon-assisted re-excitation. The QD used for this study is embedded in a CBG cavity (FWHM≈1250 GHz), which increases the emission rate of the trion line with a Purcell factor of 4.3(0.8) [80] resulting in a decay time of 465(1) ps. Further sample details in Sec. 3.1.1.

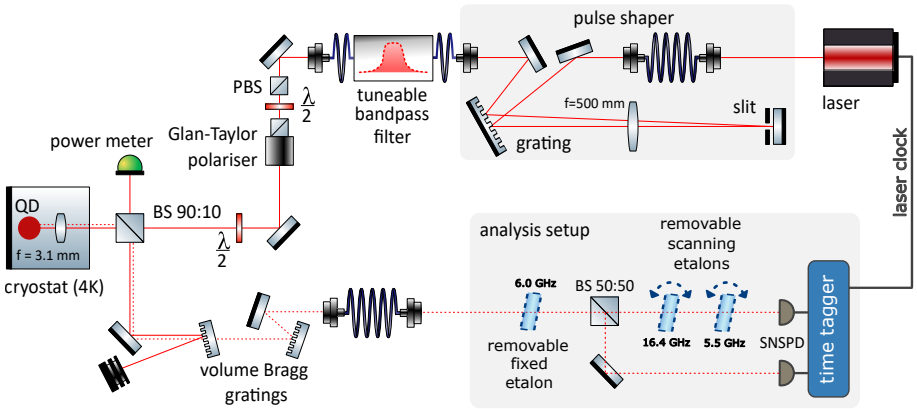


Figure 5.1: **Experimental setup including QD excitation and analysis stage.** A schematic representation of the main building blocks of the experimental setup, such as the pulsed excitation laser, pulse shaper, tuneable bandwidth filter, excitation of the QD sample, filtering of the QD transition line, and the analysis stage. Solid (dashed) red lines indicate the path of the laser beam (QD emission). polarizing beam-splitter (PBS): polarizing beam splitter, BS: beam splitter,SNSPD: superconducting nanowire single-photon detector. The figure is reprinted from [2] with permission under CC BY 4.0 license.

The analysis setup supports two types of correlation measurements: either the optical signals in both channels are identical (HBT) configuration (optionally using a narrow-band etalon with FWHM = 6.0(3) GHz before the BS), or one channel undergoes an adjustable frequency filtering, while the second channel serves as a herald (see Sec. 5.1.2). The frequency filtering is realized using two etalons (free spectral range (FSR): 125(1) GHz and 292(1) GHz; FWHM: 5.5(3) GHz and 16.4(3) GHz), each mounted on a closed-loop piezomotor-controlled rotation stage *(attocube ECR5050hs/Al/RT)* to scan their central frequency jointly. The cascaded arrangement of two etalons with non-integer-divisible FSRs greatly suppresses spectral components outside the overlapped frequency peak. The transmission peak of the joint transfer function has a FWHM of 5.3(3) GHz. The overall frequency repeatability of the analysis setup is in the few MHz range, as estimated from the

piezo stage's angle repeatability and the etalons' angle tunability. The complete setup is depicted in Fig. 5.1.

5.1.2 Temporal & spectral two-photon measurements

To uncover the temporal and spectral structure of the two-photon response, several measurements are employed. Some provide information about the joint signature, while others can distinguish which detection event was caused by a 1st photon and which by a 2nd photon[2], while ideally preserving spectral information.

The first measurement aims for the highest temporal resolution, and therefore foregoes any spectral filtering. As outlined in Sec. 3.2.2, an HBT setup augmented by a reference clock channel may be used to measure the **individual arrival-time distribution** of the 1st and the 2nd photon. Due to the inevitable timing uncertainties of the measurement apparatus, this analysis introduces a systematic assignment error that depends on the time scales in the experiment; specifically, the laser pulse duration Δt_L, the decay time τ, and the SRF measured as the correlation between the two optical channels. However, for the parameter set used in the experiment, this error is estimated to be smaller than 1% (see Supplementary Note 4 of [2]).

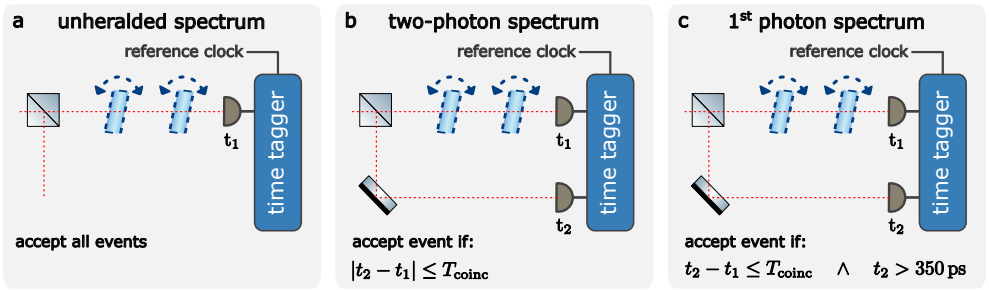


Figure 5.2: **Schematic spectro-temporal analysis setup.** A schematic representation of the logic configurations used to extract the various spectra from the same data set. The physical setup for implementing the frequency sweep on channel 1 is shown in more detail in Fig. 5.1. **(a)** Channel 2 is unused; simply, all clicks registered on channel 1 are accepted. This configuration mimics the operation of a normal spectrometer. **(b)** Postselecting events that showed a photon in channel 2 arriving within the coincidence window $T_{\text{coinc}} = 3\,\text{ns}$. This configuration yields the two-photon spectrum as all single-photon events are discarded. Because it is equally likely that the 1st photon is frequency filtered while the 2nd photon acts as the herald or vice versa, the measurement is balanced. **(c)** Accepting the photon on detector 2 only as a heralding event if it registered "late" ensures with high probability that the event is caused by a 2nd photon. Hence, a click on the frequency-filtered channel 1 within the coincidence window indicates a 1st photon.

Adding a narrow spectral filter in front of the HBT BS blurs the timing infor-

[2]A reminder to the notation introduced in Sec. 3.2.2: referring to a 1st or 2nd photon not only describes the temporal order but also implicitly states that both were generated in response to a single laser pulse as indicated through a coincidence with sufficiently small coincidence window.

mation but allows to measure the probability that two photons close to ω_0 (within the filter width) were emitted within one repetition period. This arrangement is used to demonstrate that filtering removes the re-excitation contribution for phonon-assisted driving almost entirely (see below Fig. 5.6).

Turning to the frequency-resolved measurements, Fig. 5.2 illustrates the three analysis techniques used in the paper to infer the different spectral profiles from one measurement. The adjustable frequency filter is scanned stepwise over the relevant wavelength range, and at each point, the raw time tags are recorded for a fixed integration time. In all cases, the spectrum is recovered by plotting the number of accepted events against the filter's central frequency. Which type of postselection (using channel 2) achieves which spectral subset is explained in the following.

Unheralded spectrum Ignoring all events registered in channel 2 and simply recording the count rates of channel 1 for each filter angle recovers the spectrum as it would be measured on a standard spectrometer (Fig. 5.2a).

Two-photon spectrum By defining a coincidence window T_{coinc} and accepting only clicks on channel 1 whenever a photon was also detected on channel 2, the measurement postselects two-photon events (Fig. 5.2b). Assuming the re-excitation contribution dominates the multiphoton probability, this measurement reveals the spectral two-photon signature of re-excitation. A coincidence is equally likely to be caused by a 1^{st} or 2^{nd} photon being registered in the frequency-filter channel, while the other one serves as a herald in channel 2. Therefore, the spectral weight of the 1^{st} and 2^{nd} photon is balanced for this measurement.

1^{st} photon spectrum Capturing the spectrum of only one of the two photons requires additional postselection. However, the narrowband scanning filter significantly broadens the arrival-time distribution recorded in channel 1. For our parameters, the exponential broadening is much smaller than the coincidence window T_{coinc}, but it is on the order of the laser pulse width, which makes the timing too unreliable for determining which photon was emitted first. Therefore, information from channel 1 is used only to determine coincidences, not to evaluate temporal ordering. Instead, the unaffected, high-precision time tags of channel 2 are gated based on the reference clock, discarding all clicks occurring before time t_{G}[3]. The gating ensures with high probability that only events originating from a 2^{nd} photon are accepted in channel 2. These prefiltered clicks representing 2^{nd} photons are then used to indicate coincidences with channel 1. Consequently, these detections in channel 1 must originate from spectrally filtered 1^{st} photons (assuming

[3]In the experiment, t_{G} is chosen such that the probability of detecting a 1^{st} instead of a 2^{nd} photon in channel 2 is $\approx 2.5\%$. These estimates are based on the arrival-time distributions measured separately for the same laser parameter.

negligible other multiphoton contributions), and the frequency scan recovers its spectrum.

Since the method discards all early 2^{nd} photons, the accepted coincidences represent only a subset of all two-photon events. Hence, the method implicitly assumes that the spectral profile of the 1^{st} photon is unaffected by the emission time of the 2^{nd} photon.

5.2 Manuscript reprint

Quantum emitters promise to emit exactly one photon with high probability when pumped by a laser pulse. However, even in ideal systems, re-excitation during a laser pulse causes the consecutive emission of two photons, thereby limiting single-photon purity. Although the probability and properties of re-excitation are largely determined by the optical excitation method, until now only resonant driving has been studied. Here, we demonstrate qualitative differences in the process arising from phonon-assisted excitation—a scheme standing out for its robustness and straightforward spectral suppression of scattered laser light while preserving highly indistinguishable emission. In contrast to previous studies under resonant driving, we measure not only the $g^{(2)}(0)$ as a function of pulse length but also resolve the distinct temporal and spectral shape of each of the photons, report an asymmetric two-photon spectrum and uncover correlations between the emission time and wavelength, which are unique to phonon-assisted pumping. On the fundamental level, we show that the spectrum arising from re-excitation provides direct access to the Rabi frequency of an incoherently driven quantum dot. On the application side, we use the asymmetric spectral response to selectively suppress multiphoton noise from re-excitation, ensuring a high single-photon purity regardless of the laser pulse length and thus enhancing implementations across quantum cryptography and quantum computing.

5.2.1 Introduction

Studying the rich physics of light–matter interaction not only provides valuable insights into fundamental quantum processes but also inspires the design of applications that benefit from these processes. Among the variety of quantum emitters that allow such studies, quantum dots (QDs) have proven to be a powerful platform for foundational experiments in quantum optics [63–65, 338] and applications in both quantum communication [44, 46, 73, 76, 176] and computing [58, 60, 72, 375], thanks to simultaneously high single-photon efficiencies [55], low multiphoton contribution [97, 187] and highly indistinguishable emission [56, 218]. A large part of the success stems from the high level of control gained by embedding these artificial atoms into a semiconductor host structure, which allows the integration of photonic cavities to improve collection efficiencies and provide Purcell enhancement, as well as electrical diodes [376, 377] to stabilize the charge environment and tune the emission wavelength. However, the benefits of stable semiconductor integration also come with a coupling to the surrounding environment—most notably to charge noise, spin reservoirs, and the phonon bath. When unmanaged, such interactions can impair the properties of the emitter, as evidenced by damped Rabi oscillations [253], diminished photon indistinguishability [378, 379], or reduced spin coherence times [380, 381]. Yet, when understood and used to our advantage, these very interactions also offer opportunities, including spin refocussing [88], control over the charge state [377], or phonon-mediated schemes to incoherently excite the QD [217].

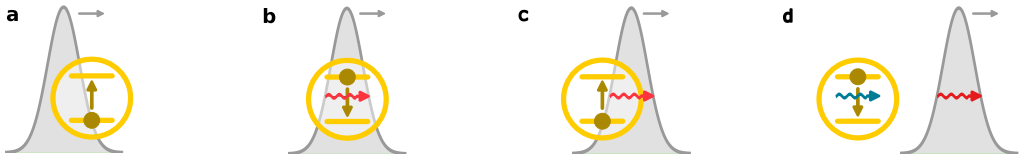


Figure 5.3: **Schematic time series of a re-excitation event.** Modelling the QD as a two-level system, the re-excitation process from a single laser pulse is described using a representative quantum trajectory, broken down into four distinct events. **(a)** Early during the pulse interaction with the laser pulse, the QD is excited and **(b)** decays shortly after, emitting a 1st single photon (red arrow). **(c)** Towards the end of the interaction with the laser pulse, the QD is excited a second time. **(d)** The laser pulse and the 1st photon have passed, when the QD spontaneously decays again and emits the 2nd photon (blue arrow). Thus, the 1st and 2nd photon are emitted consecutively in a strict time-ordered manner. The figure is reprinted from [2] under a CC-BY-4.0 license.

In recent years, longitudinal acoustic (LA) phonon-assisted pumping [234] has attracted significant interest [1, 45, 256, 330] due to its intrinsic robustness against laser instabilities [1] and the straightforward suppression of laser background via spectral filtering. The latter renders LA excitation ideal for generating polarization-entangled linear cluster states from QDs [57, 256], while applications in quantum communication additionally benefit from the phase randomization required in many cryptographic protocols and provided by the coupling to the phonon bath [45].

As a result of continuous improvements in QD fabrication, fundamental challenges—rather than technical imperfections—are now beginning to limit the achievable performance [174, 280]. Among these, the impact of re-excitation, where a single laser pulse may induce the emission of multiple photons, is particularly severe as it compromises the single-photon purity of the source. In resonantly driven two-level systems, the re-excitation probability scales proportionally to the laser pulse length and the radiative decay rate [97]. However, except for initial theoretical work that focussed on adiabatic undressing [95], the phenomenon is largely unexplored for phonon-assisted pumping. At the same time, a comprehensive understanding of the re-excitation dynamics is becoming increasingly important as recent cavity designs yield Purcell factors exceeding 25 [257], leading to emitter decay times $< 30\,\mathrm{ps}$. While the accelerated decay improves the photon indistinguishability and supports faster clock rates, it also increases the multiphoton emission caused by re-excitation.

Here, we present experimental results exploring both the temporal and spectral signatures of the re-excitation process under LA phonon-assisted pumping. We individually resolve the time and frequency modes of the first and second emitted photon and reveal significant differences compared to the resonant counterpart. Importantly, we observe a spectral shift for the first photon that is unique to phonon-assisted driving and can be interpreted as dynamic Stark tuning caused by the

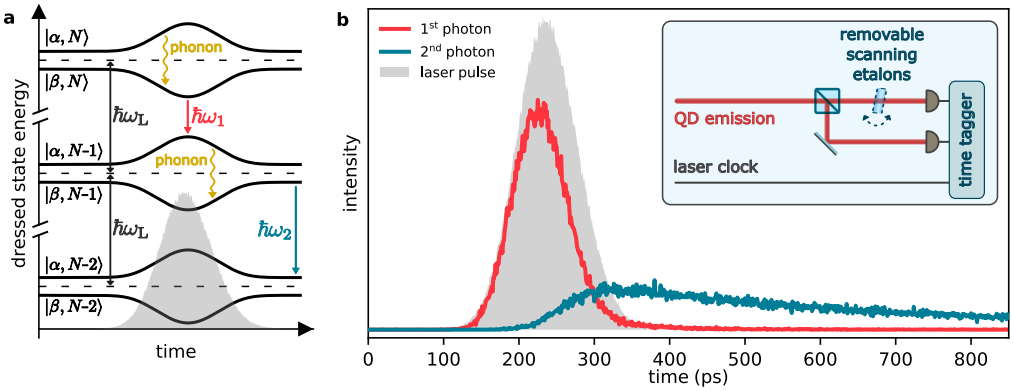


Figure 5.4: **Re-excitation under phonon-assisted driving. (a)** For one exemplary quantum trajectory, the re-excitation process under LA driving is schematically depicted in the energy diagram of the dynamical dressed states reduced to three manifolds of the laser field. The laser pulse is shown in gray. The emission of a first phonon sets the system into the lower dressed state $|\beta, N\rangle$ and, shortly after, but when the system is still dressed, a 1st photon is emitted at ω_1 by decaying into $|\alpha, N-1\rangle$. Thereafter, another phonon emission brings the system into the lower dressed state of this manifold ($|\beta, N-1\rangle$) which is adiabatically transformed into the exciton state once the laser pulse exits. The emission of a 2nd photon at ω_2 completes the process. The phonon coupling efficiency is governed by the phonon spectral density evaluated at the instantaneous Rabi frequency [263]. **(b)** Time-histogram of the 1st and 2nd photon showing the different temporal shapes. The laser pulse in gray is displayed for time reference. Inset: schematic detection setup in HBT configuration, additionally equipped with the electronic laser clock for advanced correlations. The scanning etalons indicated with dashed lines are only used for the spectrally resolved measurements. The figure is reprinted from [2] under a CC-BY-4.0 license.

excitation laser. The shift allows us to assess the Rabi frequency of an incoherently driven QD. Moreover, we demonstrate how these Stark-induced time-frequency correlations can be exploited to limit the multiphoton contribution irrespective of the laser pulse length only using spectral filtering.

5.2.2 Results

Re-excitation under phonon-assisted driving

The re-excitation process is conceptually explained in Fig. 5.3 for one exemplary quantum trajectory. Independent of the pumping scheme, when driving the QD with a short laser pulse, there is a chance that it is excited early during the interaction (a) and quickly decays by emitting a 1st photon (b) such that the QD can be excited again by the same pulse (c) and emit a 2nd photon at a later time (d). The causality of the process requires that whenever two photons were created from one laser pulse, the 1st photon must have been emitted while the pulse was

still interacting with the QD (by referring to the 1st or 2nd photon, we imply that indeed two photons were detected). Following this reasoning, temporal filtering has been proposed to restore a high single-photon purity [220, 374] but is limited to post-selection, because the fast time scales are infeasible for current electro-optic modulators.

To understand re-excitation in the phonon-assisted scheme, we briefly recall the dressed-state picture, where the QD is treated as a two-level system with transition frequency ω_0, and the laser frequency ω_L is detuned by $\delta\omega_L = \omega_L - \omega_0$ (here, we limit ourselves to $\delta\omega_L > 0$). The interaction of the QD with the external laser field is captured by a light–matter Hamiltonian in the rotating frame, which gives rise to a new set of eigenstates, the so-called dressed states $|\alpha, N\rangle$ and $|\beta, N\rangle$, where N denotes the number of laser photons [382]. At the beginning of the pulse, the ground state coincides with $|\alpha, N\rangle$ but as the laser field increases, it admixes an excitonic (ground state) component to $|\alpha, N\rangle$ ($|\beta, N\rangle$), and thereby activates phonon interactions between the dressed states [235]. When the level splitting of the dressed states—determined by the effective Rabi frequency $\Omega_{\mathrm{eff}}(\mathrm{t}) = \sqrt{\Omega^2(\mathrm{t}) + \delta\omega_{\mathrm{L}}^2}$—matches the spectral phonon density, $|\beta, N\rangle$ can be efficiently populated through phonon emission [234]. As the laser pulse exits the system, $|\beta, N\rangle$ is adiabatically transformed back to the exciton state, which completes population inversion. Recombination at a later time then produces a single photon at ω_0 [235]. However, if a photon is emitted while the laser pulse is still present, we must consider the transition $|\beta, N\rangle \rightarrow |\alpha, N-1\rangle$ in the dressed-state picture.

The complete LA re-excitation process is schematically outlined in Fig. 5.4a and we note two major differences compared to previously studied resonant systems. First, the excitation requires not only the absorption of a laser photon but also the emission of a phonon, which (depending on the phonon decay rate) might delay the excitation and thereby lower the re-excitation probability. Second, the eigenenergies of $|\alpha\rangle$ and $|\beta\rangle$ change antisymmetrically with the instantaneous laser field strength, such that the transition frequency inherits the time dependence of the driving pulse

$$\omega_{\mathrm{QD}}(t) = \omega_L - \Omega_{\mathrm{eff}}(t)\ . \tag{5.2-1}$$

This time-varying frequency shift is reminiscent of the dynamic Stark effect [383], where a laser pulse alters the energy structure of a few-level system. However, whereas in previous works the Stark shift was deliberately induced using a secondary pump laser [383–385], in the LA scheme the excitation laser itself causes a redshift of the QD emission.

Temporal and spectral signatures of re-excitation

In the experiment, we use an epitaxially grown In(Ga)As QD emitting in the telecom C-band, which is embedded in a circular Bragg grating cavity with wide-band enhancement and excited using a blue-detuned laser with adjustable pulse length. The QD transition is attributed to a charged exciton with a lifetime of

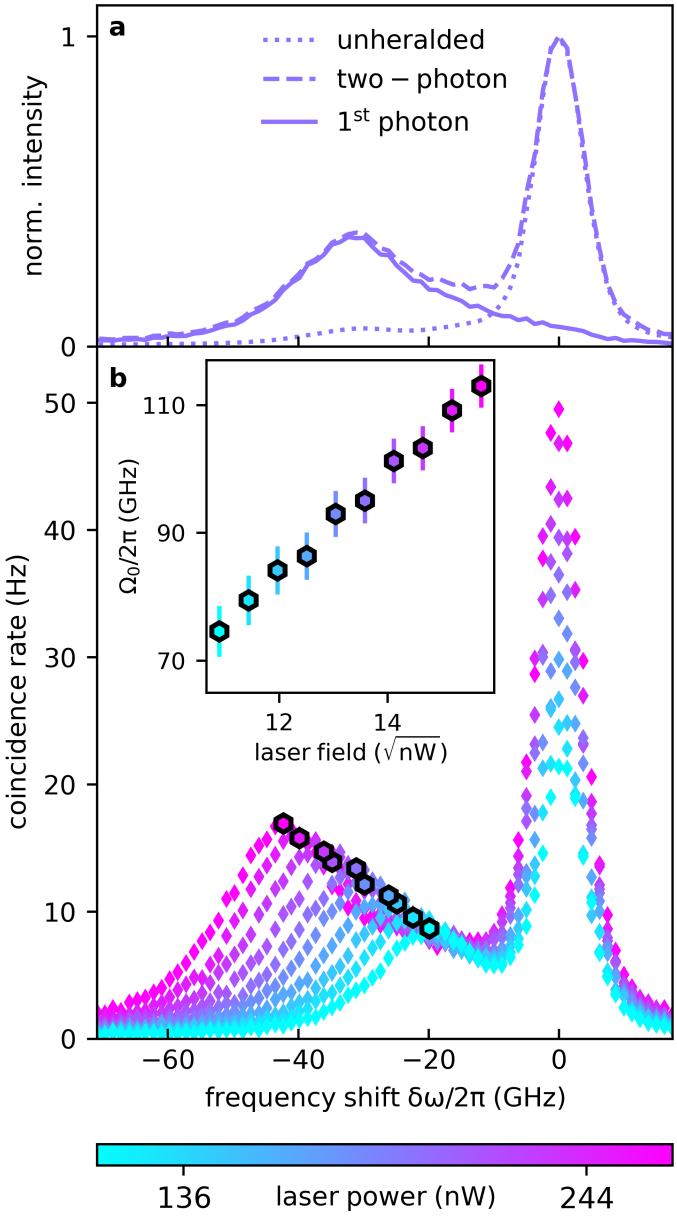


Figure 5.5: **Re-excitation spectrum for varying pump power.** A scanning frequency filter ($\Delta\omega_f/2\pi = 5.3(3)$ GHz is used to measure the QD spectrum based on different gating logics. Unheralded: no gating, all events detected on the spectrally filtered channel. Two-photon: all events on the spectrally filtered channel coinciding with a heralding photon within $\Delta\tau_c = 3$ ns. 1st photon: the heralding event is used only if it is registered > 350 ps after the laser clock. The laser pulse length is $\Delta t_L = 80(1)$ ps. The displayed data is broadened by the transfer function of the frequency filter. **(a)** The laser power is set to 185 nW and individual data points have been omitted for visual clarity. **(b)** The laser power is varied (see color bar), and the two-photon spectrum (as above) is measured. The spectral position of the side peak is marked for every power and used to calculate the Rabi frequency according to Eq. 5.2-2, which is shown in the inset as a function of the laser field. The figure is reprinted from [2] under a CC-BY-4.0 license.

$\tau_{QD} = 465(1)\,\mathrm{ps}$, indicating a Purcell factor of 4.3(8) [80]. More details on the sample and laser preparation can be found in the Methods.

Starting with the temporal analysis of the re-excitation signal, we excite the QD with 80(1) ps long Gaussian laser pulses, suppress the scattered laser, and isolate the QD line from broadband background noise. The emission is detected in a HBT configuration and correlated with the laser clock (see inset Fig. 5.4b). Instead of performing a standard second-order autocorrelation $g^{(2)}(\tau)$, we check for coincidences between the two optical channels and record the timestamp of the early and late clicks for each coincidence. Sorting this way, we separate events caused by 1^{st} photons from those caused by 2^{nd} photons, allowing us to resolve the dynamics of the two-photon emission individually. The resulting histograms, displayed in Fig. 5.4b, represent the temporal shape of the 1^{st} and 2^{nd} photon and thus provide the first direct experimental confirmation of the temporal ordering associated with the re-excitation process. Since the 1^{st} photon must be emitted when the laser is still present, its temporal profile strongly resembles that of the laser. The 2^{nd} photon, on the other hand, can be emitted during the full lifetime, which leads to the Gaussian-broadened exponential decay expected from spontaneous emission. Even though the probability distributions overlap in time, two photons originating from a single laser pulse are always emitted into orthogonal time modes (see Supplementary Information), and only averaging over many repetitions induces the statistical overlap.

Turning to the spectral analysis, we employ a tunable filter that enables us to frequency-resolve either the overall emission or, in a heralded manner, the two-photon signal (see inset Fig. 5.4b and Sec. 5.1). Recording the count rate of the filtered path while scanning the filter frequency reconstructs the spectrum averaged over all photons, whereas using the second detector as a herald post-selects on two-photon events, effectively measuring the two-photon spectrum (see Fig. 5.5a). The unheralded data features a prominent peak at ω_0 representing undisturbed single-photon emission. However, because of the length of the laser pulses, the re-excitation effect is already noticeable without any heralding and shows itself in the broad shoulder on the low-frequency side, which we attribute to the 1^{st} photon. In the heralded configuration, where the spectra of the 1^{st} and 2^{nd} photon contribute equally, the shoulder then becomes a second but broader peak. To further support the assertion of the low-frequency peak to the 1^{st} photon, we refine the gating logic and ignore all heralding events occurring before $t = 350\,\mathrm{ps}$ (compare Fig. 5.4b), that is, all photons emitted within the laser pulse duration. Counting coincidences in this way ensures that the photon in the filtered path must have been the 1^{st} photon. The resulting spectrum overlaps perfectly with the broad side peak observed in the two-photon spectrum (see Fig. 5.5a). The increased bandwidth of the 1^{st} photon is a consequence of its narrower temporal shape (recall Fig. 5.4b) and is also predicted for resonant driving [96]. The shift from the natural frequency ω_0, on the other hand, is a signature of the transition between the dressed states $|\beta, N\rangle$ and $|\alpha, N-1\rangle$ and is unique to re-excitation under phonon-assisted

pumping. Note that the cavity's influence on these measurements is minimal, since its mode is much wider than the observed frequency shift.

Measuring the Rabi frequency for incoherent excitation

Following Eq. 5.2-1, we can express the maximum frequency shift as $\delta\omega_{\max} = \delta\omega_L - \sqrt{\Omega_0^2 + \delta\omega_L^2}$. As expected from the dynamic Stark effect, the shift increases for stronger laser–QD interactions, that is, larger external fields or less detuning. Rearranging the equation as

$$\Omega_0 = \sqrt{\delta\omega_{\max}(\delta\omega_{\max} - 2\delta\omega_L)} \tag{5.2-2}$$

yields a remarkable result: we can infer the Rabi frequency—fundamental to coherent light–matter interaction in two-level systems—from an incoherent process. Instead of relying on Rabi rotations or coherent two-photon scattering (the origin of the Mollow-triplet [224]), here we exploit the signatures of the phonon-assisted re-excitation spectrum.

To experimentally verify these findings, we measure the two-photon spectrum for increasing laser power. Our results, presented in Fig. 5.5b, clearly demonstrate that the redshift of the 1$^{\text{st}}$ photon intensifies with power, while the spectral position of the 2$^{\text{nd}}$ photon is unchanged. In the parameter regime studied here, and in agreement with the dressed-state model, the shift increases approximately linearly with the external laser field. Extracting Ω_0 from the induced redshift according to Eq. 5.2-2 and plotting it against the applied laser field shows an excellent linear trend (see inset Fig. 5.5b), confirming our understanding of the process. Note that coherent two-photon scattering, as recently shown also for pulsed excitation [336, 337], would produce a similar, but importantly, symmetric spectral response. And although both processes could coexist, we show in the Supplementary Information that for the used laser parameters, efficient phonon coupling is ensured and the incoherent process dominates.

Lastly, we stress that the spectral overlap of the two photons reduces with increasing laser power. This is particularly helpful when frequency filtering the QD emission to reduce the multiphoton contribution, a method we introduce in the next section.

Restoring high single-photon purities

Recalling that for applications of QDs as SPSs the presence of an additional 1$^{\text{st}}$ photon diminishes the single-photon purity, we now leverage the strong spectro-temporal correlations found in Fig. 5.5 to our advantage. Using a frequency filter slightly wider than the spontaneous emission line at ω_0, we expect to greatly suppress the 1$^{\text{st}}$ photon. To ensure that we do not simply reduce the multiphoton noise caused by other spectrally distinct processes (such as neighboring QDs), we vary the laser pulse length across an order of magnitude while maintaining a constant laser power of 185 nW. In this way, we tune the re-excitation probability but leave other processes unaffected.

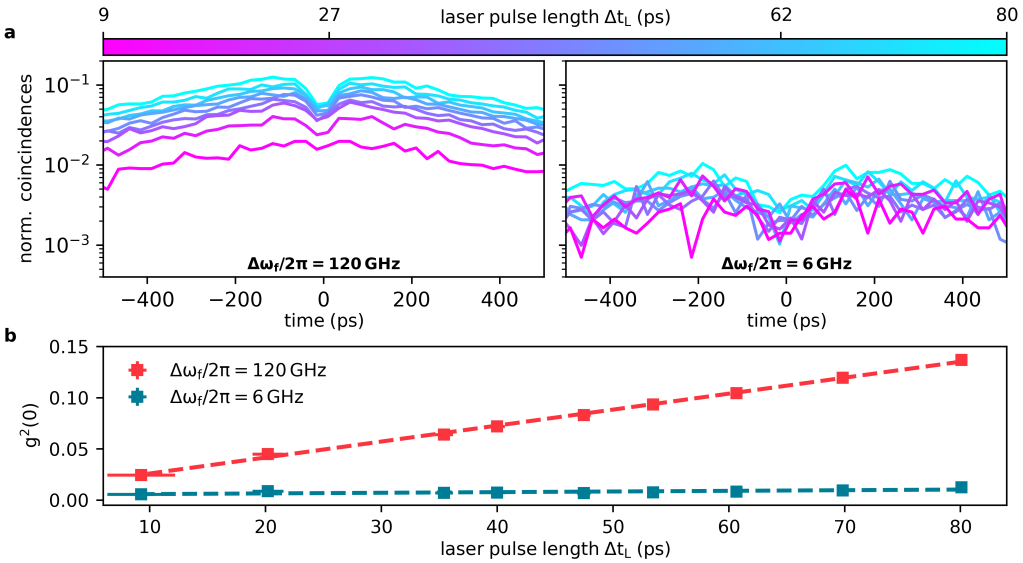


Figure 5.6: **Reducing the multiphoton contribution by frequency filtering.** For a constant laser power of 185 nW, the pulse length is varied and the second-order autocorrelation $g^{(2)}(\tau)$ is measured for each step. **(a)** The zoomed-in central peak of the coincidence counts normalized to Poissonian level is shown as a function of the delay τ for a 120(1) GHz VBG filter (left panel) and a 6.0(3) GHz etalon (right panel). **(b)** The $g^{(2)}(0)$ is evaluated by dividing the area of the correlated peak by that of the averaged uncorrelated peaks. The error bars are derived from Poissonian statistics and are too small to be visible. The uncertainty in laser pulse length arises from the deconvolution with the system response function of 26(1) ps (see Sec. 5.1). The linear fit $g^{(2)}(0, \Delta t_L) = A \cdot \Delta t_L / \tau_{QD}$ yields $\mathrm{A}_{120\mathrm{GHz}} = 0.72 \pm 0.01$ and $\mathrm{A}_{6\mathrm{GHz}} = 0.03 \pm 0.01$. The figure is reprinted from [2] under a CC-BY-4.0 license.

Fig. 5.6a compares the zoomed-in, correlated peaks around $\tau = 0$ of a standard $g^{(2)}(\tau)$ for various pulse lengths. The QD signal shown in the left panel is only isolated from broadband background, but for the measurements displayed in the right panel a 6.0(3) GHz etalon centered on ω_0 is added. Unfiltered, the multiphoton probability clearly increases for longer pulses and the $g^{(2)}(\tau)$ also features the characteristic re-excitation dip at $\tau = 0$. The depth is limited by the system response function of the detection stage (see Sec. 5.1). On the other hand, the filtered emission exhibits a drastically reduced multiphoton response and, importantly, is independent of the laser pulse length. Only for the longest pulse, a slight increase is notable. Since we maintain a constant laser power throughout the scan, the effective Rabi frequency decreases for longer pulses as $\Omega_{\mathrm{eff}} \propto 1/\sqrt{\Delta t_L}$, which results in a smaller redshift of the 1st photon and therefore in a larger spectral overlap with the spontaneous emission line. Interestingly, we also observe a (significantly wider) dip in the $g^{(2)}(\tau)$ for the filtered signal, which could be caused by charge-carrier refilling [190] or some residual transmission of the 1st photon.

This attribution is supported by time-resolved measurements (see Supplementary Information), which show that 1st photons transmitted through the etalon are emitted at the very beginning of the laser pulse (when the dynamic frequency shift is still small), resulting in a larger delay between the 1st and 2nd photon.

The quantitative comparison of the unfiltered and filtered multiphoton probability, displayed in Fig. 5.6b, demonstrates the stark difference in $g^{(2)}(0)$ when the laser pulse length is tuned. The multiphoton component for the filtered emission is virtually unaffected for all pulse durations studied here, reaching more than 1/6 of the QD decay time. On the other hand, the $g^{(2)}(0)$ of the unfiltered signal increases linearly with Δt_L, indicating that re-excitation dominates the multiphoton probability, which is in good agreement with work on resonantly driven two-level systems [97].

Finally, we note that the multiphoton response related to re-excitation scales with $\Delta t_L/\tau_{QD}$, whereas the efficiency of the proposed filtering method depends on the frequency shift of the 1st photon, which increases for shorter pulses (see above). Consequently, the filtering of multiphoton noise becomes more effective, benefiting sources with fast decay times in particular.

5.2.3 Discussion

We have presented an experimental study that provides nuanced understanding of the re-excitation dynamics and, at the same time, constitutes the first work on the phenomenon under phonon-assisted driving. In addition to verifying the strict temporal ordering of the emission predicted by theory [96], we find strong spectro-temporal correlations for the 1st emitted photon, which are induced by the dynamic Stark effect and lead to an asymmetric two-photon spectrum unique to phonon-assisted pumping. The implications of our results are both fundamental and practical. We have shown that the Rabi frequency of an incoherently driven QD can be extracted from the re-excitation spectrum, and on the application side, we have demonstrated how the spectral distinguishability of the two photons can be leveraged to efficiently remove multiphoton noise.

As advances in high-Purcell cavities continue to reduce QD lifetimes, managing the re-excitation signal with conventional methods becomes increasingly challenging, and the frequency filtering introduced here offers a compelling new approach. In applications, the laser pulses are significantly shorter than those used in this fundamental study, inducing a larger spectral shift, which allows for higher transmission without compromising the single-photon purity. Moreover, our findings reveal an additional intrinsic advantage of phonon-assisted driving when used with narrow-band cavities. During the laser–QD interaction, the dynamic frequency shift detunes the QD transition from the cavity mode, effectively canceling the Purcell enhancement. This mechanism delays photon emission until the laser pulse subsides and greatly suppresses re-excitation.

We have shown that a deeper understanding of phonon-assisted excitation offers unique advantages for applications; however, other key aspects of the QD–phonon

interactions are yet to be explored. In particular, the impact of the phonon spectral density, which could be tailored [386] using phononic cavities, presents a promising research direction that can create new opportunities in quantum communication and photonic quantum computing.

5.3 Complementary material

To further support the findings and interpretations presented in Sec. 5.2, additional measurements for other, both positive and negative, laser–QD detunings are presented and explained in Sec. 5.3.1. The time-varying emission wavelength throughout the QD's interaction with the laser pulse is investigated more deeply in Sec. 5.3.2 by analyzing the frequency-filtered TCPC signals.

The results and their explanations are reprinted with minor modifications from Supplementary Notes 3 and 4 of the original publication [2].

5.3.1 Laser–QD detuning variations

As discussed in Sec. 5.2, other signatures of dynamically dressed states were recently observed [336, 337]. These experimental findings represent the extension of the well-known Mollow-triplet [224] into the regime of pulsed driving and bear some similarity to the spectra presented here. However, the underlying processes are fundamentally different.

As explained in 2.2.3, there are four optical transitions between any two manifolds of laser-dressed states, all of which are allowed by optical selection rules [222]. Two of them are degenerate and match the frequency of the driving laser, the other two appear at $\omega_L \pm \Omega_{\text{eff}}(t)$ and are energetically only allowed for a phonon-photon or photon-photon transition. In a coherent process, the photons must appear in pairs of a low- and high-frequency photon (thus obeying the conservation of energy) and are the result of two-photon scattering from the laser field. In the absence of an asymmetric cavity or for a sufficiently large cavity mode, this leads to a symmetric spectral response.

On the other hand, incoherent interactions involving the phononic environment also facilitate single-photon transitions at those frequencies (more accurately, single-phonon single-photon transitions). However, depending on the sign of the laser–QD detuning $\delta\omega_L > 0$ ($\delta\omega_L < 0$), the emission (absorption) of a phonon is required, and therefore only one of the transitions is efficiently supported.

In Fig. 5.7, we present the spectra obtained from our QD, where the laser is either blue- or red-detuned and its power is scanned. The scattered laser is filtered with a VBG. For $\delta\omega_L/2\pi = 131\,\text{GHz}$ (Fig. 5.7a), a low-frequency side peak is clearly visible, and its separation increases with the laser field. The absence of any emission at $\omega_L + \Omega_{\text{eff}}(t)$, is also a strong indication toward the single-photon transition. On the other hand, if $\delta\omega_L/2\pi = -131\,\text{GHz}$ (Fig. 5.7b), we identify two side peaks at $\omega_L \pm \Omega_{\text{eff}}(t)$, which are much weaker in intensity than in Fig. 5.7a. Since at $T = 4\,\text{K}$ the phonon occupation is low, the incoherent, phonon-assisted single-photon transition becomes significantly less efficient for a red-detuned laser pulse and symmetric two-photon scattering dominates.

Combining these findings, we conclude that the low-frequency side peak investigated in Sec. 5.2 for $\delta\omega_L > 0$ is (almost exclusively) a result of phonon-assisted emission.

In addition to the laser power scan presented in Fig. 5.5b, Fig. 5.8 shows

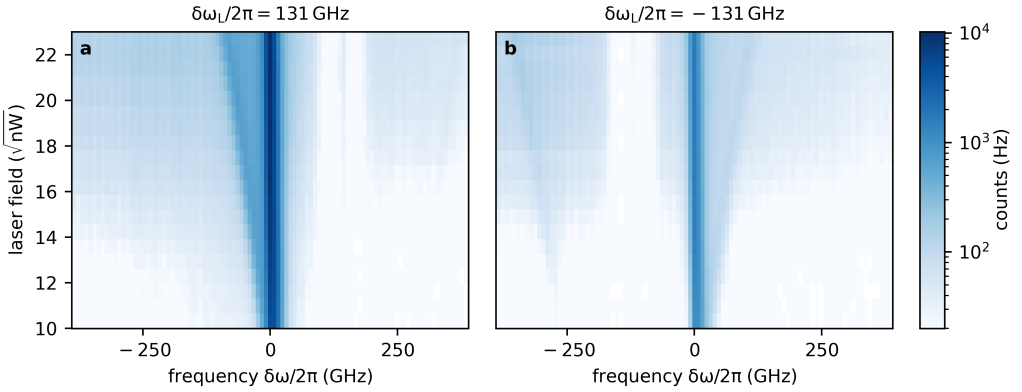


Figure 5.7: **QD spectrum for blue- and red-detuned laser at varying powers.** For 80(1) ps long laser pulses, the QD emission is measured with a spectrometer after the scattered laser is filtered using a 120(1) GHz VBG. The spectrometer resolution is ≈10 GHz per pixel. **(a)** For a 131 GHz blue-detuned laser, only the low-frequency side peak is visible, whereas the emission for a red-detuned excitation laser (**b**) features two side peaks, which are symmetric with respect to the laser frequency. The figure is reprinted from the Supplementary Material of [2] under a CC-BY-4.0 license.

how the two-photon spectrum changes when $\delta\omega_L$ is varied at a constant power of 250 nW. As expected from dressed-state emission (see Eq. 5.2-2), a smaller laser–QD detuning induces a stronger level splitting, and thus a larger frequency shift of the 1st photon, that is, the side peak.

5.3.2 Frequency-filtered time traces

The dressed-state model suggests that the level splitting, and thus the emission frequency, changes throughout the laser pulse interaction, as captured in Eq. 5.2 -1 where $\Omega_{\text{eff}}(t) = \sqrt{\Omega^2(t) + \delta\omega_L^2}$ is the effective Rabi frequency. Using the setup depicted in Fig. 5.1 without the fixed etalon but including the scanning etalons, we can measure the temporal response for different filter frequencies ω_f. Simply correlating the optical signal of the filtered HBT path with the laser clock yields a time trace for each specific filter frequency $\delta\omega_\text{f} = \omega_\text{f} - \omega_0$ as depicted in Fig. 5.9. Throughout the panels, the central frequency of the filter is shifted stepwise to larger detunings, as indicated by the color bar.

Initially, when the filter overlaps well with ω_0, we find an early rise in signal counts at the very beginning of the laser pulse (indicated in gray) that plateaus first (or even drops slightly) and is then followed by a second much stronger signal increase and a slow exponential decay. This behavior starts to change both qualitatively and quantitatively, when red-detuning the filter from ω_0. The early rise lasts longer, reaching a higher signal level, and the following drop in counts in more pronounced. The second rising-edge peaks earlier, so that the maximum

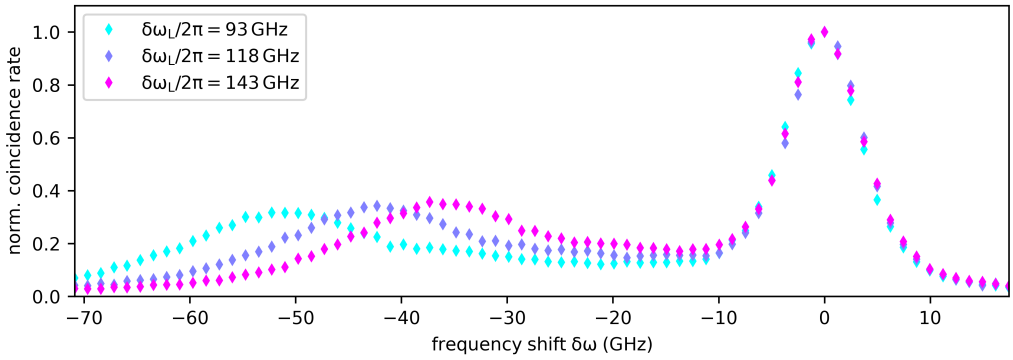


Figure 5.8: **Two-photon spectrum for varying laser–QD detunings.** For 80(1) ps long laser pulses with 250 nW, the two-photon spectrum is measured in the heralded configuration explained in the Main Text. The three measurements correspond to different laser- QD detunings, $\delta\omega_L$. For clarity, all error bars were omitted. The figure is reprinted from the Supplementary Material of [2] under a CC-BY-4.0 license.

signal strength reduces and, more importantly, the falling edge becomes more Gaussian (see e.g. darkest curve in the second panel). This trend continues as the redshift enlarges and the signal drop after the first rise slowly washes out until there is only a single smooth rise and fall in detected counts (see curve with highest peak in the fourth panel). Also note that the global peak height increases again throughout the fourth panel. In the last panel, the shape of the signal does not change anymore as it steadily declines in signal strength.

Recalling that a Gaussian laser pulse causes the QD transition frequency to first decrease until the pulse peaks and then increase again (see Eq. 5.2-1), the main features of our findings are readily explained. For most filter positions, there are two points in time when they overlap with the time-varying emission frequency; once during the rising edge of the laser field and a second time when the field strength subsides. Only the maximum frequency shift occurs for just a single point in time, namely when the laser pulse peaks.

This behavior is captured in Fig. 5.9 where for all filter settings $\delta\omega_\mathrm{f} < \delta\omega_\mathrm{max}$ (for these laser settings, $\delta\omega_\mathrm{max}/2\pi \approx -45\,\mathrm{GHz}$, see Fig. 5.8), two peaks appear in the time trace. The temporal separation of these peaks shrinks with increasing $\delta\omega_\mathrm{f}$ until they merge into one at $\delta\omega_\mathrm{f} = \delta\omega_\mathrm{max}$. Additionally, we clearly see that photons near ω_0 exhibit an exponential decay as expected from spontaneous emission, whereas photons at significantly redshifted frequencies must be emitted while the laser pulse is present and therefore inherit the temporal shape of the laser pulse.

To analyze and understand the process quantitatively, one must also incorporate the time-dependent probability of the emitter being in the excited state. For phonon-assisted excitation, involving non-Markovian processes, this is commonly done using numerical methods, including time–path integrals [387]. However, some intuition

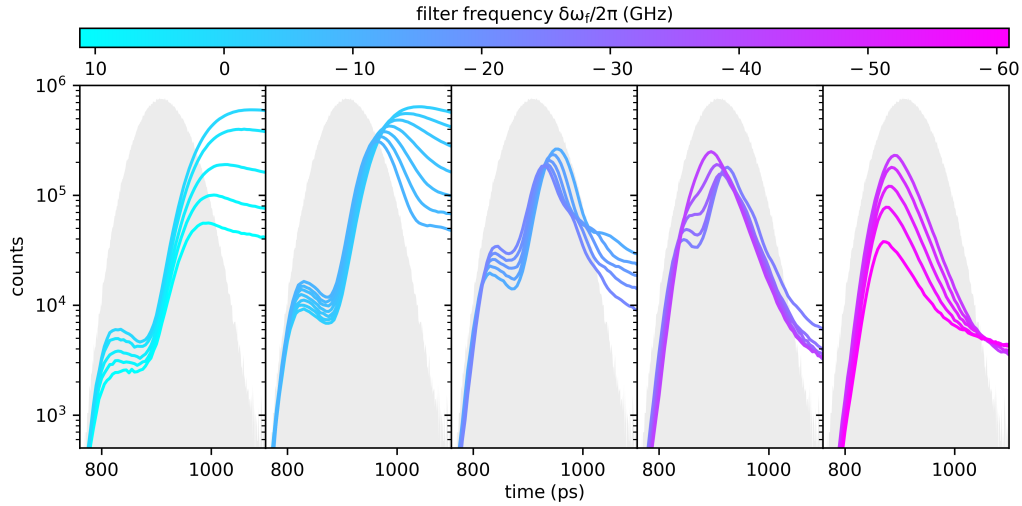


Figure 5.9: **Histogram of frequency filtered QD emission.** For 80(1) ps long laser pulses with 250 nW and $\delta\omega_L/2\pi = 125$ GHz, the QD emission is filtered using the set of scanning etalons shown in Fig. 5.1. The filtered channel is correlated with the laser clock to construct the histograms for every filter position $\delta\omega_\mathrm{f} = \omega_\mathrm{f} - \omega_0$. The curves are colored according to $\delta\omega_\mathrm{f}$ as indicated by the color bar. Due to the Fourier-transform limit, the original temporal signal is exponentially broadened by the 5.3(3) GHz wide Lorentzian filter. The errors were calculated using Poissonian statistics and are too small to be visible. The figure is reprinted from the Supplementary Material of [2] under a CC-BY-4.0 license.

about the excitation probability can be gained from considering the phonon spectral density $J(\omega)$. As indicated in 2.2.4.2, there is a Rabi frequency Ω_eff that maximizes the phonon coupling and overshooting causes less efficient phonon coupling, thus less efficient population inversion of the QD. This non-monotonic evolution of the phonon coupling throughout the laser pulse can lead to strong phonon interactions at the beginning and end of a laser pulse but weak interaction at its peak field amplitude. The reduced phonon coupling at large effective Rabi frequencies was recently also shown to revive the Rabi oscillations of a coherently driven QD [253]. While it is necessary to incorporate these effects to model the (frequency-filtered) emission probability across large parameter ranges, we note that the detuning and driving field used in Fig. 5.9 are small enough such that $\Omega_\mathrm{eff}(t)$ never exceeds the optimal phonon coupling. Therefore, the time-dependent phonon scattering rate does not feature an inflection point, and the dip (or kink) seen in Fig. 5.9 can be attributed solely to frequency-selective detection.

6

Experimental single-photon limited-knowledge proofs

This chapter presents the work "Experimental single-photon limited-knowledge proofs" and contextualizes it in light of this thesis. After the extensive investigation of LA phonon-assisted excitation presented in Chapter 4 and previous work of the group on enhancing the performance of cryptography protocols using QDs [45], we were keen to experimentally implement a single-photon primitive for which QDs outperform PDSs. At this time, Rebecca Hofer was developing a new beyond-QKD protocol with Mathieu Bozzio that allows one party to convince another that they have the solution to a hard problem without revealing the full solution. Since the security of the quantum protocol arises from QBC in the bounded-storage model—which in a practical implementation is strongly impacted by the single-photon purity of the source—this protocol promised to be an ideal match. Using the next generation of telecom C-band QDs embedded in a CBG cavity from our collaborators at the University of Stuttgart guaranteed a tenfold increase in single-photon rates while maintaining the low multiphoton noise required for this project.

My contribution: *The single-photon polarization encoding setup was designed by Michal Vyvlecka and myself, and I built, optimized, and characterized its performance. With assistance from Francesco Giorgino, I conceived and implemented electrical synchronization within the QD lab. I built the receiver's station, implemented automated polarization compensation, and wrote code to synchronize the experiment across both labs. The experimental data was collected and processed by myself. Finally, I supported Rebecca Hofer and Mathieu Bozzio in including experimental imperfection in the security analysis of the protocol. With input from Tobias Guggemos, Rebecca Hofer and I implemented the computational experiments and analyzed the results.*

This chapter presents unpublished results that are not yet accessible elsewhere. The manuscript is currently under peer review.

The manuscript entitled "Experimental single-photon limited-knowledge proofs" introduces and experimentally implements a new cryptographic protocol across

two university buildings in downtown Vienna using a QD emitting in the telecom C-band as SPS.

The primitive enables two mistrustful parties to verify that one holds the solution to a mathematically hard problem without disclosing the full solution to the other party. In this way, the verifying party (from hereon, the Verifier) does not gain sufficient knowledge of the solution, which would enable them to impersonate the first party (from hereon, the Prover) in a later interaction. Though this might sound a bit abstract at first, such *selective disclosure* schemes find application across a broad range of use cases, e.g., anonymous online voting. In contrast to previous works that relied solely on classical resources and required additional assumptions, including the existence of one-way functions or no-signaling constraints [388–394], our approach harnesses quantum states to evade these restrictions. The protocol is built on a QBC providing *i.t.* security against a bounded adversarial quantum memory, and it is designed so that the underlying problem can be reused multiple times without compromising security.

The protocol can be implemented using the standard set of single-photon BB84 states, ensuring compatibility with existing QKD quantum hardware. However, in a mistrustful setting, multiphoton noise significantly degrades protocol performance, and we demonstrate an advantage of highly pure SPSs across almost all realistic scenarios by extrapolating our source characteristics to larger distances.

In the experiment, we actively prepare the polarization state of the photons, transmit them over more than 600 m of deployed fiber, and design and realize an automated polarization compensation routine. We achieve a quantum bit error rate (QBER) of less than 0.5 % over more than 80 hours.

6.1 Methods

The preliminaries of this thesis introduce the basic principles of a QD (Sec. 2.2), its optical excitation (Sec. 2.2.4), and how to encode quantum information in the polarization of the generated single-photons (Sec. 2.1.3.2). The concept of a QBC in the noisy-storage model is explained in Sec. 2.1.4.2.

In the following, the experimental setup used to implement the protocol is outlined in more detail (see Sec. 6.1.1); in particular, we elaborate on characteristics of the SPS (Sec. 6.1.1.1), polarization encoding (Sec. 6.1.1.2), state projection (Sec. 6.1.1.3), and synchronization (Sec. 6.1.1.4). In addition, the choice and structure of the mathematically hard problem is explained in Sec. 6.1.2.

The complementary material of Sec. 6.3 contains more details about the sifted stream of experimental data (Sec. 6.3.1), an overview of the protocol's security analysis (Sec. 6.3.2), the protocol's performance for an extended parameter space (Sec. 6.3.3), as well as the problem's generation algorithm and our benchmarking method (Sec. 6.3.4). The complete security analysis is printed in Appendix B.

6.1.1 Experimental setup

The scheme of the experimental setup is shown in Fig. 6.1 and the excitation and collection section closely resembles the setup introduced in Sec. 3.1.2. Importantly, an integrated electro-optic modulator (EOM) in pulse-picking configuration (*ixBlue MXER-LN20-30dB*) is included after the pulse shaper, and it is phase-locked to the laser clock (using a FPGA, for more details see Sec. 6.1.1.4) such that only one out of 152 pulses is transmitted, effectively reducing the repetition rate to $\nu_{\mathrm{rep}} = 499.67\,\mathrm{kHz}$ (for brevity, the repetition rates are approximated throughout the section as 76 MHz and 500 kHz, respectively). The suppression contrast is $\approx 1:500$.

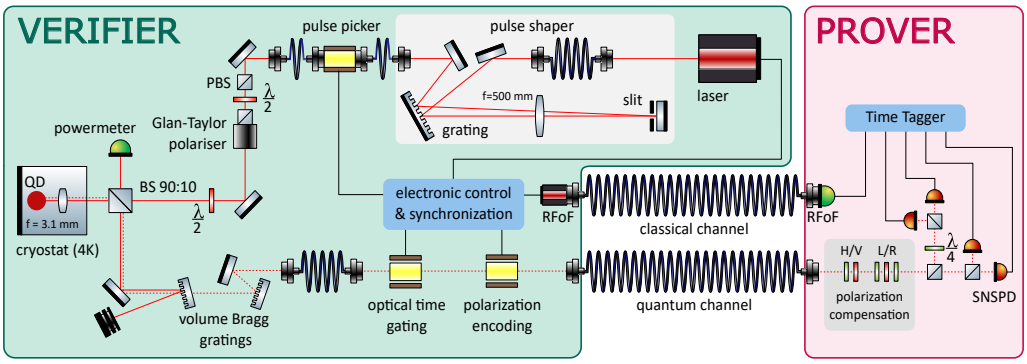


Figure 6.1: **Scheme of the experimental setup.** A schematic representation of the main building blocks of the experimental setup, such as the pulsed excitation laser, pulse shaper, pulse picker, excitation of the QD sample, filtering of the QD transition line, polarization encoding, transmission via communication channel, polarization compensation, and state projection and detection. The electronic control & synchronization block is presented in more detail Fig. 6.6. PBS: polarizing beam splitter, BS: beam splitter, radio frequency over fiber (RFoF): electronic-to-optical synchronization signal converter, SNSPD: superconducting nanowire single-photon detector.

After the single photons are generated and separated from the stray laser light as explained before (see Sec. 3.1.2), they are launched via a $f = 10\,\mathrm{mm}$ lens collimator into the free-space time-gating and polarization-encoding stage consisting of two EOMs that are phase-locked to the laser clock (for more details, see Sec. 6.1.1.4) and operate at ν_{rep}. During a gating window $\Delta t_{\mathrm{gate}} = 30\,\mathrm{ns}$, the first EOM (*EKSMA MP2*) flips the polarization from horizontal to vertical, causing the single-photon pulse to reflect at the successive PBS. Outside the gating window, the photons are dumped through the PBS's transmission port. This additional gating limits unwanted information leakage into the quantum channel with a suppression ratio $> 1:500$ (see Fig. 6.4). Within the encoding window $\Delta t_{\mathrm{enc}} = 100\,\mathrm{ns}$, a variable high-voltage (HV) pulse is applied at the second Pockels Cell (*EKSMA MP1*), changing the retardance and thus continuously rotating the polarization state of

the transmitted light from horizontal to left-circular to vertical to right-circular (see Sec. 2.1.3.2). Importantly, the amplitude of the HV pulse is modulated by an analog control signal (*Siglent SDG2000X*), allowing for the encoding of an arbitrary sequence of polarization states. Finally, the polarization-encoded single photons are coupled to an SMF using a $f = 10\,\text{mm}$ lens collimator, exhibiting an overall transmission efficiency (including fiber coupling) in the encoding stage of $\eta_{\text{enc}} = 64(4)\,\%$. For more details on the encoding stage, see Sec. 6.1.1.2.

Leaving V's trusted lab via the quantum channel—a deployed SMF in the metropolitan area of Vienna—the encoded states reach the other faculty building with an efficiency of $\eta_{\text{ch}} = 86(3)\,\%$. The link consists of two parts, a $\approx 40\,\text{m}$ long section connecting two laboratories within V's building and the second section of $\approx 640\,\text{m}$ bridging the distance to P's station. The polarization projection and single-photon detection performed by P are detailed in Sec. 6.1.1.3. Finally, P tags all photon clicks alongside the classically shared reference clock and selects the timing events corresponding to the solution bits they want to commit to.

6.1.1.1 QD characteristics

The SPS is based on an epitaxially grown In(Ga)As quantum dot coupled to a circular Bragg grating cavity as detailed in Sec. 3.1.1. This study uses the same QD as in Chapter 5, but between cooling cycles the emission wavelength has shifted to $\lambda = 1550.9\,\text{nm}$. Thanks to the broad cavity mode, the decay time is almost unchanged (see Fig. 6.2). Using the same excitation conditions as in the protocol implementation, the laser is detuned by 1.4 nm and the laser power set to $250(15)\,\text{nW}$, which is just below the saturation power for this QD as can be seen in Fig. 6.2c showing a count rate curve as expected for LA excitation when scanning the excitation power [1]. Without correcting for any loss or imperfect detection efficiency, the polarized end-to-end efficiency is $\eta^{\text{end}}_{76\text{MHz}} = 4.8(1)\,\%$.

The autocorrelation measurement $g^{(2)}(\tau)$ under the same excitation conditions is presented in Fig. 6.3a, yielding $g^{(2)}_{76\text{MHz}}(0) = 1.76(5)\,\%$ for the full integration window of $\Delta t_{76\text{MHz}} = 13.17\,\text{ns}$. For this sample, blinking is almost negligible with an optically active time of $\approx 95\,\%$. Fig. 6.3b shows the autocorrelation measurement at 500 kHz where the fiber-based pulse picker could not fully suppress the unpicked laser pulses, causing a residual, small emission probability in these time bins. Comparing the area of the uncorrelated peaks at $\tau = \pm 2\,\mu\text{s}$ (integrated over 30 ns) with the area of the correlated peak at $\tau = 0\,\mu\text{s}$, we find $g^{(2)}_{500\text{kHz}}(0) = 1.86(11)\,\%$.

In cryptography applications, the potential information leakage caused by emitting multiple photons within a single signal bin, $\Delta t_{\text{sig}} = 1/\nu_{\text{rep}}$, must be estimated and included in the security analysis. Hence, when evaluating the $g^{(2)}(0)$ of the SPS, the integration window is usually set equal to the signal bin, $\Delta t_{\text{int}} = \Delta t_{\text{sig}}$. In this way, the probability of sending two or more photons within one signal bin, quantified by p_{m}, can be upper bounded directly using $g^{(2)}(0)$ (see Sec. 2.2.2.4).

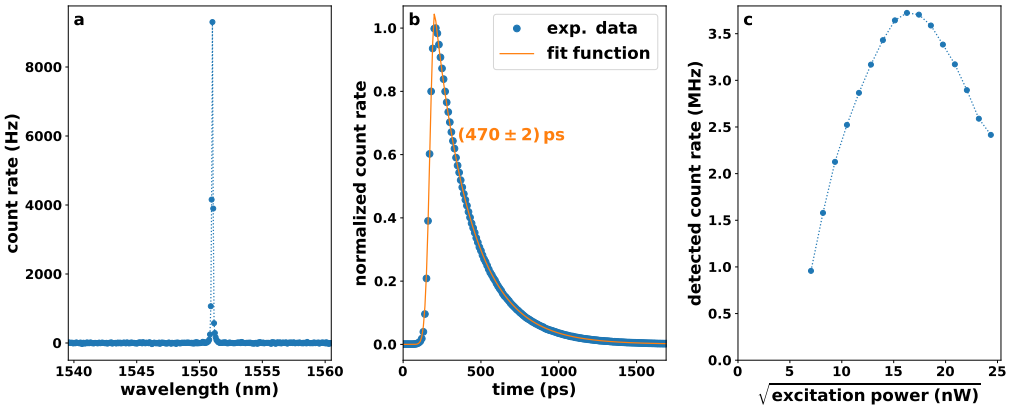


Figure 6.2: **Characteristics of the QD emission.** All measurements are performed with pulsed LA excitation at 76 MHz, a blue-shifted detuning of 1.4 nm and a laser bandwidth of 210(10) pm. For a) and b), an excitation power of 250(15) nW is used. a) The spectrum of the QD transition after the reflected laser light is suppressed. The spectrometer resolution is ≈ 70 pm per pixel. b) The mono-exponential fit function (solid orange line) yields a decay time of 470(2) ps. c) Scanning the pump power, the typical count rate behavior for LA excitation is observed. The uncorrected end-to-end efficiency of the QD in the characterization setup is 4.8(1) %.

In our implementation, however, the information leakage is drastically reduced by the additional optical gating introduced by the free-space pulse-picking Pockels cell, which transmits only within the gating window $\Delta t_{\mathrm{gate}} = 30\,\mathrm{ns}$. Therefore, we evaluate $g^{(2)}(0)$ for $\Delta t_{\mathrm{int}} = \Delta t_{\mathrm{gate}}$ and account for any additional multiphoton component outside Δt_{int} separately by analyzing the suppression of the free-space pulse picker. In Fig. 6.4, the optically gated QD signal is shown as a histogram where the red-shaded area illustrates the gating window Δt_{gate}. While the first and last peak within Δt_{gate} originate from excitations caused by imperfectly suppressed laser pulses (occurring ±13.17 ns before/after the main peak), the remaining two peaks are associated with reflections within the setup. Note that all four spurious peaks are three orders of magnitude smaller than the central emission peak. Outside the gating window, the signal is almost exclusively background-limited. The only significant contribution arises from a peak ≈ 200 ps after the main event and can be attributed to the ringing of the HV signal during switching between polarization states. When comparing all clicks detected within Δt_{gate} and the integrated events outside the window, we find a signal suppression ratio $SR > 1:500$.

To account for all information leakage outside the gating window, we introduce

$$g^{(2)}_{\mathrm{eff}}(0) = g^{(2)}_{500\mathrm{kHz}}(0) + SR \leq 2.06(11)\,\% \,, \tag{6.1-1}$$

which we then use in the security analysis to estimate the multiphoton contribution p_{m}.

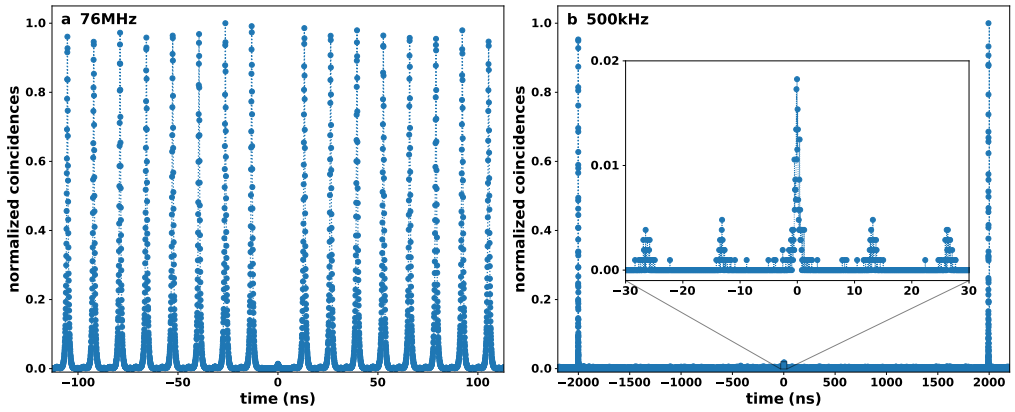


Figure 6.3: **Autocorrelation measurements of the QD.** Autocorrelation for pulsed LA excitation detuned by 1.4 nm for the standard repetition rate of 76 MHz and the protocol repetition rate of 500 kHz. The pump power is chosen just below the saturation power. a) For an integration window of 13.17 ns, we find $g^{(2)}_{76\mathrm{MHz}}(0) = 1.76(5)\,\%$. b) For an integration window of 30 ns, we find $g^{(2)}_{500\mathrm{kHz}}(0) = 1.86(10)\,\%$.

6.1.1.2 Polarization encoding

The encoding of the classical pseudo-random strings into the photons' polarization is performed by the second free-space EOM. The incoming light is linearly polarized along an axis tilted by 45° with respect to the Pockels Cell axes system, such that it excites the ordinary and extraordinary modes of the crystal equally. Tuning the crystal's retardance by applying a high voltage induces continuous rotation of the output polarization vector around the S_3 axis of the Poincaré sphere (see Sec. 2.1.3.2). Hence, given enough HV range, all polarization states on the great circle spanned by $\{|H\rangle, |L\rangle, |V\rangle, |R\rangle\}$ can be realized.

In the experiment, however, the Pockels Cell implements a residual rotation even when no nominal HV is applied. Furthermore, the available HV range (up to 2.7 kV) suffices only to perform a 270° rotation. Given these imperfections, not all four eigenstates of the rectilinear and circular bases, $\{|H\rangle, |L\rangle, |V\rangle, |R\rangle\}$, can be generated (in our case $|H\rangle$ is not accessible, see Fig. 6.5). Although the residual rotation could be corrected with a set of wave plates at fixed angles, it would complicate the alignment process of the Pockels Cell. Instead, we select a different pair of conjugate bases that satisfy the same requirements as the standard BB84 states. In particular, the set of states must limit the success probability of discrimination attacks to the theoretical minimum. Mathematically, this means that any two states picked from conjugate bases are orthogonal, while states from the same basis are antiparallel. We set up a cost function based on this condition and, from the set of all encodable polarization states, pick the four states that minimize the cost function.

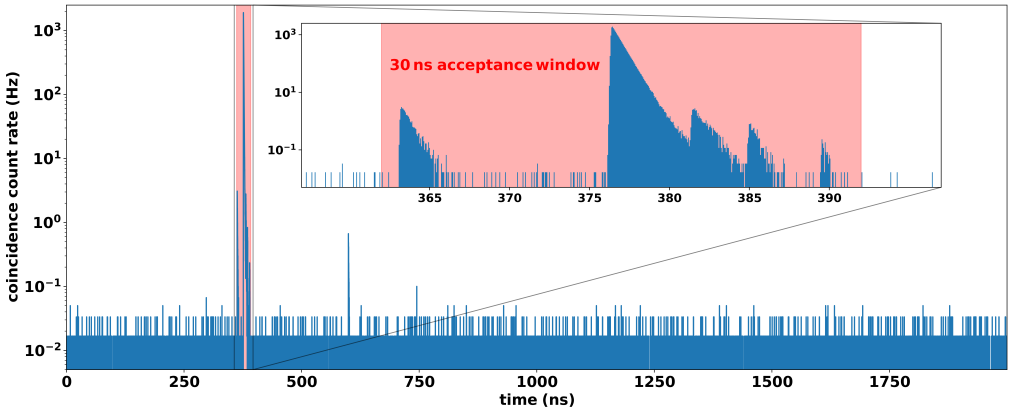


Figure 6.4: **Optical channel gating.** The signal leakage outside the gating window $\Delta t_{\text{gate}} = 30\,\text{ns}$ is measured under the same conditions as used for the verification protocol run. Correlating the clicks from an SNSPD with the reduced laser clock (500 kHz) temporally resolves the signal shape. Summing all clicks outside Δt_{gate} and dividing by the sum of clicks within Δt_{gate} yields a suppression ratio $SR > 1 : 500$. The suppression ratio is included when upper-bounding the probability of sending a multiphoton state.

In Fig. 6.5a, all encodable states are displayed on the Poincaré sphere, and the chosen optimal states are indicated. The data is acquired with classical light by stepwise increasing the control voltage and recording the output states on a polarimeter. The mean orthogonality is 0.0015(13) and the mean anti-parallelity is −0.986(6). The degree of polarization (DOP) is shown in Fig. 6.5b for the same scan, demonstrating an average DOP $> 99\%$. Even though the specified DOP accuracy of the polarimeter is only $\pm 1\%$, the results already indicate the high quality of the polarization states.

6.1.1.3 Compensation and detection

P uses a standard free-space, passive BB84 state projection stage that is complemented by a set of five wave plates to reverse any polarization rotation introduced by the channel (see Fig. 6.1). Using this arrangement, a fast, straightforward, and automated compensation routine is implemented according to the procedure described in [395]. Conceptually, the first motorized quarter-wave plate (QWP) and half-wave plate (HWP) ensure that photons sent in the rectilinear basis are projected correctly and the subsequent set of QWP, motorized HWP and QWP, where both QWPs are fixed at 45°, only correct the phase between $|H\rangle$ and $|V\rangle$ ensuring the ideal projection of the circular states but leaving the rectilinear states unchanged. Note that the correction not only compensates the drift but also removes the offset rotation caused by the residual retardance of the encoding Pockels cell, such that the ensuing projection acts on the standard BB84 states. To prevent

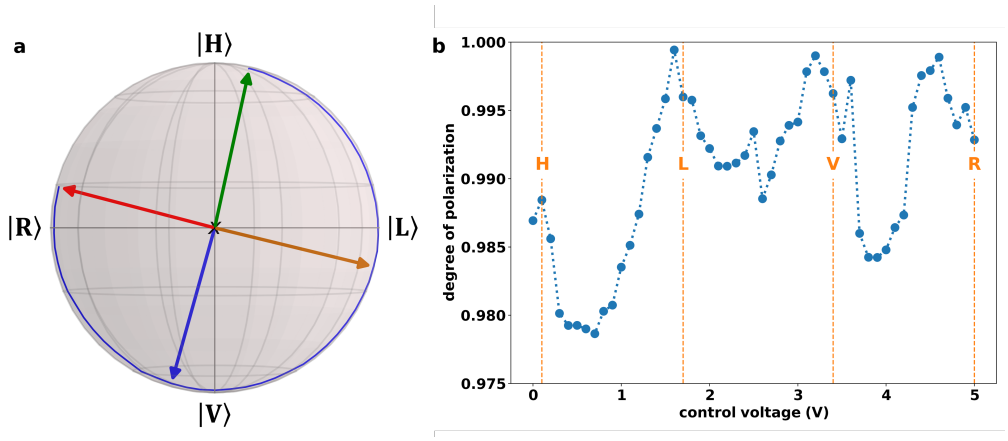


Figure 6.5: **Characterization of the polarization encoding.** The analog control voltage applied to the Pockels Cell driver is scanned, and the rotation of the output polarization is measured with a polarimeter. a) Since the polarization state at 0 V does not correspond to $|H\rangle$, the optimal set of states for the verification protocol is identified by minimizing the inner product between state vectors of conjugate bases while maximizing it for state vectors of the same basis. b) The average degree of polarization of the set is $> 99\%$.

any polarization drift on long time scales, the routing is repeated periodically (see Sec. 6.1.1.4).

The polarization corrected light is split equally into two paths, the first is projected onto the rectilinear basis using a PBS and the second onto the circular basis by passing a QWP at 45° before being separated at another PBS. All four beams are coupled to SMF using 8 mm focal length collimators. The full projection setup including the SMF coupling exhibits an average transmission of $\bar{\eta}_{\text{proj}} = 72(3)\,\%$. Analyzing the transmissions individually, a slight imbalance between the rectilinear and circular basis is noted, explicitly $\eta^H_{\text{proj}} = 76(3)\,\%$, $\eta^V_{\text{proj}} = 76(3)\,\%$, $\eta^L_{\text{proj}} = 68(2)\,\%$, $\eta^R_{\text{proj}} = 68(2)\,\%$. Separately measuring the BS's splitting ratio and fiber-coupling efficiency shows that both imperfections contribute to the imbalance. First, the fiber coupling for the rectilinear basis is $\approx 90\,\%$ but only $\approx 86\,\%$ for the circular basis. The decrease in coupling efficiency is attributed to the longer optical path of the circular projection, approaching the Rayleigh length of the collimators. Second, the BS shows a small imbalance in the splitting ratio, reflecting $\approx 2\,\%$ more than transmitting at the signal wavelength.

Since a considerable mismatch in basis detection efficiency can be exploited [154], the more efficient SNSPDs were paired with the fibers connecting the circular basis, thus decreasing the overall imbalance in click probability to $< 2\%$. The average SNSPD detection efficiency is $\bar{\eta}_{\text{det}} = 86(2)\,\%$ and the mean probability of registering a click is $\bar{\eta}_{\text{proj}} \cdot \bar{\eta}_{\text{det}} = 62(4)\,\%$.

The average dark count rate per detector is 40(10) Hz and the probability of

any detector producing a dark count within the gating window is

$$\tilde{p}_{\mathrm{dc}} = \Delta t_{\mathrm{gate}} \sum_i DCR_i = (4.8 \pm 1.2) \cdot 10^{-6} \,, \tag{6.1-2}$$

where $i = \{H, L, V, R\}$ denotes the detector.

6.1.1.4 Synchronization

Synchronization of all electronic components in the setup is performed by an FPGA (*Digilent Genesys 2*) that is sourced with the master clock from the pulsed laser. A 80 MHz low-pass filter and an amplifier pre-process the laser clock signal to only transmit the base frequency and improve the phase-locking stability of the FPGA. In this setting, a timing jitter of ≈ 30 ps is measured between the unprocessed laser clock and the re-generated clock from the FPGA. By integer division, three reduced clock rates are synthesized to match the individual requirements of the instruments.

First, a 500 kHz signal triggers the electronic pulse generator (EPG, *Polytec EPG-210B*) to produce a pulse with $FWHM = 100(10)$ ps width and ≈ 5.5 V amplitude that drives the fiber-based pulse picker. Feeding a bias voltage to the fiber EOM ensures that the light is suppressed up to its maximum extinction ratio whenever no RF pulse is applied. To overlap the electrical and laser pulse at the fiber EOM, the clock's phase can be adjusted in increments of ≈ 10 ps at the FPGA.

Second, a 20 MHz sync signal is generated by the FPGA, split into two and fed to both Pockels Cell controller (PCC) (*EKSMA pMaster 4.2*) which in turn feature an electronic circuit to phase-lock to the external clock. The PCCs generate HV pulses at 500 kHz that are overlapped with the QD light using an internal adjustable delay. The picking Pockels Cell receives only 2.6 kV pulses to flip the linear polarization by 90° but the state encoding Pockels Cell requires an additional low-voltage (0 V to 5 V) analog control signal tuning the HV amplitude (thus, the retardance).

This control signal is provided by an arbitrary waveform generator (AWG, *Siglent SDG2000X*) that is phase-locked to the third reference clock output by the FPGA with 10 MHz. The AWG samples at a rate of 1 MHz (each sample lasts for 1 µs) from a 500 kSa long sequence stored on the internal drive. The encoding sequence was generated pseudo-randomly (using the *random* module from the Python package *numpy*), picking from a set of four different voltage levels corresponding to the four polarization states. For more details on how the set voltages were selected, see Sec. 6.1.1.2. Note that each voltage plateau consists of two samples, thus matching the protocol rate of 500 kHz.

The second (internally phase-locked) channel of the AWG samples from another sequence, referred to as the trigger sequence, consisting of alternating low and high levels (0 V and 2 V) effectively generating rising edges at 500 kHz. This channel output is used to synchronize the two parties and is converted to an optical signal

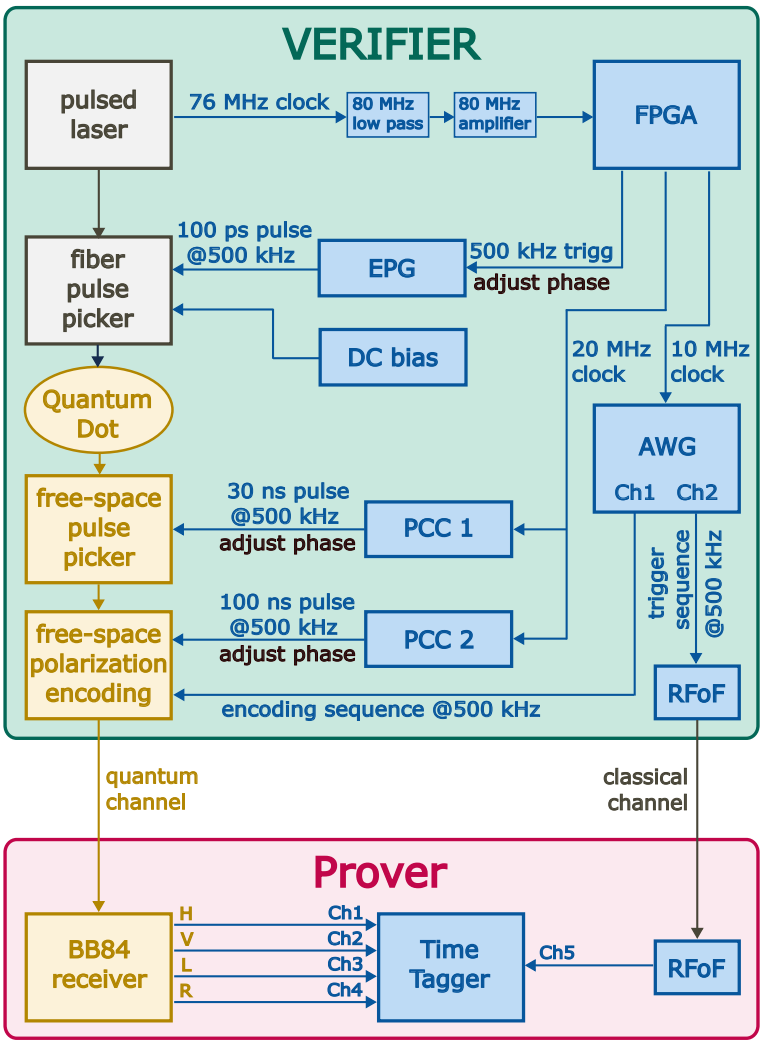


Figure 6.6: **Schematic layout of the synchronization setup.** The pulsed laser provides the master clock at 75.95 MHz to a field programmable gate array (FPGA) which performs all phase-locked integer clock division. All frequencies in the figure are rounded for simplicity. For more details, see Sec. 6.1.1.4. EPG: electronic pulse generator, PCC: Pockels Cell controller, AWG: arbitrary waveform generator, RFoF: electronic-to-optical synchronization signal converter.

using a commercial RFoF (*RFoptic*) transmitter before being sent over the classical fiber link. Using the RFoF receiver, P converts the trigger back to an electrical pulse, feeds it to the fifth channel of the time tagging device (*Swabian Instruments TimeTagger Ultra*), and can then associate each detected polarization state with its signal bin. The additional jitter introduced by the RFoF setup is small ($< 10\,\mathrm{ps}$). Both the encoding and trigger sequence feature a small marker section consisting

of ten samples at the beginning to facilitate the signal bin assignment during postprocessing.

The above synchronization scheme covers all communication between the parties during the quantum verification protocol. To compensate for the arbitrary polarization rotation introduced by the fiber link, the parties perform a separate routine that corrects the rotation in both conjugate bases. Even though these steps require cooperation, they do not impact the security of the verification. If a party acts dishonestly at that stage, it cannot reach the predefined error threshold and cannot start the protocol, resulting in denial-of-service behavior. Note that the error threshold is directly derived from the security analysis and is not set by the experimental calibration.

In our implementation, one computer in V's lab and another computer in the P's lab communicate over the university's Local Access Network (LAN) using a custom protocol based on the Transmission Control Protocol/Internet Protocol (TCP/IP). V prepares only one polarization state by changing the encoding sequence to a constant signal at the corresponding voltage and asks P to optimize the visibility of the state's basis. Once the target value has been reached, V sends another polarization state. They iterate this process until the visibility requirements for all four states have been met. For the physical implementation of the correction, see Sec. 6.1.1.3.

Once the calibration is complete, V announces to start sending random states and the quantum verification protocol begins. In pre-agreed intervals (every ≈ 3 hours in our implementation), the parties interrupt the protocol to compensate potential polarization drifts in the quantum channel.

6.1.2 Problem instance

As mentioned in the project introduction, the protocol allows a Prover to selectively disclose part of the solution to a mathematically hard problem to an untrusted Verifier who wishes to confirm their knowledge of the solution. Here, *hard* typically refers to a computational complexity, meaning that the solving time increases superpolynomially with the problem's size. Since all problems from the NP-complete complexity class may be efficiently converted to a 3-satisfiabilty (3-SAT) problem [396, 397], we adopt the natural formalism of 3-SAT, which is the conjunctive normal form (CNF)[1]. However, to conduct a security analysis and set security parameters, we must quantify the maximum number of clauses a non-satisfying solution assignment may answer correctly, and for a random 3-SAT problem without any known structure, one cannot easily find such a threshold. Therefore, we apply the PCP theorem [398, 399] stating that there exists a polynomial-time reduction mapping from any 3-SAT instance ϕ to a constraint system ψ. More specifically, using the gap amplification and reduction to MAX-E3SAT(5) described in [400–402], we obtain a MAX-E3SAT(5) instance ψ that is probabilistically checkable. This

[1] An example of a CNF is given in the reprint in Sec. 6.2

is a 3-SAT instance in which each variable appears in exactly 5 clauses. For ψ it follows

$$\begin{aligned} &\text{Correctness (C):} \quad \phi \in \text{SAT} \Rightarrow \text{OPT}(\psi) = 1, \\ &\text{Soundness (S):} \quad \phi \notin \text{SAT} \Rightarrow \text{OPT}(\psi) \leq 1 - \frac{\frac{1}{8} - \delta}{8600}, \end{aligned} \tag{6.1-3}$$

where $0 < \delta < \frac{1}{8}$ and $\text{OPT}(\psi)$ is the maximum fraction of clauses of ψ that any assignment can satisfy, i.e. $\text{OPT}(\psi) = \max_{\tilde{s}} f_{\tilde{s}}(\psi)$ where $f_{\tilde{s}}(\psi)$ denotes the fraction of clauses satisfied by assignment $\tilde{s}$. In essence, correctness is the probability that V accepts a correct proof, while soundness upper-bounds the probability that V accepts an incorrect proof. Including the impact of the bounded quantum memory, we derive a bound for the soundness related to the 3-SAT problem (Eq. 6.2-2). The full chain of reductions from which one arrives at the gap-problem from Eq. (6.1-3) and how the soundness follows is described in Appendix B.

6.2 Manuscript reprint

Convincing a party that one possesses the solution to a hard mathematical problem without disclosing the full solution is a challenge of broad cryptographic relevance. Typically, existing schemes assume either the existence of one-way functions or a stringent monitoring of the parties' spacetime coordinates, which is difficult to enforce in practice. Here, we propose the notion of limited-knowledge proofs, an information-theoretically secure quantum alternative that lets a Prover reveal only a small portion of a secret while still convincing a Verifier of its validity. Secure against quantum adversaries holding a perfect memory of a hundred qubits, we allow for multiple verification rounds without compromising the initial problem's hardness. From a practical standpoint, our scheme can be implemented with highly attenuated laser states or down-conversion sources. However, we show that employing a SPS emitting directly in the telecom C-band can boost its performance by several orders of magnitude, even at moderate channel loss. By coupling quantum resources with an investigation of solution leakage in hard computational problems, our results offer a quantum primitive with applications to authentication in the context of QKD.

6.2.1 Introduction

Suppose you hold the solution to a notoriously hard mathematical problem and wish to convince another party of this fact without leaking the solution or any information that would compromise the problem's hardness. Beyond personal motivations for maintaining such secrecy, this challenge presents widespread cryptographic applications, from proving one's identity to a potentially untrusted party without revealing one's sensitive information (eg: when authenticating oneself to a potentially spoofed webpage) to proving one's eligibility for an online vote without revealing one's identity, or convincing an auctioneer that a bid is high enough without revealing its actual amount.

To achieve these, *selective disclosure* schemes can be used, allowing a user to reveal only parts of a verifiable credential to an authority, without exposing their entire set of credentials [403–405]. Typically, security in this setting may be achieved through the primitive of *bit commitment*, in which a Prover locks their secret inside a safe such that it is hidden from any other party and cannot be modified once the safe has been locked (even by the Prover themselves). At a later time, the Prover may publicly disclose the key to a Verifier wishing to check the commitment. For our scenario of interest, imagine a longer secret split across several safes, each with its own key. During verification, the Prover reveals only the keys to some randomly chosen safes, so that a Verifier can check their contents without learning the remnant of the secret contained in the other, locked safes. Using classical resources, commitment schemes rely either on computational assumptions such as the existence of one-way functions [388–390] or on no-signaling constraints that require precise monitoring of multiple Provers' spacetime coordinates [391–394]. While the latter assumption offers security independent of an attacker's

computational power (labeled *information-theoretic (i.t.)* security), enforcing it in practice implies a trusted Global Positioning System (GPS), which presents vulnerabilities to spoofing attacks and time-stamp manipulation [406].

In this work, we introduce the notion of *limited-knowledge proofs*, a partial disclosure protocol that relies on an *i.t.*-secure quantum alternative to the aforementioned solutions. Based on a commitment scheme secure against attackers holding a quantum memory of bounded size [106, 147, 148, 407, 408], a Prover can convince a Verifier that they hold the solution to a hard problem by revealing only a constant fraction of the problem size at each verification round. Our *i.t.*-secure scheme essentially hides the solution of an NP-complete problem (to which all other problems in the NP class can be efficiently reduced [396, 397]) and allows for multiple verification rounds while maintaining the inherent hardness of the original problem. This can be applied to scenarios in which a very small part of a secret may be disclosed for verification while still allowing for a finite number of reuses of the secret, such as for the initial authentication step of QKD.

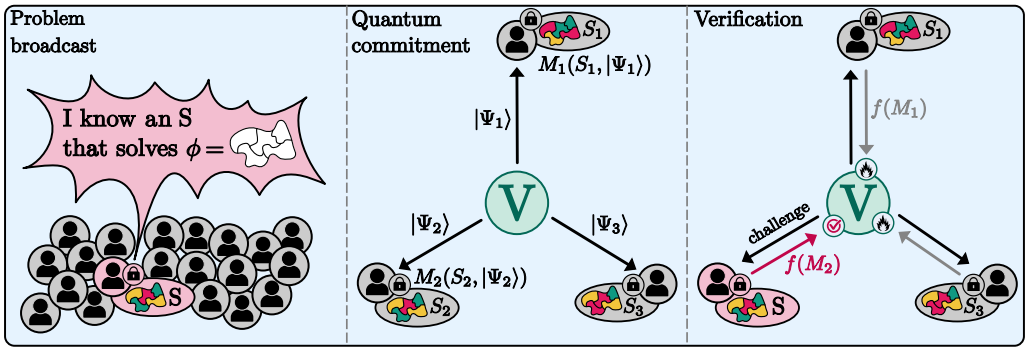


Figure 6.7: **Limited-knowledge proof using quantum commitments.** During the *problem broadcast*, an honest Prover announces that they have a satisfying assignment S to an NP-complete problem ϕ. In the *quantum commitment* phase, a Verifier (V) sends a string of quantum states $|\Psi_i\rangle$ to each party P_i claiming to be the honest Prover. In response, the honest Prover measures the quantum states in a basis dictated by their solution S_i. After a pre-agreed time, V initiates the *verification* by sending a series of challenges to each party used to identify the legitimate Prover. Based on their classical answer $M_i\,(S_i, |\Psi_i\rangle)$, V may check the commitment and verify that S_i satisfies ϕ.

By operating a QD SPS emitting directly in the telecom C-band [80], we implement the scheme over a deployed optical fiber linking two buildings in central Vienna. While PDSs such as attenuated laser states may be used, photonic states emitted by QDs provide a number of benefits in mistrustful quantum cryptography [45, 121], combining low multiphoton contribution with the possibility of erasing coherence in the photon-number basis [45, 46, 215]. The first feature provides an advantage in terms of distance performance, while the second circumvents the need for active phase randomization. In our scheme, we show that SPSs

provide a scaling advantage for all practical values of channel loss (not just those corresponding to large distances) and extend the protocol's distance limit beyond what is achievable with PDSs. By an appropriate choice of classical post-processing, our experiment closes practical loopholes exploiting the back-reporting of detection time-tags present in previous demonstrations of mistrustful quantum cryptography [121, 154], including in the noisy-storage model [352, 409].

6.2.2 Limited-knowledge proofs

6.2.2.1 Problem choice

Our protocol uses logical formulas in CNF to encode the problem whose solution should be hidden from the Verifier. We denote a 3-CNF formula as ϕ, consisting of a conjunction of clauses which each contain three variables from the set $\{y_l\}_{l=1...L}$. For example, ϕ may take the form

$$\phi = (y_3 \vee y_{17} \vee \neg y_9) \wedge \ldots \wedge (\neg y_1 \vee y_9 \vee \neg y_{14}) \,. \tag{6.2-1}$$

This is the natural language for a 3-SAT problem that asks whether there exists a variable assignment $s \in \{true, false\}^L$ to a given 3-CNF formula, such that the formula is satisfied. Each clause is satisfied if at least one of its variables (or its negation) evaluates to *true*, while ϕ is satisfied if all clauses evaluate to *true*.

3-SAT is proven to be NP-complete, which implies that any other problem in NP can be efficiently converted to a 3-SAT instance and a 3-CNF formula [396, 397]. We call an instance *regular* if every variable appears in the same number of clauses, and *balanced* if each variable appears equally often positively and negatively.

6.2.2.2 Protocol

As depicted in Fig. 6.7, our limited-knowledge proof scheme consists of three main phases: *problem broadcast*, *quantum commitment*, and *verification*.

In *problem broadcast*, a Prover P publicly announces a 3-SAT instance ϕ encoded as a 3-CNF formula. P is challenged to convince a Verifier V that $\phi \in SAT$, i.e. ϕ is satisfiable. As outlined in the Sec. 6.1.2, V applies a reduction to arrive at a regular and balanced MAX-E3SAT(5) instance ψ that is probabilistically checkable [398, 399]. This is a 3-SAT instance in which each variable appears in exactly 5 clauses, with a total of L variables and M clauses, and satisfies the following property: if $\phi \in SAT$, then all clauses of ψ are satisfiable, if $\phi \notin SAT$, then at most $1 - \tilde{\delta}$ clauses are satisfiable, for $\tilde{\delta} > 0$.

In *quantum commitment*, P receives a random BB84-type quantum state $\otimes_{n=1}^{N} |X_n\rangle_{\Theta_n}$ of length N from V for each variable $l \in \{1, ..., L\}$. They measure each state in a basis $\tilde{\Theta} \in_R \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowleft\}^N$ and record the timings $\tilde{t}_{true} = \{\tilde{t}_j \,|\tilde{\Theta}_j = \leftrightarrow\updownarrow\}$ and $\tilde{t}_{false} = \{\tilde{t}_j \,|\tilde{\Theta}_j = \circlearrowleft\circlearrowleft\}$ (the tilde always marks variables resulting from P's actions). The outcome string of this measurement is labeled $\tilde{X}$. To complete the commitment, P sends the concatenation of all labels $\tilde{t}_{true}$ if the solution assignment

$s_l = true$ or $\tilde{t}_{false}$ if $s_l = false$ back to V. In this way, $\tilde{t}_{true}$ and $\tilde{t}_{false}$ are used to commit, while $\tilde{X}$ is used to verify the commitment.

In *verification*, V asks P to reveal the variables for k arbitrarily selected clauses, corresponding to a subset of variables $\mathcal{L} \subseteq \{1, ..., L\}$ of size $|\mathcal{L}| \approx 3k$. P must announce the solution bit $s_{l'} = true/false$ for each $l' \in \mathcal{L}$. V then checks whether all challenged clauses are satisfied by the reported solution bits. Additionally, to ensure that P has honestly committed, V identifies the set of timings within $\tilde{t}_{s_{l'}}$ for which P's measurement basis matches their own encoding basis, yielding the set $t_{s_{l'}}$. For those states, P's measurement outcomes must correspond to V's encoded bits—yielding the set $\{\tilde{X}_{s_{l'}}\}$—except for errors that can be corrected with the syndromes computed and sent by V. P and V concatenate the bitstrings $\{\tilde{X}_{s_{l'}}\}$ and $\{X_{s_{l'}}\}$ into the global strings $\tilde{X}_{\mathrm{v}}$ and X_{v}, respectively. P sends $\tilde{X}_{\mathrm{v}}$ to V who accepts the proof if $\tilde{X}_{\mathrm{v}} = X_{\mathrm{v}}$, and rejects otherwise.

6.2.2.3 Correctness, soundness and limited knowledge

The protocol's *correctness*, i.e. the probability that a correct proof is accepted by V, should be lower-bounded by $1 - \varepsilon_{\mathrm{c}}$, where ε_{c} is a small deviation parameter that accounts for potential failures after the PCP reduction or during the error correction if that step is required.

Against a dishonest Prover P^*, the overall *soundness* ε_{s} denotes the probability that an incorrect proof is accepted by V, and can be expressed in terms of two security bounds $\varepsilon_{\mathrm{s,sat}} + \varepsilon_{\mathrm{s,q}} = \varepsilon_{\mathrm{s}}$. If the instance ϕ is not satisfiable, we know from the PCP theorem that no more than a fraction $1 - \tilde{\delta}$ of clauses can be satisfied in the corresponding MAX-E3SAT(5) instance ψ (see Sec. 6.1.2). Using this result, our first term $\varepsilon_{\mathrm{s,sat}}$ relates to the security of the SAT instance, essentially upper-bounding the probability that V challenges only satisfied or stored clauses in k rounds:

$$\left(1 - \tilde{\delta} + 3 \cdot \frac{y_{\mathrm{store}}}{L}\right)^k < \varepsilon_{\mathrm{s,sat}}. \tag{6.2-2}$$

Here, y_{store} denotes the number of variables stored by P^*, and accounts for attacks in which an initially unsatisfied clause is turned into a satisfied one by storing the quantum states associated with these variables (delaying measurements in such a way allows P^* to change their committed values in order to satisfy the clause). However, since the adversary's quantum storage is assumed to be limited in size, y_{store} can be upper-bounded. As detailed later, although increasing k improves the achievable $\varepsilon_{\mathrm{s,sat}}$, it also increases the leakage of solution variables, thus affecting the number of secure reuses of the SAT instance (for a summary of the parameters, see Tab. 6.1). The derivation of the second term $\varepsilon_{\mathrm{s,q}}$ relating to the quantum commitment scheme is detailed in the Sec. 6.3.2.

Against a dishonest Verifier V^*, our *limited knowledge* property fixes a maximum number $G_{\max}$ of verification rounds such that

(i) no more than $B = O(1)$ bits of information are leaked across all rounds,

(ii) no more than a fraction $q_{\max}$ of the problem's clauses are challenged.

With this constant leakage, NP-hardness is preserved: if knowing up to B solution bits let one always decide SAT more easily, then we could try all $\sum_{j\leq B}\binom{L}{j}2^j = poly(n)$ possibilities for which j bits are leaked and their values, thereby simulating the advantage in polynomial time *without* actual leakage, which would contradict NP-hardness.

6.2.3 Experimental demonstration

6.2.3.1 Setup

Following Fig. 6.8, we realize the protocol between two university buildings connected by a 680 m-long optical fiber link with a transmission of $\eta_{\text{ch}} = 83(3)\,\%$. V generates single photons using a QD embedded in a circular Bragg grating cavity emitting at 1550.9 nm [80], pumped by a 1.4 nm blue-shifted pulsed laser. The longitudinal acoustic phonon-assisted (LA) excitation [235] produces phase-randomized single-photon states [1, 45, 219] with high polarized fiber-coupled emission efficiency $\eta_{\text{sys}} = 8.5(17)\,\%$ and low multiphoton contribution, extracted from the autocorrelation function at zero delay $g^{(2)}(0) = 1.86(11)\,\%$. Furthermore, the QD emission is easily spectrally filtered from pump light reflections, and the scheme is robust against laser instabilities [1].

Based on the random input strings X and Θ, V prepares each photon in a different polarization state using a free-space electro-optic modulator (EOM), and sends them to P via the deployed quantum channel. To compensate for photon loss and to increase the protocol's security, each variable is encoded not once but N times (using different random polarization states). P projects the encoded states passively onto the rectilinear or circular basis and detects the photons using SNSPDs. As a countermeasure to attacks exploiting the back-reporting of invalid timetags [154], any detection round yielding more than one click is randomly mapped onto one of the four detector outcomes.

In our implementation, we detect $\approx 4.1 \cdot 10^9$ photons out of $N_{\text{enc}} = N \cdot L \approx 1.46 \cdot 10^{11}$ total encoding cycles, signifying an end-to-end efficiency of $2.813(2)\,\%$. Evaluating the individual error rates per encoding setting, we find an average QBER of $0.48(24)\,\%$. The complete set of experimental parameters is presented in Tab. 6.1.

6.2.3.2 Practical performance

As a reminder, our scheme must account for two types of cheating strategies from P^*: they may convince V that an instance ϕ is satisfiable even though $\phi \notin$ SAT either by presenting some compelling witness, or by directly attacking the quantum commitment. The success probability of the first is controlled by the number of challenges k issued by V: for sufficiently large k, an unsatisfiable instance will, with high probability, contain at least one challenged clause that cannot be answered consistently. The security of the quantum commitment, on the other hand, is

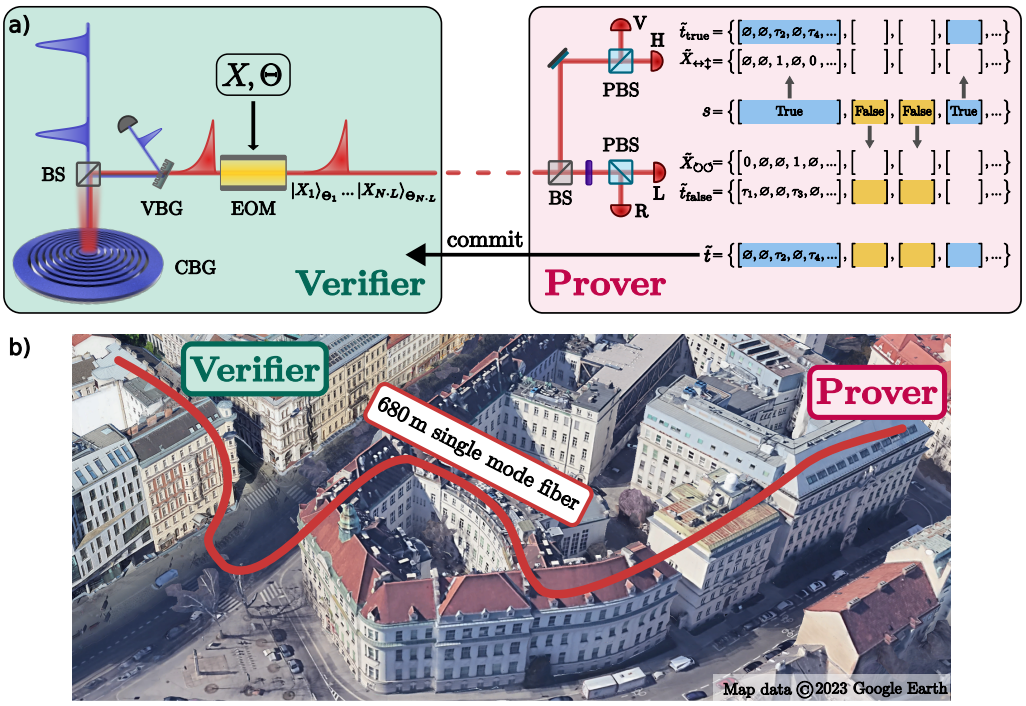


Figure 6.8: **Experimental implementation of the quantum commitment phase.** V generates single photons by exciting a quantum dot sitting in a circular Bragg grating (CBG) with a blue-detuned pulsed laser. After suppressing the laser with a volume Bragg grating (VBG), the classical random strings X and Θ are bit-wise encoded onto the single photons' polarization using a free-space electro-optic modulator (EOM). Per solution variable s_l, N states are encoded, such that in total $N_{\text{enc}} = N \cdot L$ states are sent via the quantum channel, where L is the number of variables of the problem ϕ. P passively projects all states and records the outcomes and arrival times using superconducting nanowires single-photon detectors and a time-tagging device. Signal bins for which no photon was detected (due to imperfect single-photon generation or channel loss) remain unassigned. Then, P commits to each solution variable s_l, by forming the timing subset $\tilde{t}_{true} = \{\tilde{t}_j | \tilde{\Theta}_j = \leftrightarrow\updownarrow\}$ ($\tilde{t}_{false} = \{\tilde{t}_j \,|\tilde{\Theta}_j = \circlearrowleft\circlearrowleft\}$) if $s_l = false$ ($s_l = true$) and back-reporting the concatenated list of arrival times for all challenged variables. b) Satellite image showing the two university buildings linked by the quantum channel consisting of a 680 m long optical fiber.

governed by the number of states per variable N.

Additionally, the protocol must ensure that P does not reveal too much information to V^*. Suppose an instance ψ (i.e., the reduced form of ϕ) is verified G times. In each run, V^* challenges k clauses, so across all runs, at most kG distinct clauses are challenged, corresponding to revealing a total of B distinct variables. As shown in the Sec. 6.3.2, as long as the total leakage satisfies $B = O(1)$, the

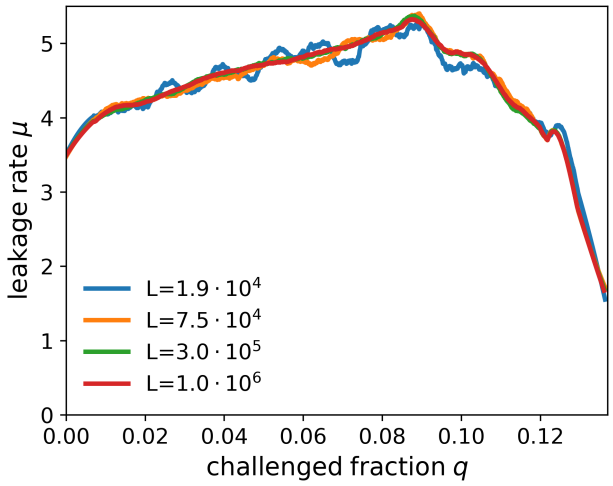


Figure 6.9: **Maximum leakage rate for V^*.** For MAX-E3SAT(5) problems randomly generated using our *cnfgen* algorithm, the optimal leakage rate μ that V^* can achieve is plotted for various problems with L variables against the challenged fraction $q = C/M$ (C is the total number of challenged clauses, $M = 5L/3$ is the number of clauses of the problem instance). The leakage rate is calculated as the derivative $\mu = \mathrm{d}B/\mathrm{d}C$ (B is the total number of leaked variables including indirect leakage). Since μ exhibits strong statistical fluctuations per challenge, a smoothing Savitzky–Golay filter was applied to improve clarity. The leakage rate is upper-bounded for all q, and $\mu(q)$ behaves very similarly for all L, indicating the independence of the problem size. The numeric simulations are performed in the limit of $k = 1$ (that is, one challenge at a time), which overestimates the power of V^*, since they receive new information after every challenge before deciding which clause to challenge next. The Verifier strategy is outlined in detail in the Sec. 6.3.4. When $q \approx 14\,\%$ of all clauses have been challenged, the full solution has been revealed. The absence of a spike in the leakage rate μ indicates that there is no particular challenge that would make the problem trivial.

underlying gap-problem from Eq. (6.1-3) remains NP-hard.

In practice, condition (i) from the *"Correctness, soundness and limited knowledge"* section motivates an instance-independent leakage threshold B_0: V^* may repeatedly challenge the same problem instance as long as the cumulative information leaked about its satisfying assignment does not exceed B_0 bits. To study how leakage accumulates, we simulate a strategy that is designed to maximize V^*'s information gain. Given a balanced 3-5CNF formula generated with our *cnfgen* algorithm (see Sec. 6.3.4), where every variable appears in five distinct clauses, V^* repeatedly samples clauses and selects those that are expected to reveal the most information, either directly or indirectly. Indirect leakage occurs whenever two variables in a clause evaluate to *false*: the clause structure then forces the third variable to be *true* even if it is never explicitly queried. On the one hand, this

simulation allows us to estimate the number of times an instance can be verified before the cumulative leakage reaches the global budget B_0. On the other hand, in Fig. 6.9, we quantify the average information revealed per challenge by tracking the leakage rate $\mu = \mathrm{d}B/\mathrm{d}C$ (C is the total number of challenges) to define $q_{\max}$ for condition (*ii*). Analyzing μ over multiple challenges, we find that the number of newly leaked variables is almost always larger than the number of variables directly disclosed by P, i.e. $\mu > |\mathcal{L}|/k \approx 3$.

However, μ does *not* increase monotonically as one might expect. Instead, once a large part of the problem is known, V^* must start to challenge clauses that are partially known but still not trivially solvable, and in doing so, both direct leakage and the chance to discover a variable indirectly decline steadily. Importantly, our computational experiments suggest that the performance of this strategy is independent of the problem size L.

Combining these results, we conclude that an instance ψ may be reused safely as long as less than B_0 bits have been leaked and less than $q_{\max}$ of all clauses have been challenged. Using $q_{\max}$, we then calculate how many times ψ can be used securely as $G = 5/3k \cdot L \cdot q_{\max}$, indicating that G increases at least linearly with the size of the problem instance (k is also a function of L, see Fig. 6.15). Recalling that the security against P^* scales with N, it is clear that the reusability and quantum security are in direct competition for more encoding states ($N_{\mathrm{enc}} = N \cdot L$), and any implementation must balance them while maintaining practical acquisition times.

Based on the value of N_{enc} used in our experiment, Fig. 6.10a illustrates this trade-off. The adversary's maximal quantum storage size $S_{\max}(N)$ is calculated using our security analysis based on the experimental parameters (see Tab. 6.1). We find a lower limit if N is too small to secure the protocol against even a memory of a single qubit ($S_{\max} < 1$), but on the other side, if N is chosen too large, the problem size L must necessarily be reduced so much that even a single use could reveal the full solution to V^* ($G < 1$). For a fixed N_{enc}, we may express the number of secure uses as $G(N)$, and rearrange to explicitly formulate the security trade-off as $S_{\max}(N(G))$: the protocol is secure for any combination of S and G that satisfies $S < S_{\max}(N(G))$.

To visualize how the protocol performs over larger distances, we chose a pair of security guarantees S and G (indicated by the star in Fig. 6.10a) achieving five secure uses, while securing the bit commitment against a quantum memory with fidelity arbitrarily close to 1 and a size of 117 qubits, which is orders of magnitude beyond state-of-the-art [410–412]. Note that even if the adversary acquires many copies of such a quantum memory, one can increase the number of quantum states per variable sent during the protocol execution.

6.2.3.3 Scaling with distance

To study the impact of communication distance on our scheme, we recall that our quantum dot source generates phase-randomized states [1, 45, 219], so P^* may exploit either one of two strategies to decommit on a challenged variable: storing

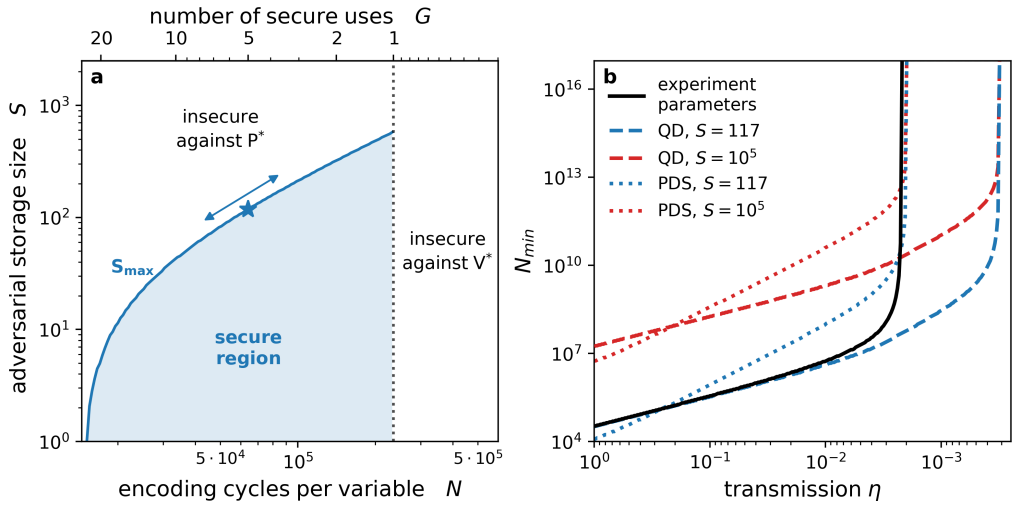


Figure 6.10: **Limited-knowledge performance.** ***a) Quantum security bound.*** Using the experimental parameters $\eta = 51\,\%$, $[p_0, p_1, p_\mathrm{m}] = [0.946, 0.054, 3.1 \cdot 10^{-5}]$, detector dark count probability $\tilde{p}_\mathrm{dc} = 5 \cdot 10^{-6}$ and QBER of $0.5\,\%$, the security against an dishonest Prover with a perfect quantum memory ($r = 1$) of total storage size S is calculated as function of the encoding cycles per variable N. For a fixed $N_\mathrm{enc} = N \cdot L$ (here, $N_\mathrm{enc} = 1.46 \cdot 10^{11}$), the problem size L directly depends on N. The curve S_max indicates the storage that an the adversary would need store a full variable, $y_\mathrm{store} = 1$, and the shaded region below is secure. Under these conditions, the number of challenges k required to achieve the soundness $\varepsilon_\mathrm{s,sat}$ is determined (see Sec. 6.3.2) and the number of secure uses is inferred as $G = 5Lq_\mathrm{max}/3k$, where $q_\mathrm{max} = 12\,\%$. Please refer to Tab. 6.1 for the full set of protocol parameters. ***b) Quantum security over transmission.*** Maintaining a given security for increasing channel loss requires a larger N_min. Note that the transmission η not only includes channel loss but also imperfect detection. All lines indicate the lower bound on N, the area above is secure. The solid black line shows the bound using the parameters from the experiment and security bounds as indicated by the marker in **a)**. To compare the impact of the different parameters, N_min is displayed for the photon-number optimized PDS and QD statistics at $QBER = 0.5\%$ and $\tilde{p}_\mathrm{dc} = 10^{-7}$ for two storage sizes. For more parameter variations, see Sec. 6.3.3.

single-photon states in their quantum memory and only measuring them after being challenged, or exploiting the multiphoton states of an imperfect source to project one photon in each basis. Similar to other quantum protocols, there is a maximum loss that can be tolerated before perfect cheating is possible solely based on the source's multiphoton component p_m. Thus, reaching further requires reducing p_m, even at the cost of also lowering p_1. For PDSs, photon-number optimization is routinely performed by attenuating it in a trusted manner, and the same approach has also been shown to benefit an imperfect SPS at high channel loss [1, 42].

Including this optimization, Fig. 6.10b compares how the photon statistics

of our QD emitter and those of a PDS perform with increasing channel loss. In the low to moderate loss regime, the multiphoton component of the QD is small enough to not impact security, and the quantum memory poses the main threat. Hence, the minimum number of encoded states $N_{\min}$ must increase only linearly with loss ($N_{\min} \propto 1/\eta$) to ensure that the expected number of detected photons remains large enough to overflow the memory such that any attempt to decommit would be noticed. At some distance, however, p_{m} becomes too large even when using a QD source and the emission must be pre-attenuated, resulting in a quadratic increase $N_{\min} \propto 1/\eta^2$. Importantly, this photon-number optimization is necessary for PDSs even at low loss (causing a quadratic scaling throughout), and we identify a performance advantage of our SPS over PDSs for all but for very low-loss and low-QBER implementations. For a system efficiency $\eta_{\mathrm{sys}} \geq 16\%$, readily over-achieved by QD sources emitting in the near-infrared [55, 56], single-photon emitters outperform PDSs in all parameter regimes (see Sec. 6.3.3).

Finally, the impact of detector dark counts, causing the protocol to abort at high loss, is significantly mitigated by the larger single-to-multi photon ratio provided by QDs, pushing the maximum loss almost by an order of magnitude. Independent of the source type, securing the protocol against larger storage sizes S requires approximately linearly more encoding cycles per variable, $N_{\min} \propto S$. While low error rates improve the protocol's performance by reducing $N_{\min}$ (see Sec. 6.3.3), we note that our protocol is remarkably robust to noise and can tolerate a QBER of up to 7.2%.

6.2.4 Discussion

In this work, we have introduced a limited-knowledge proof that embeds one of the most well-studied NP-complete problems into an information-theoretically secure quantum commitment scheme. By analyzing its practical security and demonstrating its robustness to noise and loss-dependent attacks, we have shown that the amount of disclosed information can be bounded by a constant, while the underlying quantum commitment is secure against adversaries holding a perfect memory of more than a hundred qubits.

To assess the scalability of our approach, we note that the $QBER \approx 0.5\%$ achieved here is remarkably low for active polarization encoding, especially for an intra-city fiber link. When implementing the protocol on a larger scale, replacing the free-space EOM with a waveguide-based alternative can significantly enhance the encoding rate by a factor of $10^3 - 10^4$ [80, 413, 414]. Additionally, the protocol's high noise tolerance of $QBER \leq 7.2\,\%$ easily accommodates the elevated QBER of integrated schemes, making it suitable for combining it with QKD, even when using commercial equipment. While the limited-knowledge proof is secure for both Poisson-distributed and SPSs, a low multiphoton component significantly reduces the number of required states, especially in the high-loss regime. Further, in the presence of detector dark counts, they enable larger communication distances. These findings demonstrate that highly pure SPSs are particularly well-

suited for mistrustful quantum primitives [121], as simulated in [45] and recently experimentally demonstrated for another quantum-cryptographic task [46].

On the computational side, we have initiated a study of how information leakage propagates when random 3-5CNF formulas are repeatedly queried. To tighten our bound and thus increase the number of secure reuses G, a more comprehensive analysis of MAX-E3SAT(5) instances obtained from a *full* reduction starting from arbitrary 3-SAT formulas, together with a systematic exploration of all possible leakage patterns for such instances, should be investigated. Furthermore, identifying alternative reductions to SAT instances that yield a larger soundness gap than MAX-3SAT(5) problems should significantly enhance the performance. The larger the soundness gap, the fewer the number of clauses that a Verifier must challenge, thus allowing for a larger number of secure instance reuses. Finally, we have numerically shown that, for the dishonest Verifier strategy employed in our protocol, the total information leaked can be bounded. Since NP-complete problems are only known to be hard in the worst case, a natural next step is to adapt the construction to settings based on average-case hard problems, making our limited knowledge protocol directly applicable to authentication schemes for example.

6.3 Complementary material

6.3.1 Experimental data

To commit to a solution assignment and later address the challenges from V, P has to communicate the timing, basis, and outcome for a subset of the detected states. Therefore, the honest Prover must analyze the stream of detection events based on the shared reference clock. They accept only detections within the gating window Δt_{gate} with respect to the reference; all other events are discarded. For all signal bins with a single click, basis and measurement outcome are straightforwardly assigned. If two or more clicks occurred within a single signal bin, a random basis and outcome are assigned to the bin. If there is no detection within a signal bin, it is considered as photon loss. All signal bins that have been assigned a basis and measurement outcome are considered valid. The number of valid signal bins per second is referred to as signal rate ν_{sig}. If not enough detection events are recorded within a predefined time interval, P aborts the protocol. Else, P selects the signal bins used for the commitment according to the solution as described in Box 1.

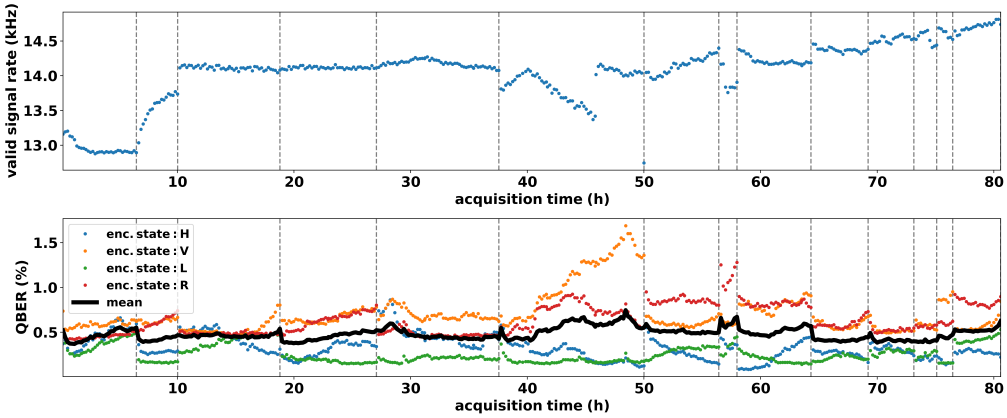


Figure 6.11: **Analysis of the full detected data.** The data used for the verification protocol was fully sifted to analyze the experimental results using larger statistics. Note that this is not done for the protocol itself (see Sec. 6.3.1). The overall dataset was acquired over $\approx$ 4 days, with multiple measurements indicated by the Grey dashed lines. Each data point corresponds to an averaging window of $\approx$ 11 minutes.

In the verification protocol, only a small subset of variables is challenged, and consequently, only a few detected states are openly revealed. V accepts the verification only if all clauses are satisfied and P's MAC matches their own. A QBER, as known from QKD, is not computed during the verification protocol. For a comprehensive analysis of the experiment, however, we also evaluate the QBER for all detected states and distinguish between error probabilities for the four polarization states. Note that this would leak the full solution s of the problem

instance ϕ and is **not** done in the protocol itself.

Fig. 6.11 shows the signal rate ν_{sig} and the individual QBERs for the complete transmission duration of ≈ 80h. In the first data set, the QD source was not yet fully optimized, resulting in slightly lower signal rates. The remaining fluctuations in ν_{sig} are ascribed to changes in the pump power due to varying mode-locking conditions and subsequent spectral filtering in the 4f pulse shaper. The averaged signal rate is $\bar{\nu}_{\text{sig}} = 14.1(4)\,\text{kHz}$. Fig. 6.11b displays the QBER achieved in the presented experiment for all encoded input states. Slow polarization drifts in the fiber link have been periodically compensated by an automated routine (see Sec. 6.1.1.4). Averaging over all states and the full acquisition duration, we find a mean QBER of 0.48(24).

Based on the classically measured efficiencies of the complete setup—complemented by the specified SNSPD detection efficiencies—the system source efficiency (as defined in Sec. 2.2.2.1) of the QD is estimated. This efficiency corresponds to the probability of coupling one laser-filtered, polarized single photon per signal bin (i.e., per excitation pulse) into an SMF

$$\eta_{\text{sys}} = \frac{\eta_{\text{end-2-end}}}{\eta_{\text{ch}}\eta_{\text{enc}}\bar{\eta}_{\text{det}}\bar{\eta}_{\text{proj}}} = \frac{N_{\text{signals}}/N_{\text{bins}}}{\eta_{\text{ch}}\eta_{\text{enc}}\bar{\eta}_{\text{det}}\bar{\eta}_{\text{proj}}} = 8.5(17)\,\%\,. \tag{6.3-1}$$

The error for the total signal count is based on Poissonian statistics. The overall error is calculated by combining the errors of the individual efficiencies using Gaussian error propagation.

6.3.2 Security analysis

The complete protocol is recalled in the structured representation shown in Box. 1 and an overview of all the parameter values used for the experimental implementation is given in Tab. 6.1. In the following the security against a dishonest Prover P^* (Sec. 6.3.2.1) and dishonest V^* (Sec. 6.3.2.2) is derived. The full security analysis can be found in Appendix B.

6.3.2.1 Security against P^*

The overall security of our scheme combines two security bounds $\varepsilon_{\text{s,sat}}$ and $\varepsilon_{\text{s,q}}$. The first stems from the soundness gap, introduced via the PCP reduction: if ϕ is not satisfiable, then after the reduction, no more than $(1 - \frac{\frac{1}{8}-\delta}{8600})$ of the clauses can be satisfied by any assignment [400–402]. In the worst case, recalling that in a MAX-E3SAT(5) instance each variable appears exactly in 5 distinct clauses, the probability that V challenges only satisfied or stored clauses in k rounds is given by

$$\left(1 - \frac{\frac{1}{8} - \delta}{8600} + \frac{3 \cdot y_{\text{store}}}{L}\right)^k < \varepsilon_{\text{s,sat}}. \tag{6.3-2}$$

Box 1 Graphical representation of the limited-knowledge protocol.

R_LK(M: #clauses, L: #variables, k: #chall. clauses, N: #states/variable, ε: sec. par.)

	Prover P		**Verifier** V	
		Problem broadcast:		
3-SAT:	ϕ			
Solution:	$\overline{s} = \{\overline{s}_l\}^{\overline{L}}; \overline{s}_l \in \{true, false\}$	$SAT\ instance : \phi \longrightarrow$	$\psi \leftarrow R(\phi)$	
			$k, N \leftarrow params(\varepsilon)$	
Solution:	$s = \{s_l\}^{L}; s_l \in \{true, false\}$	$\longleftarrow E3SAT(5)\ instance : \psi$		
		Quantum Commitment:		
		`(i) Commit Step:`		
foreach	$\mathbf{l} \in \{\mathbf{1}, \ldots, \mathbf{L}\}$			
M. Basis:	$\tilde{\Theta} \in_R \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowright\}^N$	$\longleftarrow \otimes_{n=1}^{N} \lvert X_n\rangle_{\Theta_n}$	$X \in_R \{0,1\}^N; \Theta \in_R \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowright\}^N$	*Prepare States*
Measure:	$\tilde{X} \xleftarrow{\tilde{\Theta}} \otimes_{n=1}^{N} \lvert X_n\rangle_{\Theta_n}$			
Define Sets:	$\tilde{t}_{true} = \{\tilde{t}_j \mid \tilde{\Theta}_j = \leftrightarrow\updownarrow\}$			
	$\tilde{t}_{false} = \{\tilde{t}_j \mid \tilde{\Theta}_j = \circlearrowleft\circlearrowright\}$	$\tilde{t}_{s_l} \longrightarrow$		
	fail abort after Δ_t			
endforeach	$\mathbf{l} \in \{\mathbf{1}, \ldots, \mathbf{L}\}$			
		`(ii) Challenge Step:`		
			$m \xleftarrow{k} \phi$.`clauses`	*Chall. k Clauses*
Solutions of	$\mathcal{L} \subset \{1, \ldots, L\}$	$\longleftarrow m, \mathcal{L}$	$\mathcal{L} \xleftarrow{\lvert\mathcal{L}\rvert} m$.`variables`	*Chall.* $\lvert\mathcal{L}\rvert$ *Variables*
foreach	$\mathbf{l}' \in \mathcal{L}$			
Reveal $s_{l'}$		$s_{l'} \longrightarrow$	$t_{s_{l'}} = \{t_j \mid \Theta_j = \tilde{\Theta}_j \wedge t_j \in \tilde{t}_{s_{l'}}\}$	*Timings*
			$X_{s_{l'}} = \{X_{t_j} \mid t_j \in t_{s_{l'}}\}$	
	$\tilde{X}_{s_{l'}} = \{\tilde{X}_{t_j} \mid t_j \in t_{s_{l'}}\}$	$\longleftarrow Syn(X_{s_{l'}}), t_{s_{l'}}$		
endforeach	$\mathbf{l}' \in \mathcal{L}$			
use syndromes	$\hat{X}_{\mathrm{v}} = \hat{X}_{s_{1'}} \lVert \ldots \lVert \hat{X}_{s_{\lvert\mathcal{L}\rvert'}}$		$X_{\mathrm{v}} = X_{s_{1'}} \lVert \ldots \lVert X_{s_{\lvert\mathcal{L}\rvert'}}$	
		Verification Phase:		
		$\hat{X}_{\mathrm{v}} \longrightarrow$	$\hat{X}_{\mathrm{v}} = X_{\mathrm{v}}$	*compute*
			return $\hat{X}_{\mathrm{v}} \stackrel{?}{=} X_{\mathrm{v}}$ *and* $m \stackrel{?}{=}$ **true**	*evaluate*

Table 6.1: Protocol parameters. All experimental parameters are inferred from measurements, whereas the security parameter r is fixed. The remaining parameters are derived from the constraints provided in the security analysis.

Experimental parameters	
N_{enc}	$1.46 \cdot 10^{11}$
η	51%
$\{p_0, p_1, p_{\text{m}}\}$	$\{0.946, 0.054, 3.1 \cdot 10^{-5}\}$
p_{dc}	$5 \cdot 10^{-6}$
QBER	0.5%
Instance parameters	
L	2 276 985
y_{store}	1
k	90 991
Secure reuses	5
Bounded quantum storage	
r	$\to 1$
$\nu_{\varepsilon_{\text{s,sat}}}$	0.14
$S_{\max}$	117
Security parameters	
$\varepsilon_{\text{s,sat}}$	$3 \cdot 10^{-1}$
$\varepsilon_{\min}$, ε_{det}	$5 \cdot 10^{-3}$
ε_{key}	$2^{-7} \approx 0.0078$
$\varepsilon_{\text{s,q}}$	$2.28 \cdot 10^{-2}$

Here, $\varepsilon_{\text{s,sat}}$ is the probability that an adversary can cheat without being caught. This gives a threshold

$$y_{\min} = \frac{L}{3} \cdot \left(\varepsilon_{\text{s,sat}}^{1/k} - \left(1 - \frac{\frac{1}{8} - \delta}{8600} \right) \right), \tag{6.3-3}$$

for which the verification is secure with probability $(1 - \varepsilon_{\text{s,sat}})$. Note that for fixed δ and $\varepsilon_{\text{s,sat}}$, the number of challenges should be lower-bounded by

$$k_{\min} = \frac{\ln \varepsilon_{\text{s,sat}}}{\ln\left(1 - \frac{\frac{1}{8} - \delta}{8600} \right)} \tag{6.3-4}$$

Conversely, for a fixed storage size S, P^* may try to increase y_{store} to boost the likelihood of the event *"no unsatisfied or un-stored clause is challenged"*, at the cost of increasing their uncertainty about X_{v}, the string needed for passing the verification check. In other words, increasing y_{store} decreases their storage rate

$$\nu_{\varepsilon_{\text{s,sat}}} = \frac{S}{y_{\text{store}} \cdot N^{(1)}_{\text{detect,min}}}, \tag{6.3-5}$$

where $N^{(1)}_{\text{detect,min}}$ is the number of correctly measured single-photon states per variable later defined in Eq. (6.3-10).

We choose the number of states per variable such that, whenever P^* stores more than $y_{\min}$ variables (or their corresponding quantum states), their storage rate becomes too small to allow a reliable readout of those variables. Consequently, if $y_{\text{store}} > y_{\min}$, the probability guessing X_{v} is upper-bounded by ε_{key} (see Eq. (6.3-12)). This bound holds except with probability $(1 - \frac{\frac{1}{8}-\delta}{8600})^k$, corresponding to the event that only satisfied clauses are challenged. The total soundness is therefore $\varepsilon_s = \varepsilon_{\text{s,sat}} + \varepsilon_{\text{s,q}}$, where $\varepsilon_{\text{s,q}} = \varepsilon_{\text{key}} + \varepsilon_{\min} + 2\varepsilon_{\text{det}}$, which, for the parameters used here, evaluates to $\varepsilon_s \approx 1/3$ (see Tab. 6.1).

With complementary probability $1 - (1 - \frac{\frac{1}{8}-\delta}{8600})^k$, at least one challenged clause will be unreliably stored or un-satisfied by the adversary. In this setting, their best-case scenario is one in which all challenged clauses are un-stored and satisfied (i.e. no variable is measured in the incorrect basis) *except* for one stored variable. Let $\mathcal{V}$ denote the set of stored, challenged variables and $X^{\mathcal{V}}$ the corresponding string of measurement outcomes for which P and V chose the same basis. We determine how large N must be so that the verification string X_{v} retains 7-bit security except with probability $\varepsilon_{\min} + 2\varepsilon_{\text{det}}$, for the case $|\mathcal{V}| = 1$. To do so, we use techniques introduced in [106, 408] to bound the adversary's uncertainty about the verification string X_{v} quantified by the following min-entropy relation:

$$H_\infty(X_{\text{v}} \mid K_{\text{v}}\, Q_{\text{v}}) \geq -\log p_{\text{guess}}(X_{\text{v}} \mid K_{\text{v}}\, Q_{\text{v}}). \tag{6.3-6}$$

Here, p_{guess} denotes the probability that an adversary correctly guesses X_{v} by performing an optimal measurement on the quantum register Q_{v}, conditioned on classical side-information K_{v}. Note that X_{v} is defined as the concatenation of measurement outcomes from states corresponding to the challenged variables $l' \in \mathcal{L}$, measured in the basis required to satisfy the clauses.

Let us now divide $X^{\mathcal{V}}$ into two substrings $X^{\mathcal{V}}_{\text{correct}}$ and $X^{\mathcal{V}}_{\text{incorrect}}$ containing measurement outcomes from correct and incorrect basis choices (associated with correct and incorrect assignments, respectively). Since we assume that all but $|\mathcal{V}|$ variables are measured in the correct basis, we must bound the adversary's uncertainty about at least one of the two basis substrings. We can then relate the lower bound on the smooth min-entropy (with smoothing parameter $\varepsilon_{\min}$) of the secret input X_{v} conditioned on classical side-information to that of the substring $\{X^{\mathcal{V}}_{\text{D}}\}_{D\in_R\{\text{correct,incorrect}\}}$:

$$H_\infty^{\varepsilon_{\min}/2}(X_{\text{v}}|K_{\text{v}}\Theta_{\text{SK}}) = H_\infty^{\varepsilon_{\min}/2}(X^{\mathcal{V}}_D|D\; K_{\mathcal{V}}\Theta_{\mathcal{V}}) \tag{6.3-7}$$

We calculate the lower bound on $X^{\mathcal{V}}_D/X_{\text{v}}$ using the uncertainty relation from [352], along with the min-splitting theorem, which yields

$$H_\infty^{\varepsilon_{\min}/2}(X_{\text{v}}|K_{\text{v}}\Theta_{\text{SK}}) \geq \frac{|\mathcal{V}| \cdot N^{(1)}_{\text{detect,min}} \cdot c_{BB84}}{2} - 1, \tag{6.3-8}$$

where

$$c_{BB84} = \max_{s\in(0,1]} \frac{1}{s}\left[1+s-\log(1+2^s)\right] - \frac{1}{s\cdot|\mathcal{V}|\cdot N^{(1)}_{\text{detect,min}}} \log\left(\frac{8}{\varepsilon_{\min}}\right). \tag{6.3-9}$$

Furthermore, we assume that multiphoton states do not contribute to the adversary's uncertainty. Considering trusted photon-number generation probabilities $\{p_0, p_1, p_{\text{m}}\}$ and accounting for a non-unit channel efficiency η_{ch}, the number of single-photon states per variable satisfies

$$\begin{aligned} N^{(1)}_{\text{detect}} \geq & \left(\left(\frac{(1-\tilde{\sigma}_{1-})\cdot\tilde{p}_1}{2} + \frac{(1-\tilde{\sigma}_{\text{m}-})\cdot\tilde{p}_{\text{m}}}{2} \right.\right. \\ & \left. + \frac{(1-\tilde{\sigma}_{\text{dc}-})\cdot\tilde{p}_{\text{dc}}}{2}\right) \\ & \left. - \left((1+\sigma_{\text{m}+})\cdot p_{\text{m}} + (1+\tilde{\sigma}_{\text{dc}+})\cdot\tilde{p}_{\text{dc}}\right)\right)\cdot N \\ & = N^{(1)}_{\text{detect,min}} \end{aligned} \tag{6.3-10}$$

except with probability ε_{det}, where $\sigma/\tilde{\sigma}$ accounts for finite-size statistics (detailed in Appendix B). Here, $\{\tilde{p}_0, \tilde{p}_1, \tilde{p}_{\text{m}}\}$ are the photon-number probabilities undergoing channel losses and $\tilde{p}_{\text{dc}}$ are the expected detector dark counts. In fact, Eq. (6.3-8) determines the rate at which an adversary needs to send states through their quantum storage [106], i.e.

$$R = \frac{1}{y_{\text{store}}\cdot N^{(1)}_{\text{detect,min}}}\left(\frac{|\mathcal{V}|\cdot N^{(1)}_{\text{detect,min}}\cdot c_{BB84}}{2} - 1 - \log\left(\frac{2}{\varepsilon_{\min}}\right)\right). \tag{6.3-11}$$

Together with Lemma II.2. in [106], this gives a lower bound on the uncertainty about X_{v} conditioned on quantum and classical side-information. Lastly, we use the strong converse property of the depolarizing channel to obtain

$$\begin{aligned} & H^{\varepsilon_{\min}}_{\infty}(X_{\text{v}}|K_{\text{v}}\,\Theta_{\text{SK}}\,\mathcal{F}(Q_{in})Syn(X^{\mathcal{V}}_D)) \geq \\ & \nu_{\varepsilon_{\text{s,sat}}}\cdot y_{\text{store}}\cdot N^{(1)}_{\text{detect,min}}\cdot\gamma^{\mathcal{N}}\left(\frac{R}{\nu_{\varepsilon_{\text{s,sat}}}}\right) \\ & - 1.2\cdot\frac{|\mathcal{V}|\cdot N_{\text{detect,min}}}{2}\cdot h(\tilde{p}_{\text{error}_+}) \\ & \geq -\log(\varepsilon_{\text{key}}), \end{aligned} \tag{6.3-12}$$

where $N_{\text{detect,min}}$ is the minimum expected number of states per variable s_l that P^* detects and measures in the correct basis, except with probability ε_{det}. This final expression determines the bit security of the verification string X_{v}. $\varepsilon_{\text{s,q}} = \varepsilon_{\min} + 2\varepsilon_{\text{det}} + \varepsilon_{\text{key}}$ is the probability that an adversary obtains the correct verification string, which incorporates errors arising from smoothing the min-entropy, the expected single-photon detection events per variable and the probability of guessing X_{v} (the first two contributing $5 \cdot 10^{-3}$, and the latter 2^{-7}).

6.3.2.2 Security against V^*

Commitment phase

V^* does not initially know the satisfying assignment for the input instance ψ, so P's measurement basis choice—and hence their committed bit—appears random to V^*. A hypothetical memory-unbounded P could indeed create a *dummy* commitment that may later be opened as either bit without altering V^*'s view. Specifically, for each variable y_l, this P could declare a random timing set $\tilde{t}_{s_l} \subset t_l$ of size at least $N_{\text{detect,min}}$, store the corresponding quantum states, and await V^*'s challenges. Only after the challenges are revealed does P measure the states and output $\tilde{X}_{s_l}$. The only difference between the honest and hypothetical unbounded P is their measurement basis choice, which V^* cannot learn since they do not have access to P's laboratory. This provides the hiding property of our scheme.

In the presence of experimental imperfections, the only information obtained by V^* before the challenge phase is the timing string $\tilde{t}_{s_l}$ provided by P in the commit phase, revealing the rounds for which the states were measured in the correct commitment basis. If $\tilde{t}_{s_l}$ for $s_l = true$ is distinguishable from $\tilde{t}_{s_l}$ for $s_l = false$ (i.e. the set $\tilde{t}_{s_l}$ cannot be chosen uniformly at random for either bit value), then V^* could distinguish between the two values. This could occur if detectors do not behave symmetrically or if P explicitly back-reports the type of click configuration that occurred on their detectors: by sending brighter polarized pulses, V^* can indeed remotely probe the measurement basis chosen by P [154]. To prevent this leakage, any detection round yielding more than one click is uniformly mapped by P onto one of the four detector outcomes.

Verification phase

From Eq. (6.1-3), it is NP-hard to distinguish $\text{OPT}(\psi) = 1$ (YES case) from $\text{OPT}(\psi) \leq 1 - \frac{\frac{1}{8}-\delta}{8600}$ (NO case), since this is equivalent to deciding whether $\phi \in \text{SAT}$. We now establish a *hardness under bounded leakage* property: even if an algorithm is given the values of $B = O(1)$ variables of ψ (i.e., $O(1)$ bits of information about a satisfying assignment), this does not yield a polynomial-time procedure that can reliably distinguish the YES from the NO cases on all instances.

Assume there exist constants $B_0, \beta > 0$ and a polynomial-time routine V^* such that for *every* MAX-E3SAT(5) instance ψ, whenever $\text{OPT}(\psi) = 1$, there is a set

$S(\psi)$ of at most B_0 distinct variables and values satisfying

$$f_{V^*(\psi,S(\psi))}(\psi) \;\geq\; 1 - \frac{\frac{1}{8} - \delta}{8600} + \beta,$$

where $V^*(\psi, S(\psi))$ is the assignment produced by V^* when given the instance ψ and its leaked set $S(\psi)$. Intuitively, this assumption requires that there exists a polynomial-time algorithm V^* which, whenever ψ is a YES instance, can exploit a leakage of at most B_0 distinct variables to produce an assignment that exceeds the hardness gap by a fixed value β. We show in Appendix B that if such a V^* existed, then one could construct a polynomial-time algorithm that *distinguishes* the two cases YES and NO *without any leakage at all*, thus contradicting NP-hardness.

We simulate the expected number of variables that V^* can learn after verifying the proof G times. First, we generate a *planted, balanced 3-5CNF* instance from a random variable assignment $s \in_R \{true, false\}^L$ using our algorithm *cnfgen* (see Sec. 6.3.4). In each verification round $g \in \{1, \ldots, G\}$, V^* samples k clauses, which are gathered in a set m_g. Let $\mathcal{L}_g \subseteq \{1, \ldots, L\}$ be the set of variables that appear in m_g and let $\mathcal{B}$ be the set of variables known so far. Define

$$\mathcal{B}_g = \mathcal{L}_g \setminus \mathcal{B},$$

the variables newly revealed in round g whose values are learned *directly* from the sampled clauses. Update $\mathcal{B} \leftarrow \mathcal{B} \cup \mathcal{B}_g$.

Next, identify variables that become *indirectly* known given the updated set $\mathcal{B}$. A variable is indirectly revealed if it appears in a clause together with two variables whose values are already known and both evaluate to *false* under s, since in this case the third variable must evaluate to *true* for the clause (and hence the instance) to remain satisfied. Let $\mathcal{L}'_g$ denote the set of variables that are newly revealed in this way during round g, and define

$$\mathcal{B}'_g = \mathcal{L}'_g \setminus \mathcal{B},$$

Update $\mathcal{B} \leftarrow \mathcal{B} \cup \mathcal{B}'_g$. The total set of variables learned in round g is

$$\tilde{\mathcal{B}}_g = \mathcal{B}_g \cup \mathcal{B}'_g.$$

After G rounds, the total number of known variables is $B = |\mathcal{B}|$.

To upper-bound leakage, we adopt an adaptive (adversarial) sampling model that selects one clause at a time rather than k clauses at once. This strengthens the adversary, since each choice can exploit the currently known set $\mathcal{B}$ to maximize indirect leakage. Concretely, when a clause m_i reveals a new set of variables $\{y_j\}$, we identify clauses $\tilde{m}_i$ in which these $\{y_j\}$ occur and in which they evaluate to *false* under s. Let $\{\tilde{y}_i\}$ denote the set of not-yet-known variables in $\tilde{m}_i$. In the next round, V^* challenges a clause in which most of $\{\tilde{y}_i\}$ occur. With this strategy, they learn on average $\mu(q)$ variables per challenge (see Fig. 6.9).

The purpose of this simulation is to estimate how much information about the planted solution is disclosed when the same instance is verified over multiple rounds with constant soundness, using this adversarial sampling model. However, note that we do not perform the full series of reductions, but instead generate CNF instances that share *key structural properties* of the reduced instances as would occur in an actual implementation of the protocol. The exact instances produced by the full reduction chain may indeed exhibit additional structural features, such as locally sparse or dense regions of clause-variable connectivity, which we do not attempt to reproduce. However, our sampling method is constructed to maximize potential leakage by exploiting structural connections in a way that an adversary in the real protocol cannot, since the protocol forces them to sample k clauses simultaneously rather than optimizing per clause. This gives our simulated adversary strictly more freedom than any adversary interacting with the actual protocol. For this reason, we believe that our leakage simulation overestimates the capabilities of an adversary and provides an upper bound on the leakage.

Let $\overline{B_0} = O(1)$ be an a priori upper-bound on the total number of solution bits we are willing to leak. The simulation yields $B = \mathfrak{B}(G, k, L)$ as a function of the number of rounds G, the number of challenged clauses k, which is tied to the soundness gap in Eq. (6.1-3), and L, the size of the instance. We then define

$$G_{\max,L} = \max\{G : \mathfrak{B}(G, k, L) \leq \overline{B_0}\} \tag{6.3-13}$$

Verifying an instance of size L up to $G_{\max,L}$ rounds, so that at most $\overline{B_0} = O(1)$ bits are leaked, preserves NP-hardness of the verified problem.

6.3.3 Extended parameter variations

In addition to the results presented in Sec. 6.2, we now explore verification performance across a wide range of parameters to develop intuition for our protocol and provide an overview of how different photonic sources would perform under these conditions.

In particular, the protocol's noise and loss tolerance determine its practicality for any implementation outside a laboratory environment. Therefore, we analyze how the maximally tolerable QBER changes with varying transmission η, and number of encoding cycles N and present the results in Fig. 6.12 for the photon-number distribution of our QD. Note that η includes not only the channel loss but all other imperfections occurring outside of V's trusted lab, for example, imperfect state detection. The white region indicates where a QBC is no longer possible. Choosing any other $QBER > 0$ provides a secure lower bound on N as shown in Fig. 6.9. Here, no additional trusted attenuation was included to optimize the photon statistics. We verify that increasing N improves the noise tolerance for any η since the enhanced detection statistics for larger N allow for correcting for more noise. However, the improvement in noise tolerance is limited, and the QBER cannot exceed $\approx 7.2\%$ for any value of η.

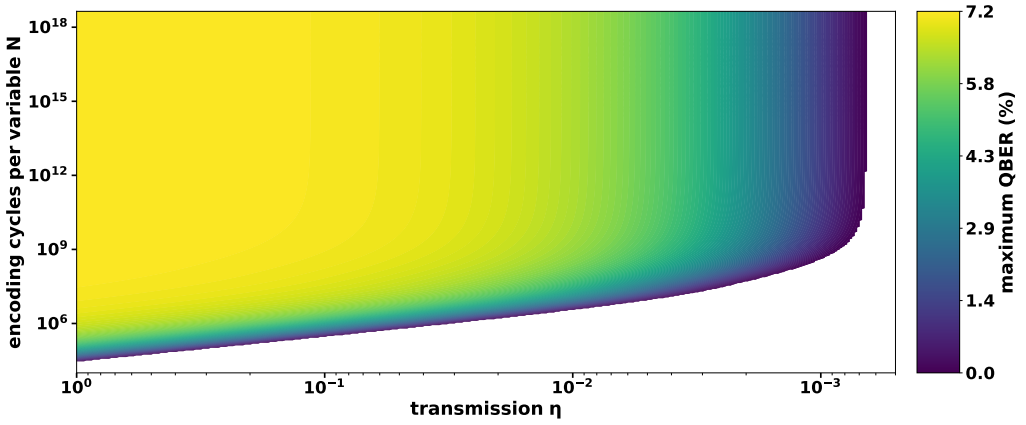


Figure 6.12: **Maximally tolerable quantum bit error rate.**. For each pair of transmission η and encoding cycles per variable N, the maximal QBER that still achieves a secure protocol is calculated. For $QBER > 7.2\%$, the protocol can no longer be secured, independent of transmission η or number of encoding cycles N. The simulations used the same parameters as for Fig. 6.9, that is $[p_0, p_1, p_\mathrm{m}] = [0.946, 0.054, 3.1 \cdot 10^{-5}]$, problem size $L = 2276985$, quantum memory storage size $S = 117$, quantum memory depolarization parameter $r = 1 - 10^{-5}$, dark count probability $\tilde{p}_\mathrm{dc} = 10^{-7}$, $QBER = 5 \cdot 10^{-3}$, $\varepsilon_{sc} = 3 \cdot 10^{-1}$, $\varepsilon_\mathrm{corr} = 1.67 \cdot 10^{-3}$, $\varepsilon_\mathrm{sound} = 1.25 \cdot 10^{-3}$. No photon number optimization with trusted attenuation was performed here.

Fig. 6.13 shows how the verification performance varies as a function of transmission when scanning the experimental parameter. As in Sec. 6.2, we use the photon-number distribution emitted by our QD for the SPSand include the loss caused by the state encoding setup ($\eta_\mathrm{enc} = 64\%$, see Sec. 6.1.1). Changing the acceptable QBER as in Fig. 6.13a has a two-fold impact. First, both types of sources need more encoding cycles for any given transmission and exhibit lower maximum loss thresholds. Second, we find that SPSs are robuster against noise, in particular in the low-loss regime. This is signified by the smaller increase in N_min compared to PDSs. For $QBER > 4.5\%$, PDSs lose their only advantage and are also outperformed for high-transmission implementations. Finally, we note that the increase in N_min with rising QBER is significant. For example, tolerating $QBER = 5\%$ instead of 0.5% must be compensated by encoding ≈ 500 times more cycles when using a PDS but < 10 times more when using out QD.

The impact of the dark count probability (see Fig. 6.13b) strongly resembles the behavior known from QKD [1] and other protocols. Reducing $\tilde{p}_\mathrm{dc}$ allows for longer communication distances, since more loss can be tolerated before the probability of detecting a signal photon approaches that of registering a dark count click. Furthermore, as the performance difference between QDs and PDSs increases with loss, SPSs benefit even more from reducing $\tilde{p}_\mathrm{dc}$ than PDSs.

In Fig. 6.13c, the system efficiency of the QD source is increased to project how technological advancements will impact the verification. In the low-loss regime,

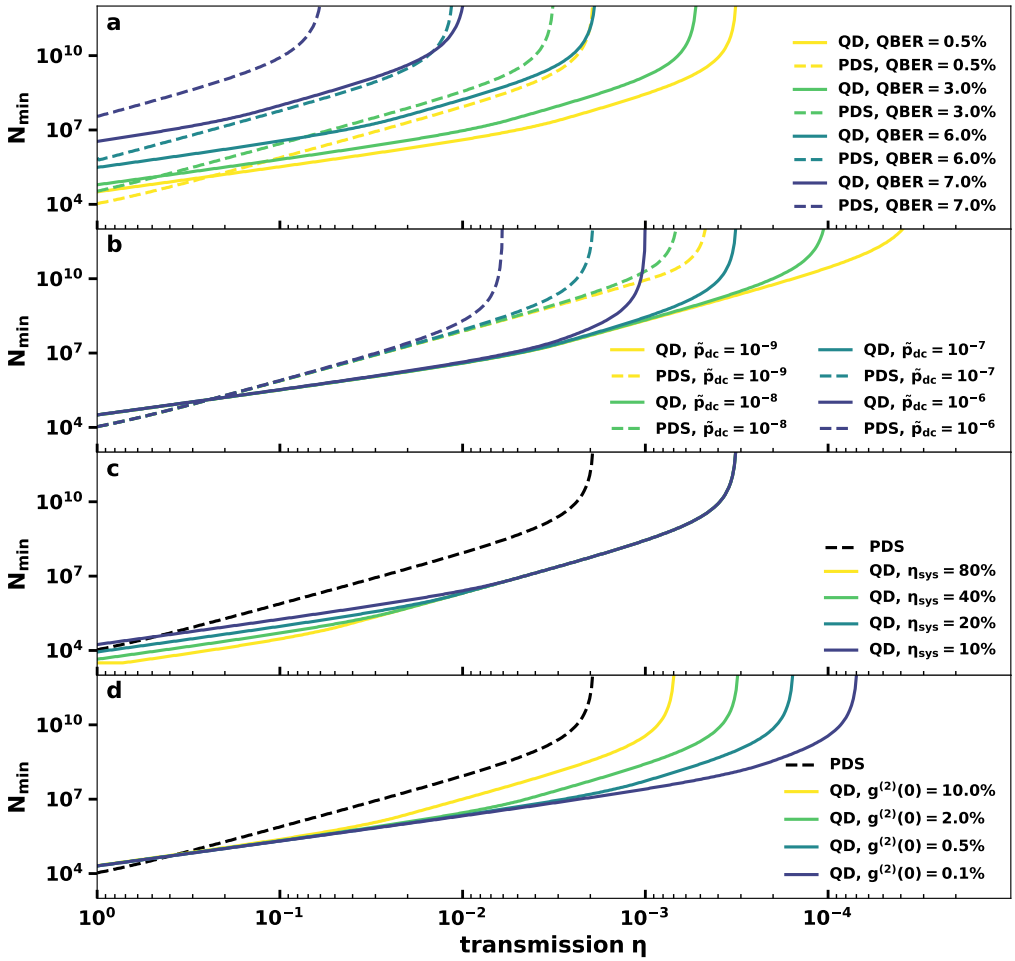


Figure 6.13: **Verification performance under parameter variations.** The minimally required number of encoding cycles for a secure verification as a function of transmission. Where not indicated otherwise, the parameters are: photon-number statistics of the QD $[p_0, p_1, p_{\mathrm{m}}] = [0.946, 0.054, 3.1 \cdot 10^{-5}]$, problem size $L = 2276985$, quantum memory storage size $S = 117$, quantum memory depolarization parameter $r = 1 - 10^{-5}$, dark count probability $\tilde{p}_{\mathrm{dc}} = 10^{-7}$, $QBER = 5 \cdot 10^{-3}$, $\varepsilon_{sc} = 3 \cdot 10^{-1}$, $\varepsilon_{\mathrm{corr}} = 1.67 \cdot 10^{-3}$, $\varepsilon_{\mathrm{sound}} = 1.25 \cdot 10^{-3}$. The photon-number statistics of QD and PDS are optimized for each transmission value using trusted attenuation. a) Variation of the tolerable QBER. b) Variation of the dark count probability. c) Variation of the QD system efficiency for fixed $g^{(2)}(0) = 2.07\%$. d) Variation of the QD multiphoton system contribution for fixed system efficiency $\eta_{\mathrm{sys}} = 8.5\%$. The resulting photon-numbers for c) and d) are calculated according to Sec. 2.2.2.4.

$N_{\min}$ can be significantly reduced when improving the system efficiency. However, already for intermediate loss the advantage reduces as the source must be attenuated to limit the multiphoton component of the quantum state. Enhancing the single-photon purity of the source, on the other hand (see Fig. 6.13d) has little effect for low loss but defers trusted attenuation toward higher loss and thus continuously improves the performance of typical sources for $\eta < 1\%$. This behavior of source efficiency and single-photon purity is also known from QKD with SPSs [1, 42].

6.3.4 3-5 CNF generation and benchmarking

To estimate the maximum amount of disclosed variables in a verification round, when V^* challenges k clauses of ψ, we simulate the interaction and propagation of information on random satisfiable 3-5CNF instances that we generated with our *cnfgen* algorithm 1. The generator works as follows. We first fix a problem size L, the number of distinct boolean variables in our instance. We then sample a random truth assignment $s \in_R \{true, false\}^L$, which serves as a planted solution to the 3-5CNF instance.

For each selected variable y_ℓ in a clause, we randomly choose to include it either as a positive literal y_ℓ or as a negated literal $\neg y_\ell$. After forming a clause with three literals, we evaluate whether it is satisfied by the planted assignment s. A clause is satisfied if at least one literal evaluates to *true* under s. If the clause is already satisfied, we include it in the instance without modification. If the clause is not satisfied by s, we select one of its three literals uniformly at random and flip it to be consistent with s: if the variable is assigned *true*, we include it as a positive literal; if it is assigned *false*, we include it as a negated literal. This guarantees that the clause becomes satisfied. Repeating this process for all M clauses, we construct a CNF formula in which every variable appears in exactly five clauses, but not twice in the same clause, and where every clause contains three variables that is guaranteed to be satisfiable by the known assignment.

We now model the behavior of a *dishonest* V^* who aims to learn the satisfying assignment of the problem instance $\psi = R(\phi)$. Hence, V^* assumes that P is honest and that $\phi \in SAT$. Under these assumptions, V^* will try to exploit the structure of the SAT instance to induce *indirect leakage* of variables. More specifically, when V^* finds a clause for which they already know two variables, which happen to evaluate both to *false*, V^* can infer that the last variable must evaluate to *true*. Otherwise, this clause and therefore the whole instance would be unsatisfied. This is the key element that allows V^* to increase the leakage rate μ beyond the limit of direct leakage of $|\mathcal{L}|/k \leqslant 3$, and the strategy of the dishonest Verifier is to select challenges that maximize the indirect inference.

As mentioned in the Sec. 6.2, we allow V^* to challenge one clause at a time and update their knowledge base before each new challenge. Clearly, this overestimates the power of V^* compared to the actual protocol, where V^* must select k clauses at once and cannot fully optimize the choice.

Algorithm 1 *cnfgen*(L)

Require: $5L$ divisible by 3
$M \leftarrow \frac{5L}{3}$
Y $\leftarrow [1, 2, \ldots, L]$ ▷ variable indices
cnf $\leftarrow$ CNF()
▷ Step 1: sample a planted assignment
for all $y \in$ Y **do**
    s$[y] \leftarrow\in_R \{true, false\}$
end for
▷ Step 2: enforce exactly 5 occurrences per variable
for all $y \in$ Y **do**
    deg$[y] \leftarrow 5$
end for
▷ Step 3: build M clauses with 3 distinct variables each
for $i = 1$ **to** M **do**
    available $\leftarrow \{\, y \in$ Y : deg$[y] > 0 \,\}$
    m $\leftarrow$ Sample_without_replacement(available, 3)
    for all $y \in$ m **do**
        deg$[y] \leftarrow$ deg$[y] - 1$
    end for
    clause $\leftarrow$ empty list
    ▷ Step 4: choose random signs for literals
    for all $y \in$ m **do**
        $\ell \leftarrow\in_R \{y, \neg y\}$
        append ℓ to clause
    end for
    clause_satisfied $\leftarrow$ False
    for all $\ell \in$ clause **do**
        $y \leftarrow |\ell|$
        value $\leftarrow$ s$[y]$
        if ($\ell > 0$ **and** value $=$ *true*) **or** ($\ell < 0$ **and** value $=$ *false*) **then**
            clause_satisfied $\leftarrow$ True
            break
        end if
    end for
    if not clause_satisfied **then**
        idx $\leftarrow\in_R \{1, 2, 3\}$
        $y \leftarrow$ m[idx]
        if s$[y] =$ *true* **then**
            clause[idx] $\leftarrow y$
        else
            clause[idx] $\leftarrow \neg y$
        end if
    end if
    append clause to cnf
end for
return cnf, s

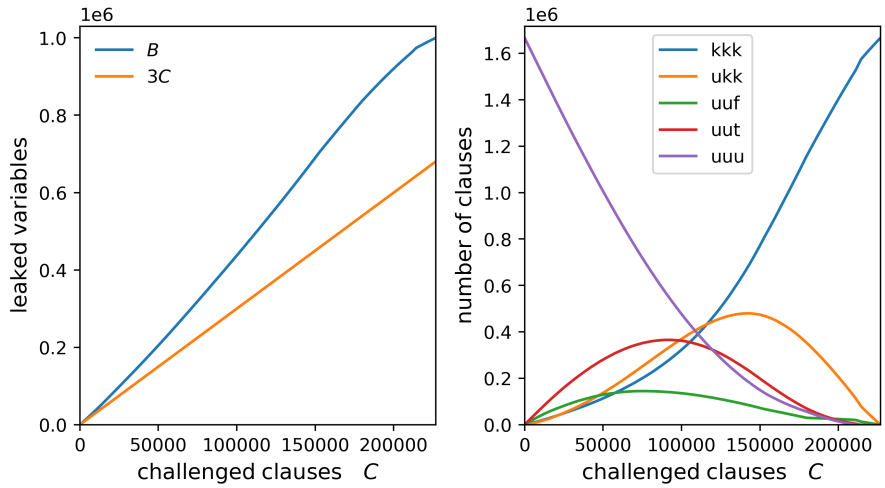


Figure 6.14: **Performance of V^*.** For a randomly generated 3-5CNF instance containing $L = 10^6 + 2$ variables and $M = 5L/3$ clauses, V^*'s strategy explained in the text is applied. The left panel shows the simulated total leakage B compared to the direct leakage bound $|\mathcal{L}| = 3C$. The bound ignores the fact that at some point no fully unknown clauses are left, as it purely serves as a reference for the simulated leakage. Modeling the interaction using $k = 1$, V^* can select each challenge based on the full knowledge they obtained so far about the solution, which overestimates the power compared to $k > 1$ as in the protocol. Each clause is assigned one of the possible five states $\{uuu, uut, uuf, ukk, kkk\}$ (see text), and the status is updated after every challenge.

Throughout their interaction with P, V^* will update the status of each clause depending on the current knowledge about this clause. The possible states are

- *'uuu'*: all three variables are unknown
- *'uut'*: two variables are unknown, the other is known and evaluates to *true*
- *'uuf'*: two variables are unknown, the other is known and evaluates to *false*
- *'ukk'*: one variable is unknown, two are known
- *'kkk'*: all three variables are known

Initially, all clauses are fully unknown, i.e., their status is *'uuu'*. To generate indirect inferences, *'uuf'* is the most relevant status. If one of the two unknown variables of this clause is challenged, there is a chance of $p_{\text{ind}} = 0.5$ that the variable evaluates to *false* and thus causes the remaining variable to be indirectly revealed. On the other hand, clauses that are *'uut'* or *'ukk'* cannot produce any indirect inferences since at least one variable already evaluates to *true*. Since each clause consists of 3 variables and each of those appears in exactly 4 more clauses, V^* might be able to

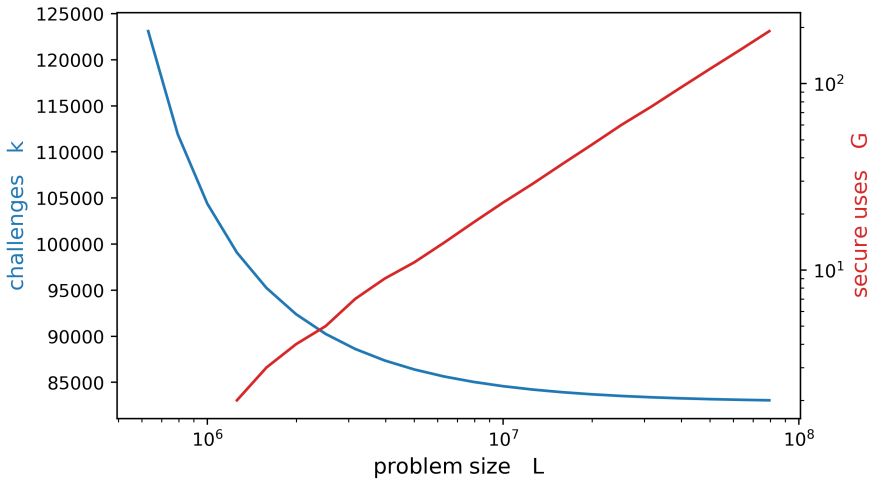


Figure 6.15: **Conditions for a secure protocol.** The number of challenges k that V must send depends on the security parameters and the size of the problem instance, as given by Eq. 6.3-14. Based on this k, and the maximum fraction of clauses that can be challenged $q_{\max}$, the number of secure uses of a specific problem instance is calculated as given in Eq. 6.3-15. Here, $\varepsilon_{\mathrm{s,sat}} = 0.3$, $\delta = 10^{-6}$ and $q_{\max} = 0.12$.

choose challenges such that not just one, but up to $N_{uuf} \leqslant 12$ indirect inferences are triggered. In general, the expected indirect leakage is $B_{\mathrm{ind}} = p_{\mathrm{ind}} N_{uuf}$ and the total leakage in one challenge is $B = B_{\mathrm{dir}} + B_{\mathrm{ind}}$. Importantly, if any variable was leaked indirectly, V^* checks if this triggered another *'uuf'* clause to reveal yet another variable such that the leakage caused by a single challenge, $\mu = \frac{dB}{dC}$, is theoretically unbounded.

Using this approach, the dishonest Verifier calculates the expected B for each possible clause and chooses the one that maximizes B. In case of a tie, a random maximal clause is selected. In the first rounds, when very few variables are known, this strategy effectively samples randomly from the set of *'uuu'* clauses, but as more knowledge about the solution is accumulated, the pool of *'uuf'* clauses builds up, and on average, more indirect leakage is generated every challenge. This is captured by the leakage rate μ shown in Fig. 6.9.

How the status of all clauses changes when challenging the whole problem step by step is depicted exemplarily for a problem instance with $L = 10^6 + 2$ in Fig. 6.14. Towards the end, at around challenge number $1.8 \cdot 10^5$, we can see that the number of clauses in *'uuf'* is almost constant for about $4 \cdot 10^4$ challenges before it is finally depleted. This indicates that when most of the problem is already known, challenging a random *'uuu'* clause ($B = 3 + 0$) is more beneficial than targeting some indirect leakage by challenging a clause where at least one variable is already known, that is $B = 2 + 1 \cdot p_{\mathrm{ind}}$.

The described dishonest Verifier model only looks a single challenge ahead

instead of analyzing the full problem structure. The latter requires complex clustering routines and significantly more computing power, especially if the strategy needs to be updated after each challenge. However, we believe that allowing V^* to choose every challenge separately rather than in bundles of k (where $k > 8 \cdot 10^4$ in our implementation) grants even more power to V^*.

As mentioned in Sec. 6.2 and also shown in the left panel of Fig. 6.14, the leakage rate remains bounded during the continued challenging. We therefore conclude that the problem does not get easier all of a sudden, but may be used until a significant portion of all clauses have been challenged. We define this limit as $q_{\max} = 12\,\%$, where $q = C/M$ and M is the total number of clauses.

To decommit, the dishonest Prover must be able to store at least one variable in their quantum memory (remember that, depending on N and the experimental parameters, this might require a memory capable of storing hundreds of qubits perfectly). Hence, rearranging Eq. 6.3-3 and setting $y_{\text{store}} = 1$, we can infer how many challenges k must be asked by V to achieve the protocols soundness parameter

$$k = \frac{\log(\varepsilon_{\text{s,sat}})}{\log\left(\frac{3}{L} + 1 - \frac{1/8-\delta}{8600}\right)} \tag{6.3-14}$$

Fig. 6.15 illustrates how k scales with L and indicates that it is asymptotically constant in L for large problems.

Knowing how many challenges are allowed per use and having bounded the maximum number of challenges a particular instance can tolerate without fully revealing the solution assignment, we can infer how often a problem may be used safely

$$\begin{aligned} kG &\leqslant q_{\max} M \\ G_{\max=} &= q_{\max} \frac{5}{3} \frac{L}{k} \,. \end{aligned} \tag{6.3-15}$$

Since k is independent of L for large L, $G_{\max} \propto L$ as shown in Fig. 6.15.

7 Conclusion

Since the advance of QDs as sources of quantum light, they have become not only a system of choice for applications requiring the highest performance [44, 55], but also matured to be a prime platform for exploring fundamentals of quantum optics [63, 64, 302, 338, 415]. Over the last two decades, technology and applications have advanced in tandem: on the technology side, growth quality and control over emitter properties through micro- and nanostructures have improved significantly [56, 296, 416], while on the application side, the range of use-cases is constantly extended and optimized [44, 46, 76, 241]. Newly discovered phenomena create new opportunities, while the requirements of various applications set priorities, shaping the progress in system design. This thesis reflects on the reciprocity by presenting work at the intersection: studying quantum optics with QDs and using them as sources for quantum technologies, specifically quantum-secured communication.

7.1 Phonon-assisted excitation for single-photon applications

The pivotal point of this work is the LA phonon-assisted excitation of a QD, exploring both details of the light–matter interaction and showing how to leverage its advantages for quantum cryptography. The excitation–phonon coupling gives rise to several favorable features when employing QDs as a SPS for quantum-secured applications, and this thesis presented a comprehensive study of the excitation scheme in light of cryptography. Prominent examples include its straightforward suppression of the pump field via spectral filtering and the absence of PNC [219], which alleviates the need for phase scrambling otherwise required by many protocols [45].

In Chapter 4, the photon-number statistics have been studied experimentally across a large parameter space, demonstrating that system efficiency and multiphoton noise exhibit broad maxima and minima, respectively, ensuring good SPS performance for a wide range of laser settings. These results indicate the scheme's robustness against laser fluctuations, as theoretically predicted [95] and previously shown for the system efficiency [235, 260, 264], but not for single-photon purity. The behavior is a direct consequence of the large spectral bandwidth of the

exciton–phonon coupling $J(\omega)$ compared to the sharp resonance conditions found in RF or TPE.

Additionally, we showed that the photon statistics cannot be trivially optimized, as no laser setting maximizes both system efficiency and single-photon purity. Thus, parameter optimization requires a figure of merit that weights the photon-number probabilities according to the application. Notably, the ideal operating point changes with the protocol parameters, particularly channel loss, as demonstrated by the example of QKD. Finally, this study showed that the decoy-state technique provides little benefit for highly pure SPSs such as QDs.

To maintain an experimentally manageable parameter space, the pulse length is unchanged throughout the first work. However, for resonant excitation, it was shown that some multiphoton contribution is inevitably created in a re-excitation event due to the finite pump pulse lengths [96, 97]. Rethinking this process for phonon-mediated excitation in the dressed-state picture motivated the investigations presented in Chapter 5. By developing an experimental technique to extract the individual temporal shapes of the first and second photons produced in such a re-excitation process, we demonstrated their ordering for the first time directly. Extending the analysis to the spectral domain, we recovered the two-photon and first-photon spectra, thereby revealing not only differences in their spectral profiles, as expected for resonant driving, but also a redshifted central wavelength for the first photon. The results are fully consistent with the dressed-state model, and the frequency shift can also be understood as an AC–Stark effect induced by the driving laser. These findings are interesting from a fundamental perspective, since the two-photon spectrum allows one to infer the Rabi frequency, even for an incoherently excited 2LS, and from the application side, as the spectro-temporal correlations—that is, the temporally evolving emission wavelength and line shape—enable selective suppression of the multiphoton component. We demonstrated that spectral filtering can restore high single-photon purities for pulse lengths reaching 15 % of the emitter's decay time, adding another benefit of using phonon-assisted excitation to the list. Importantly, even for pulses as long as 80 ps (which would be uncommonly long for an application), the reported frequency shift is 45 GHz, allowing for filtering using off-the-shelf components. For more typical laser pulses of 20 ps, the shift is expected to double, improving the selectivity of the filtering further.

7.2 Single-photon cryptography

The last work presented in Chapter 6 applies the extended understanding of phonon-assisted excitation and harnesses its advantages to implement a mistrustful quantum-cryptographic protocol in a real-world environment. We developed a primitive, coined *limited-knowledge proofs*, that allows verifying the solution to a hard problem while only partially disclosing it. In this way, problem and

solution can be reused multiple times without compromising the security guarantees. Based on a QBC *i.t.*-secure under the bounded-storage assumption [106, 147], our approach avoids no-signaling constraints, which have previously been used to secure commitments [391–394] but are difficult to enforce [406]. The protocol's practicality is further underlined by its loss tolerance, a noise tolerance of more than 7 %, and its implementability with standard BB84 states. These characteristics highlight its compatibility with existing QKD hardware and infrastructure, which—combined with its reusability property arising from the selective disclosure—signifies its relevance to the outstanding challenge of authenticating QKD links.

While our limited-knowledge proofs can be realized with WCS, we demonstrate that QDs outperform PDSs in all practically relevant regimes. Additionally, our distance projections show that the low multiphoton noise of our QD source achieves almost an order of magnitude more loss tolerance. These results are among the first confirming that SPSs enhance performance in quantum cryptography beyond what is possible with Poissonian photon statistics [44, 46], with particularly large improvements expected for mistrustful settings [45].

7.3 Challenges and perspectives

Phonon-assisted excitation offers numerous advantages for the efficient, practical, and flexible generation of quantum light, and the scheme's success is accompanied by growing interest. One noteworthy application benefiting from the technique–besides single-photon cryptography discussed at length–is the generation of spin-mediated cluster states [57, 61, 256, 417]. These highly versatile resource states—for example, for measurement-based quantum computing [418] and multiparty cryptography [121]—consist of a stream of successively emitted photons that are entangled in polarization. Thus, laser suppression must not rely on polarization filtering, which excludes schemes such as RF or ARP in confocal configurations [419]. In addition to this specific but important use-case, we anticipate that the unique features of phonon-assisted excitation, most notably the AC–Stark-shifted emission frequency and the phonon scattering characteristics governed by the exciton–phonon coupling $J(\omega)$, will benefit SPSs more generally. We therefore expect their relevance to increase as technological advances continue.

For instance, in times when Purcell enhancement accelerates emitter decay times ever more, and multiphoton noise attributed to fabrication imperfections steadily reduces, re-excitation accounts for a significant portion of multiphoton events. However, as the laser pulse length is practically lower-limited by bandwidth requirements arising from laser suppression, alternative approaches to maintain high single-photon purities become necessary, and we believe that the straightforward and efficient frequency filtering introduced in Chapter 5 provides a powerful solution. In combination with narrowband cavities, such as micro pillars, the scheme's benefits multiply as the AC–Stark effect shifts the emission line out of the cavity resonance, canceling the Purcell enhancement and effectively suppressing early photon emission

(see Sec. 5.2).

Extending the analysis to consider not only the cavity-enhanced, time-dependent photon rates but also the time-dependent phonon scattering rates (see Sec. 2.2.4.2) promises to create additional opportunities to enhance the emission process. Ultimately, one could manipulate $J(\omega)$ to tailor the light–matter interaction so that, at no point during the laser pulse, the phonon and photon decay rates are simultaneously high, thereby preventing re-excitation altogether.

Given the ubiquity of phonons and the success of phonon-assisted driving schemes, exploring the potential of these techniques for other quantum emitters represents another compelling pathway [420–422].

Shifting the focus to quantum-secured communication, recent improvements in QD performance experimentally manifested the advantage of QDs over PDSs in QKD [44]. Widening the scope to beyond-QKD applications [121] showed that these profit even more from the low multiphoton noise: they enable secure protocols in parameter regimes inaccessible to PDSs [46] and also provide improved scalings (see Chapter 6). Nonetheless, building large networks to securely exchange quantum information [69, 73, 354, 423, 424], as envisioned for distributed or blind quantum computing [59, 425, 426] and multiparty computations [121], still poses many more challenges, and most likely, such a quantum internet will be hybrid, playing on the advantages of both PDSs and SPSs. Combining not only various photonic sources but also creating efficient light–matter interfaces is another major undertaking, requiring much work on interconnectivity [427]. Here, QDs offer rich opportunities; coupling photonic and matter qubits [72, 88, 416] and enabling deterministic photon–photon interactions via strong engineered nonlinearities [62–64, 66].

However, other architectures achieve better scalability, qubit connectivity, or two-qubit gate fidelities [9, 428], incentivizing further hardware diversification. As photons are widely regarded as ideal carriers of quantum information over long distances, the ability to efficiently interface single photons with one another or with matter qubits is essential. To this end, FT limited light, quantum frequency conversion [429–431], and bandwidth matching [432, 433] across multiple orders of magnitude [434] are necessary. While many QD systems routinely achieve the first [55, 56, 435], the latter remain challenging, and although impressive advances have been made, most implementations still fall short of the target performance.

List of Acronyms

ARP	Adiabatic rapid passage	57
BS	beam splitter	37
CBG	circular Bragg grating	4
CW	continuous wave	49
DBR	distributed Bragg reflector	33
DOP	degree of polarization	129
DOS	density of states	60
DV	discrete-variable	15
EOM	electro-optic modulator	125
FPGA	field-programmable gate array	72
FSR	free spectral range	105
FSS	fine structure splitting	35
FT	Fourier-transform	38
FWHM	full width at half maximum	38
HBT	Hanbury Brown and Twiss experiment	37
HOM	Hong–Ou–Mandel	39
HWP	half-wave plate	129
HV	high-voltage	125
LA	longitudinal acoustic	
LO	longitudinal optical	46
MBE	molecular-beam epitaxy	32
MMB	metamorphic buffer layer	33
MOVPE	metalorganic vapor-phase epitaxy	32
PBS	polarizing beam-splitter	105
PCC	Pockels Cell controller	131
PDS	Poisson-distributed source	2
PhC WG	photonic crystal waveguide	60
PNC	photon-number coherence	17
PNS	photon-number splitting	19

QBC quantum bit commitment . . . . .
QBER quantum bit error rate 124
QD quantum dot . . . . .
QKD quantum key distribution . . . . .
QWP quarter-wave plate 129
RF resonance fluorescence 48
RFoF radio frequency over fiber 125
SK secure key 19
SKR secure key rate 84
SMF single-mode fiber 23
SNR signal-to-noise ratio 4
SNSPD superconducting nanowire single-photon detector 72
SPDC spontaneous parametric down-conversion 2
SPS single-photon source 2
SRF system response function 72
STIRAP stimulated Raman adiabatic passage 57
sTPE stimulated two-photon excitation 58
SUPER Swing-UP of the quantum EmitteR population 48
TA transversal acoustic 46
TCPC time-correlated photon counting 63
CNF conjunctive normal form 133
3-SAT 3-satisfiabilty 133
3LS three-level system 38
2LS two-level system 38
TPE resonant two-photon excitation 57
TPI two-photon interference 39
USD unambiguous state discrimination 19
VBG volume Bragg grating 71
WCS weak coherent state . . . . .

Bibliography

[1] M. Vyvlecka, L. Jehle, C. Nawrath, F. Giorgino, M. Bozzio, R. Sittig, M. Jetter, S. L. Portalupi, P. Michler, and P. Walther. *Robust excitation of C-band quantum dots for quantum communication*. *Applied Physics Letters* 123(17): p. 174001. 2023.

[2] L. Jehle, L. M. Hansen, P. I. Sund, T. W. Sandø, R. Joos, M. Jetter, S. L. Portalupi, M. Bozzio, P. Michler, and P. Walther. *Asymmetric two-photon response of an incoherently driven quantum emitter*, 2025.

[3] S. Wiesner. *Conjugate coding*. *ACM Sigact News* 15(1): pp. 78–88. 1983.

[4] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala. *Evidence for the utility of quantum computing before fault tolerance*. *Nature* 618(7965): pp. 500–505. 2023.

[5] H. Aghaee Rad, T. Ainsworth, R. N. Alexander, B. Altieri, M. F. Askarani, R. Baby, L. Banchi, B. Q. Baragiola, J. E. Bourassa, R. S. Chadwick, I. Charania, H. Chen, M. J. Collins, P. Contu, N. D'Arcy, G. Dauphinais, R. De Prins, D. Deschenes, I. Di Luch, S. Duque, P. Edke, S. E. Fayer, S. Ferracin, H. Ferretti, J. Gefaell, S. Glancy, C. González-Arciniegas, T. Grainge, Z. Han, J. Hastrup, L. G. Helt, T. Hillmann, J. Hundal, S. Izumi, T. Jaeken, M. Jonas, S. Kocsis, I. Krasnokutska, M. V. Larsen, P. Laskowski, F. Laudenbach, J. Lavoie, M. Li, E. Lomonte, C. E. Lopetegui, B. Luey, A. P. Lund, C. Ma, L. S. Madsen, D. H. Mahler, L. Mantilla Calderón, M. Menotti, F. M. Miatto, B. Morrison, P. J. Nadkarni, T. Nakamura, L. Neuhaus, Z. Niu, R. Noro, K. Papirov, A. Pesah, D. S. Phillips, W. N. Plick, T. Rogalsky, F. Rortais, J. Sabines-Chesterking, S. Safavi-Bayat, E. Sazhaev, M. Seymour, K. Rezaei Shad, M. Silverman, S. A. Srinivasan, M. Stephan, Q. Y. Tang, J. F. Tasker, Y. S. Teo, R. B. Then, J. E. Tremblay, I. Tzitrin, V. D. Vaidya, M. Vasmer, Z. Vernon, L. F. S. S. M. Villalobos, B. W. Walshe, R. Weil, X. Xin, X. Yan, Y. Yao, M. Zamani Abnili, and Y. Zhang. *Scaling and networking a modular photonic quantum computer*. *Nature* 638(8052): pp. 912–919. 2025.

[6] A. D. King, A. Nocera, M. M. Rams, J. Dziarmaga, R. Wiersema, W. Bernoudy, J. Raymond, N. Kaushal, N. Heinsdorf, R. Harris, K. Boothby, F. Altomare, M. Asad, A. J. Berkley, M. Boschnak, K. Chern, H. Christiani, S. Cibere, J. Connor, M. H. Dehn, R. Deshpande, S. Ejtemaee, P. Farre, K. Hamer, E. Hoskinson, S. Huang, M. W. Johnson, S. Kortas, E. Ladizinsky, T. Lanting, T. Lai, R. Li, A. J. R. MacDonald, G. Marsden, C. C. McGeoch, R. Molavi, T. Oh, R. Neufeld, M. Norouzpour, J. Pasvolsky, P. Poitras, G. Poulin-Lamarre, T. Prescott, M. Reis, C. Rich, M. Samani, B. Sheldan, A. Smirnov, E. Sterpka, B. T. Clavera, N. Tsai, M. Volkmann, A. M. Whiticar, J. D. Whittaker, W. Wilkinson, J. Yao, T. J. Yi, A. W. Sandvik, G. Alvarez, R. G. Melko, J. Carrasquilla, M. Franz, and M. H. Amin. *Beyond-classical computation in quantum simulation*. *Science* 388(6743): pp. 199–204. 2025.

[7] D. A. Abanin et al. *Observation of constructive interference at the edge of quantum ergodicity*. *Nature* 646(8086): pp. 825–830. 2025.

[8] K. Alexander et al. *A manufacturable platform for photonic quantum computing*. *Nature* 641(8064): pp. 876–883. 2025.

[9] N.-C. Chiu, E. C. Trapp, J. Guo, M. H. Abobeih, L. M. Stewart, S. Hollerith, P. L. Stroganov, M. Kalinowski, A. A. Geim, S. J. Evered, S. H. Li, X. Lyu, L. M. Peters, D. Bluvstein, T. T. Wang, M. Greiner, V. Vuletić, and M. D. Lukin. *Continuous operation of a coherent 3, 000-qubit system*. *Nature* 646(8087): pp. 1075–1080. 2025.

[10] É. Gouzien and N. Sangouard. *Factoring 2048-bit RSA Integers in 177 Days with 13 436 Qubits and a Multimode Memory*. *Phys. Rev. Lett.* 127: p. 140503. 2021.

[11] C. Hong, Z. Pei, Q. Wang, S. Yang, J. Yu, and C. Wang. *Quantum attack on RSA by D-Wave Advantage: a first break of 80-bit RSA*. *Science China Information Sciences* 68(2): 2025.

[12] W. K. Wootters and W. H. Zurek. *A single quantum cannot be cloned*. *Nature* 1982(299): pp. 802–803. 1982.

[13] 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(7861): pp. 37–40. 2021.

[14] M. Pompili, S. L. N. Hermans, S. Baier, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, S. Wehner, and R. Hanson. *Realization of a multinode quantum network of remote solid-state qubits*. *Science* 372(6539): pp. 259–264. 2021.

[15] M. Guo, S. Liu, W. Sun, M. Ren, F. Wang, and M. Zhong. *Rare-earth quantum memories: The experimental status quo*. *Frontiers of Physics* 18(2): 2023.

[16] E. Bersin, M. Sutula, Y. Q. Huan, A. Suleymanzade, D. R. Assumpcao, Y.-C. Wei, P.-J. Stas, C. M. Knaut, E. N. Knall, C. Langrock, N. Sinclair, R. Murphy, R. Riedinger, M. Yeh, C. Xin, S. Bandyopadhyay, D. D. Sukachev, B. Machielse, D. S. Levonian, M. K. Bhaskar, S. Hamilton, H. Park, M. Lon čar, M. M. Fejer, P. B. Dixon, D. R. Englund, and M. D. Lukin. *Telecom Networking with a Diamond Quantum Memory*. *PRX Quantum* 5: p. 010303. 2024.

[17] H. Mamann, T. Nieddu, F. Hoffet, M. Bozzio, F. G. de Loubresse, I. Kerenidis, E. Diamanti, A. Urvoy, and J. Laurat. *Quantum cryptography integrating an optical quantum memory*. *Science Advances* 11(38): eadx3223. 2025.

[18] B. M. Terhal. *Quantum error correction for quantum memories*. *Rev. Mod. Phys.* 87: pp. 307–346. 2015.

[19] A. Einstein, B. Podolsky, and N. Rosen. *Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?* *Phys. Rev.* 47: pp. 777–780. 1935.

[20] J. S. Bell. *Physics long island city. NY* 1(195): p. 1. 1964.

[21] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. *Proposed Experiment to Test Local Hidden-Variable Theories*. *Phys. Rev. Lett.* 23: pp. 880–884. 1969.

[22] S. J. Freedman and J. F. Clauser. *Experimental Test of Local Hidden-Variable Theories*. *Phys. Rev. Lett.* 28: pp. 938–941. 1972.

[23] A. Aspect, P. Grangier, and G. Roger. *Experimental Tests of Realistic Local Theories via Bell's Theorem*. *Phys. Rev. Lett.* 47: pp. 460–463. 1981.

[24] D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eibl, H. Weinfurter, and A. Zeilinger. *Experimental quantum teleportation*. *Nature* 390(6660): pp. 575–579. 1997.

[25] S. L. N. Hermans, M. Pompili, H. K. C. Beukers, S. Baier, J. Borregaard, and R. Hanson. *Qubit teleportation between non-neighbouring nodes in a quantum network*. *Nature* 605(7911): pp. 663–668. 2022.

[26] J.-W. Pan, D. Bouwmeester, H. Weinfurter, and A. Zeilinger. *Experimental Entanglement Swapping: Entangling Photons That Never Interacted*. *Phys. Rev. Lett.* 80: pp. 3891–3894. 1998.

[27] D. Bouwmeester, J.-W. Pan, M. Daniell, H. Weinfurter, and A. Zeilinger. *Observation of Three-Photon Greenberger-Horne-Zeilinger Entanglement*. *Phys. Rev. Lett.* 82: pp. 1345–1349. 1999.

[28] M. Eibl, N. Kiesel, M. Bourennane, C. Kurtsiefer, and H. Weinfurter. *Experimental Realization of a Three-Qubit Entangled W State*. *Phys. Rev. Lett.* 92: p. 077901. 2004.

[29] J.-W. Pan, D. Bouwmeester, M. Daniell, H. Weinfurter, and A. Zeilinger. *Experimental test of quantum nonlocality in three-photon Greenberger–Horne–Zeilinger entanglement*. *Nature* 403(6769): pp. 515–519. 2000.

[30] C. H. Bennett and G. Brassard. *Quantum cryptography: Public key distribution and coin tossing*. *Theoretical Computer Science* 560: pp. 7–11. 1984.

[31] A. Pappa, A. Chailloux, E. Diamanti, and I. Kerenidis. *Practical quantum coin flipping*. *Phys. Rev. A* 84: p. 052305. 2011.

[32] Y. Liu, Y. Cao, M. Curty, S.-K. Liao, J. Wang, K. Cui, Y.-H. Li, Z.-H. Lin, Q.-C. Sun, D.-D. Li, H.-F. Zhang, Y. Zhao, T.-Y. Chen, C.-Z. Peng, Q. Zhang, A. Cabello, and J.-W. Pan. *Experimental Unconditionally Secure Bit Commitment*. *Phys. Rev. Lett.* 112: p. 010504. 2014.

[33] A. Rubenok, J. A. Slater, P. Chan, I. Lucio-Martinez, and W. Tittel. *Real-World Two-Photon Interference and Proof-of-Principle Quantum Key Distribution Immune to Detector Attacks*. *Phys. Rev. Lett.* 111: p. 130501. 2013.

[34] Y. Liu, T.-Y. Chen, L.-J. Wang, H. Liang, G.-L. Shentu, J. Wang, K. Cui, H.-L. Yin, N.-L. Liu, L. Li, X. Ma, J. S. Pelc, M. M. Fejer, C.-Z. Peng, Q. Zhang, and J.-W. Pan. *Experimental Measurement-Device-Independent Quantum Key Distribution*. *Phys. Rev. Lett.* 111: p. 130502. 2013.

[35] F. Xu, X. Ma, Q. Zhang, H.-K. Lo, and J.-W. Pan. *Secure quantum key distribution with realistic devices*. *Reviews of Modern Physics* 92(2): 2020.

[36] V. Zapatero, T. van Leent, R. Arnon-Friedman, W.-Z. Liu, Q. Zhang, H. Weinfurter, and M. Curty. *Advances in device-independent quantum key distribution*. *npj Quantum Information* 9(1): 2023.

[37] Z.-S. Yuan, Y.-A. Chen, B. Zhao, S. Chen, J. Schmiedmayer, and J.-W. Pan. *Experimental demonstration of a BDCZ quantum repeater node*. *Nature* 454(7208): pp. 1098–1101. 2008.

[38] Z.-D. Li, R. Zhang, X.-F. Yin, L.-Z. Liu, Y. Hu, Y.-Q. Fang, Y.-Y. Fei, X. Jiang, J. Zhang, L. Li, N.-L. Liu, F. Xu, Y.-A. Chen, and J.-W. Pan. *Experimental quantum repeater without quantum memory*. *Nature Photonics* 13(9): pp. 644–648. 2019.

[39] G. Brassard, N. Lütkenhaus, T. Mor, and B. C. Sanders. *Limitations on Practical Quantum Cryptography*. *Phys. Rev. Lett.* 85: pp. 1330–1333. 2000.

[40] W.-Y. Hwang. *Quantum Key Distribution with High Loss: Toward Global Secure Communication*. *Phys. Rev. Lett.* 91: p. 057901. 2003.

[41] H.-K. Lo, X. Ma, and K. Chen. *Decoy State Quantum Key Distribution*. *Phys. Rev. Lett.* 94: p. 230504. 2005.

[42] E. Waks, C. Santori, and Y. Yamamoto. *Security aspects of quantum key distribution with sub-Poisson light*. *Physical Review A* 66(4): 2002.

[43] A. Beveratos, R. Brouri, T. Gacoin, A. Villing, J.-P. Poizat, and P. Grangier. *Single Photon Quantum Cryptography*. *Phys. Rev. Lett.* 89: p. 187901. 2002.

[44] Y. Zhang, X. Ding, Y. Li, L. Zhang, Y.-P. Guo, G.-Q. Wang, Z. Ning, M.-C. Xu, R.-Z. Liu, J.-Y. Zhao, G.-Y. Zou, H. Wang, Y. Cao, Y.-M. He, C.-Z. Peng, Y.-H. Huo, S.-K. Liao, C.-Y. Lu, F. Xu, and J.-W. Pan. *Experimental Single-Photon Quantum Key Distribution Surpassing the Fundamental Weak Coherent-State Rate Limit*. *Phys. Rev. Lett.* 134: p. 210801. 2025.

[45] M. Bozzio, M. Vyvlecka, M. Cosacchi, C. Nawrath, T. Seidelmann, J. C. Loredo, S. L. Portalupi, V. M. Axt, P. Michler, and P. Walther. *Enhancing quantum cryptography with quantum dot single-photon sources*. *npj Quantum Information* 8(1): 2022.

[46] D. A. Vajner, K. Kaymazlar, F. Drauschke, L. Rickert, M. von Helversen, H. Liu, S. Li, H. Ni, Z. Niu, A. Pappa, and T. Heindel. *Single-photon advantage in quantum cryptography beyond QKD*. *Nature Communications* 17(1): 2026.

[47] P. Michler, A. Kiraz, C. Becher, W. V. Schoenfeld, P. M. Petroff, L. Zhang, E. Hu, and A. Imamoglu. *A Quantum Dot Single-Photon Turnstile Device*. *Science* 290(5500): pp. 2282–2285. 2000.

[48] C. Kurtsiefer, S. Mayer, P. Zarda, and H. Weinfurter. *Stable Solid-State Source of Single Photons*. *Phys. Rev. Lett.* 85: pp. 290–293. 2000.

[49] C. Santori, M. Pelton, G. Solomon, Y. Dale, and Y. Yamamoto. *Triggered Single Photons from a Quantum Dot*. *Phys. Rev. Lett.* 86: pp. 1502–1505. 2001.

[50] C. Santori, D. Fattal, J. Vučković, G. S. Solomon, and Y. Yamamoto. *Indistinguishable photons from a single-photon device*. *Nature* 419(6907): pp. 594–597. 2002.

[51] M. Pelton, C. Santori, J. Vuc kovi ć, B. Zhang, G. S. Solomon, J. Plant, and Y. Yamamoto. *Efficient Source of Single Photons: A Single Quantum Dot in a Micropost Microcavity*. *Phys. Rev. Lett.* 89: p. 233602. 2002.

[52] A. Kuhn, M. Hennrich, and G. Rempe. *Deterministic Single-Photon Source for Distributed Quantum Networking*. *Phys. Rev. Lett.* 89: p. 067901. 2002.

[53] A. Zrenner, E. Beham, S. Stufler, F. Findeis, M. Bichler, and G. Abstreiter. *Coherent properties of a two-level system based on a quantum-dot photodiode*. *Nature* 418(6898): pp. 612–614. 2002.

[54] R. M. Stevenson, R. J. Young, P. Atkinson, K. Cooper, D. A. Ritchie, and A. J. Shields. *A semiconductor source of triggered entangled photon pairs*. *Nature* 439(7073): pp. 179–182. 2006.

[55] X. Ding, Y.-P. Guo, M.-C. Xu, R.-Z. Liu, G.-Y. Zou, J.-Y. Zhao, Z.-X. Ge, Q.-H. Zhang, H.-L. Liu, L.-J. Wang, M.-C. Chen, H. Wang, Y.-M. He, Y.-H. Huo, C.-Y. Lu, and J.-W. Pan. *High-efficiency single-photon source above the loss-tolerant threshold for efficient linear optical quantum computing*. *Nature Photonics* 19(4): pp. 387–391. 2025.

[56] N. Tomm, A. Javadi, N. O. Antoniadis, D. Najer, M. C. Löbl, A. R. Korsch, R. Schott, S. R. Valentin, A. D. Wieck, A. Ludwig, and R. J. Warburton. *A bright and fast source of coherent single photons*. *Nature nanotechnology* 16(4): pp. 399–403. 2021.

[57] N. Coste, D. A. Fioretto, N. Belabas, S. C. Wein, P. Hilaire, R. Frantzeskakis, M. Gundin, B. Goes, N. Somaschi, M. Morassi, A. Lemaître, I. Sagnes, A. Harouri, S. E. Economou, A. Auffeves, O. Krebs, L. Lanco, and P. Senellart. *High-rate entanglement between a semiconductor spin and indistinguishable photons*. *Nature Photonics* 17(7): pp. 582–587. 2023.

[58] H. Cao, L. M. Hansen, F. Giorgino, L. Carosini, P. Zahálka, F. Zilk, J. C. Loredo, and P. Walther. *Photonic Source of Heralded Greenberger-Horne-Zeilinger States*. *Phys. Rev. Lett.* 132: p. 130604. 2024.

[59] N. Maring, A. Fyrillas, M. Pont, E. Ivanov, P. Stepanov, N. Margaria, W. Hease, A. Pishchagin, A. Lemaître, I. Sagnes, T. H. Au, S. Boissier, E. Bertasi, A. Baert, M. Valdivia, M. Billard, O. Acar, A. Brieussel, R. Mezher, S. C. Wein, A. Salavrakos, P. Sinnott, D. A. Fioretto, P.-E. Emeriau, N. Belabas, S. Mansfield, P. Senellart, J. Senellart, and N. Somaschi. *A versatile single-photon-based quantum computing platform*. *Nature Photonics* 18(6): pp. 603–609. 2024.

[60] Y. Meng, M. L. Chan, R. B. Nielsen, M. H. Appel, Z. Liu, Y. Wang, N. Bart, A. D. Wieck, A. Ludwig, L. Midolo, A. Tiranov, A. S. Sørensen, and P. Lodahl. *Deterministic photon source of genuine three-qubit entanglement*. *Nature communications* 15(1): p. 7774. 2024.

[61] P. Laccotripes, T. Müller, R. M. Stevenson, J. Skiba-Szymanska, D. A. Ritchie, and A. J. Shields. *Spin-photon entanglement with direct photon emission in the telecom C-band*. *Nature communications* 15(1): p. 9740. 2024.

[62] P. Türschmann, H. L. Jeannic, S. F. Simonsen, H. R. Haakh, S. Götzinger, V. Sandoghdar, P. Lodahl, and N. Rotenberg. *Coherent nonlinear optics of quantum emitters in nanophotonic waveguides*. *Nanophotonics* 8(10): pp. 1641–1657. 2019.

[63] H. Le Jeannic, A. Tiranov, J. Carolan, T. Ramos, Y. Wang, M. H. Appel, S. Scholz, A. D. Wieck, A. Ludwig, N. Rotenberg, L. Midolo, J. J. García-Ripoll, A. S. Sørensen, and P. Lodahl. *Dynamical photon–photon interaction mediated by a quantum emitter*. *Nature Physics* 18(10): pp. 1191–1195. 2022.

[64] N. Tomm, S. Mahmoodian, N. O. Antoniadis, R. Schott, S. R. Valentin, A. D. Wieck, A. Ludwig, A. Javadi, and R. J. Warburton. *Photon bound state dynamics from a single artificial atom*. *Nature physics* 19(6): pp. 857–862. 2023.

[65] A. Tiranov, V. Angelopoulou, C. J. van Diepen, B. Schrinski, O. A. D. Sandberg, Y. Wang, L. Midolo, S. Scholz, A. D. Wieck, A. Ludwig, A. S. Sørensen, and P. Lodahl. *Collective super- and subradiant dynamics between distant optical quantum emitters*. *Science (New York, N.Y.)* 379(6630): pp. 389–393. 2023.

[66] A. González-Tudela, A. Reiserer, J. J. García-Ripoll, and F. J. García-Vidal. *Light–matter interactions in quantum nanophotonic devices*. *Nature Reviews Physics* 6(3): pp. 166–179. 2024.

[67] R. J. Warburton. *Single spins in self-assembled quantum dots*. *Nature Materials* 12(6): pp. 483–493. 2013.

[68] P. Senellart, G. Solomon, and A. White. *High-performance semiconductor quantum-dot single-photon sources*. *Nature nanotechnology* 12(11): pp. 1026–1039. 2017.

[69] P. Lodahl. *Quantum-dot based photonic quantum networks*. *Quantum Science and Technology* 3(1): p. 013001. 2018.

[70] S. Hepp, M. Jetter, S. L. Portalupi, and P. Michler. *Semiconductor Quantum Dots for Integrated Quantum Photonics*. *Advanced Quantum Technologies* 2(9): p. 1900020. 2019.

[71] R. Trivedi, K. A. Fischer, J. Vučković, and K. Müller. *Generation of Non-Classical Light Using Semiconductor Quantum Dots*. *Advanced Quantum Technologies* 3(1): p. 1900007. 2020.

[72] R. Uppu, L. Midolo, X. Zhou, J. Carolan, and P. Lodahl. *Quantum-dot-based deterministic photon-emitter interfaces for scalable photonic quantum technology*. *Nature nanotechnology* 16(12): pp. 1308–1317. 2021.

[73] Y. Yu, S. Liu, C.-M. Lee, P. Michler, S. Reitzenstein, K. Srinivasan, E. Waks, and J. Liu. *Telecom-band quantum dot technologies for long-distance quantum networks*. *Nature nanotechnology* 18(12): pp. 1389–1400. 2023.

[74] J. H. Weber, B. Kambs, J. Kettler, S. Kern, J. Maisch, H. Vural, M. Jetter, S. L. Portalupi, C. Becher, and P. Michler. *Two-photon interference in the telecom C-band after frequency conversion of photons from remote quantum emitters*. *Nature Nanotechnology* 14(1): pp. 23–26. 2018.

[75] M. Pont, G. Corrielli, A. Fyrillas, I. Agresti, G. Carvacho, N. Maring, P.-E. Emeriau, F. Ceccarelli, R. Albiero, P. H. Dias Ferreira, N. Somaschi, J. Senellart, I. Sagnes, M. Morassi, A. Lemaître, P. Senellart, F. Sciarrino, M. Liscidini, N. Belabas, and R. Osellame. *High-fidelity four-photon GHZ states on chip*. *npj Quantum Information* 10(1): 2024.

[76] T. Strobel, M. Vyvlecka, I. Neureuther, T. Bauer, M. Schäfer, S. Kazmaier, N. L. Sharma, R. Joos, J. H. Weber, C. Nawrath, W. Nie, G. Bhayani, C. Hopfmann, C. Becher, P. Michler, and S. L. Portalupi. *Telecom-wavelength quantum teleportation using frequency-converted photons from remote quantum dots*. *Nature communications* 16(1): p. 10027. 2025.

[77] F. B. Barnes, R. G. Pousa, C. L. Morrison, Z. X. Koong, J. Ho, F. Graffitti, J. Jeffers, D. K. L. Oi, B. D. Gerardot, and A. Fedrizzi. *Decoy-state quantum key distribution over 227 km with a frequency-converted telecom single-photon source*, 2025.

[78] Cornelius Nawrath, Raphael Joos, Sascha Kolatschek, Stephanie Bauer, Pascal Pruy, Florian Hornung, Julius Fischer, Jiasheng Huang, Ponraj Vijayan, Robert Sittig, Michael Jetter, Simone Luca Portalupi, and Peter Michler. *Bright Source of Purcell-Enhanced, Triggered, Single Photons in the Telecom C-Band*. *Advanced Quantum Technologies* 6(11): p. 2300111. 2023.

[79] L. Wells, T. Müller, R. M. Stevenson, J. Skiba-Szymanska, D. A. Ritchie, and A. J. Shields. *Coherent light scattering from a telecom C-band quantum dot*. *Nature communications* 14(1): p. 8371. 2023.

[80] R. Joos, S. Bauer, C. Rupp, S. Kolatschek, W. Fischer, C. Nawrath, P. Vijayan, R. Sittig, M. Jetter, S. L. Portalupi, and P. Michler. *Coherently and Incoherently Pumped Telecom C-Band Single-Photon Source with High Brightness and Indistinguishability*. *Nano letters* 24(28): pp. 8626–8633. 2024.

[81] R. J. Warburton. *A Self-assembled Quantum Dot as Single Photon Source and Spin Qubit: Charge Noise and Spin Noise*. In: P. Michler, *Quantum Dots for Quantum Information Technologies*. Springer International Publishing, Cham. Pp. 287–323. 2017.

[82] H. Drexler, D. Leonard, W. Hansen, J. P. Kotthaus, and P. M. Petroff. *Spectroscopy of Quantum Levels in Charge-Tunable InGaAs Quantum Dots*. *Phys. Rev. Lett.* 73: pp. 2252–2255. 1994.

[83] R. J. Warburton, B. T. Miller, C. S. Dürr, C. Bödefeld, K. Karrai, J. P. Kotthaus, G. Medeiros-Ribeiro, P. M. Petroff, and S. Huant. *Coulomb interactions in small charge-tunable quantum dots: A simple model*. *Phys. Rev. B* 58: pp. 16221–16231. 1998.

[84] C. Schimpf, S. Manna, S. F. C. da Silva, M. Aigner, and A. Rastelli. *Entanglement-based quantum key distribution with a blinking-free quantum dot operated at a temperature up to 20 K*. *Advanced Photonics* 3(6): p. 065001. 2021.

[85] K. J. Vahala. *Optical microcavities*. *Nature* 2003(424): pp. 839–846. 2003.

[86] H. Wang, Y.-M. He, T.-H. Chung, H. Hu, Y. Yu, S. Chen, X. Ding, M.-C. Chen, J. Qin, X. Yang, R.-Z. Liu, Z.-C. Duan, J.-P. Li, S. Gerhardt, K. Winkler, J. Jurkat, L.-J. Wang, N. Gregersen, Y.-H. Huo, Q. Dai, S. Yu, S. Höfling, C.-Y. Lu, and J.-W. Pan. *Towards optimal single-photon sources from polarized microcavities*. *Nature Photonics* 13(11): pp. 770–775. 2019.

[87] E. M. Purcell, H. C. Torrey, and R. V. Pound. *Resonance Absorption by Nuclear Magnetic Moments in a Solid*. *Phys. Rev.* 69: pp. 37–38. 1946.

[88] L. Zaporski, N. Shofer, J. H. Bodey, S. Manna, G. Gillard, M. H. Appel, C. Schimpf, S. F. Covre da Silva, J. Jarman, G. Delamare, G. Park, U. Haeusler, E. A. Chekhovich, A. Rastelli, D. A. Gangloff, M. Atatüre, and C. Le Gall. *Ideal refocusing of an optically active spin qubit under strong hyperfine interactions*. *Nature Nanotechnology* 18(3): pp. 257–263. 2023.

[89] M. H. Appel, A. Ghorbal, N. Shofer, L. Zaporski, S. Manna, S. F. C. da Silva, U. Haeusler, C. Le Gall, A. Rastelli, D. A. Gangloff, and M. Atatüre. *A many-body quantum register for a spin qubit*. *Nature Physics* 21(3): pp. 368–373. 2025.

[90] I. Schwartz, D. Cogan, E. R. Schmidgall, Y. Don, L. Gantz, O. Kenneth, N. H. Lindner, and D. Gershoni. *Deterministic generation of a cluster state of entangled photons*. *Science (New York, N.Y.)* 354(6311): pp. 434–437. 2016.

[91] H.-K. Lo and J. Preskill. *Security of quantum key distribution using weak coherent states with nonrandom phases*. *Quantum Info. Comput.* 7(5): pp. 431–458. 2007.

[92] P. Michler and S. L. Portalupi. *Semiconductor Quantum Light Sources: Fundamentals, Technologies and Devices*. De Gruyter, Berlin, Boston. 2024.

[93] S. Lüker and D. E. Reiter. *A review on optical excitation of semiconductor quantum dots under the influence of phonons*. *Semiconductor Science and Technology* 34(6): p. 063002. 2019.

[94] F. Kappe, Y. Karli, G. Wilbur, R. G. Krämer, S. Ghosh, R. Schwarz, M. Kaiser, T. K. Bracht, D. E. Reiter, S. Nolte, K. C. Hall, G. Weihs, and V. Remesh. *Chirped Pulses Meet Quantum Dots: Innovations, Challenges, and Future Perspectives*. *Advanced Quantum Technologies*: 2024.

[95] M. Cosacchi, F. Ungar, M. Cygorek, A. Vagov, and V. M. Axt. *Emission-Frequency Separated High Quality Single-Photon Sources Enabled by Phonons*. *Physical review letters* 123(1): p. 017403. 2019.

[96] K. A. Fischer, L. Hanschke, J. Wierzbowski, T. Simmet, C. Dory, J. J. Finley, J. Vučković, and K. Müller. *Signatures of two-photon pulses from a quantum two-level system*. *Nature Physics* 13(7): pp. 649–654. 2017.

[97] L. Hanschke, K. A. Fischer, S. Appel, D. Lukin, J. Wierzbowski, S. Sun, R. Trivedi, J. Vučković, J. J. Finley, and K. Müller. *Quantum dot single-photon sources with ultra-low multi-photon probability*. *npj Quantum Information* 4(1): 2018.

[98] G. S. Vernam. *Cipher Printing Telegraph Systems For Secret Wire and Radio Telegraphic Communications*. *Transactions of the American Institute of Electrical Engineers* XLV: pp. 295–301. 1926.

[99] R. L. Rivest, A. Shamir, and L. Adleman. *A method for obtaining digital signatures and public-key cryptosystems*. *Commun. ACM* 21(2): pp. 120–126. 1978.

[100] P. W. Shor. *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*. *SIAM Journal on Computing* 26(5): pp. 1484–1509. 1997.

[101] W. Diffie and M. Hellman. *New directions in cryptography*. *IEEE Transactions on Information Theory* 22(6): pp. 644–654. 1976.

[102] Y. Liu, W.-J. Zhang, C. Jiang, J.-P. Chen, C. Zhang, W.-X. Pan, D. Ma, H. Dong, J.-M. Xiong, C.-J. Zhang, H. Li, R.-C. Wang, J. Wu, T.-Y. Chen, L. You, X.-B. Wang, Q. Zhang, and J.-W. Pan. *Experimental Twin-Field Quantum Key Distribution over 1000 km Fiber Distance*. *Phys. Rev. Lett.* 130: p. 210801. 2023.

[103] J. Yin, Y.-H. Li, S.-K. Liao, M. Yang, Y. Cao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, S.-L. Li, R. Shu, Y.-M. Huang, L. Deng, L. Li, Q. Zhang, N.-L. Liu, Y.-A. Chen, C.-Y. Lu, X.-B. Wang, F. Xu, J.-Y. Wang, C.-Z. Peng, A. K. Ekert, and J.-W. Pan. *Entanglement-based secure quantum cryptography over 1,120 kilometres*. *Nature* 582(7813): pp. 501–505. 2020.

[104] H.-K. Lo and H. F. Chau. *Is Quantum Bit Commitment Really Possible?* *Phys. Rev. Lett.* 78: pp. 3410–3413. 1997.

[105] H.-K. Lo and H. Chau. *Why quantum bit commitment and ideal quantum coin tossing are impossible*. *Physica D: Nonlinear Phenomena* 120(1): pp. 177–187. 1998.

[106] R. Konig, S. Wehner, and J. Wullschleger. *Unconditional Security From Noisy Quantum Storage*. *IEEE Transactions on Information Theory* 58(3): pp. 1962–1984. 2012.

[107] A. Kent. *Unconditionally Secure Bit Commitment by Transmitting Measurement Outcomes*. *Phys. Rev. Lett.* 109: p. 130501. 2012.

[108] J. Kaniewski, M. Tomamichel, E. Hänggi, and S. Wehner. *Secure Bit Commitment From Relativistic Constraints*. *IEEE Transactions on Information Theory* 59(7): pp. 4687–4699. 2013.

[109] S. Storz, J. Schär, A. Kulikov, P. Magnard, P. Kurpiers, J. Lütolf, T. Walter, A. Copetudo, K. Reuer, A. Akin, J.-C. Besse, M. Gabureac, G. J. Norris, A. Rosario, F. Martin, J. Martinez, W. Amaya, M. W. Mitchell, C. Abellan, J.-D. Bancal, N. Sangouard, B. Royer, A. Blais, and A. Wallraff. *Loophole-free Bell inequality violation with superconducting circuits*. *Nature* 617(7960): pp. 265–270. 2023.

[110] J. Altepeter, E. Jeffrey, and P. Kwiat. *Photonic State Tomography*. In: ed. by P. Berman and C. Lin. Vol. 52. Advances In Atomic, Molecular, and Optical Physics. Academic Press, 2005, pp. 105–159.

[111] V. Bu žek and M. Hillery. *Quantum copying: Beyond the no-cloning theorem*. *Phys. Rev. A* 54: pp. 1844–1852. 1996.

[112] M. Fox. *Quantum optics: an introduction*. Oxford University Press., 2006.

[113] R. J. Glauber. *Coherent and Incoherent States of the Radiation Field*. *Phys. Rev.* 131: pp. 2766–2788. 1963.

[114] R. J. Glauber. *The Quantum Theory of Optical Coherence*. *Phys. Rev.* 130: pp. 2529–2539. 1963.

[115] U. M. Titulaer and R. J. Glauber. *Correlation Functions for Coherent Fields*. *Phys. Rev.* 140: B676–B682. 1965.

[116] C. Couteau, S. Barz, T. Durt, T. Gerrits, J. Huwer, R. Prevedel, J. Rarity, A. Shields, and G. Weihs. *Applications of single photons in quantum metrology, biology and the foundations of quantum physics*. *Nature Reviews Physics* 5(6): pp. 354–363. 2023.

[117] Q. Li, K. Orcutt, R. L. Cook, J. Sabines-Chesterking, A. L. Tong, G. S. Schlau-Cohen, X. Zhang, G. R. Fleming, and K. B. Whaley. *Single-photon absorption and emission from a natural photosynthetic complex*. *Nature* 619(7969): pp. 300–304. 2023.

[118] I. Aharonovich, D. Englund, and M. Toth. *Solid-state single-photon emitters*. *Nature Photonics* 10(10): pp. 631–641. 2016.

[119] P. W. Shor and J. Preskill. *Simple Proof of Security of the BB84 Quantum Key Distribution Protocol*. *Phys. Rev. Lett.* 85: pp. 441–444. 2000.

[120] B. E. A. Saleh and M. C. Teich. *Fundamentals of Photonics*. John Wiley Sons Inc., 1991.

[121] M. Bozzio, C. Crépeau, P. Wallden, and P. Walther. *Quantum cryptography beyond key distribution: Theory and experiment*. *Rev. Mod. Phys.* 97: p. 045006. 2025.

[122] O. Sattath. *Uncloneable Cryptography*. *Commun. ACM* 66(11): pp. 78–86. 2023.

[123] A. Broadbent and S. Lord. *Uncloneable Quantum Encryption via Oracles*. In: *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*. Ed. by S. T. Flammia. Vol. 158. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020, 4:1–4:22. ISBN: 978-3-95977-146-7.

[124] M. S. Bullock, C. N. Gagatsos, S. Guha, and B. A. Bash. *Fundamental Limits of Quantum-Secure Covert Communication Over Bosonic Channels*. *IEEE Journal on Selected Areas in Communications* 38(3): pp. 471–482. 2020.

[125] B. A. Bash, A. H. Gheorghe, M. Patel, J. L. Habif, D. Goeckel, D. Towsley, and S. Guha. *Quantum-secure covert communication on bosonic channels*. *Nature communications* 6: p. 8626. 2015.

[126] J.-Q. Qin, J.-T. Wang, Y.-L. Yu, and X.-B. Wang. *General theory of quantum fingerprinting network*. *Phys. Rev. Res.* 3: p. 033039. 2021.

[127] R. Amiri, A. Abidin, P. Wallden, and E. Andersson. *Efficient Unconditionally Secure Signatures Using Universal Hashing*. In: *Applied Cryptography and Network Security*. Ed. by B. Preneel and F. Vercauteren. Cham: Springer International Publishing, 2018, pp. 143–162. ISBN: 978-3-319-93387-0.

[128] G. Brassard, C. Crepeau, R. Jozsa, and D. Langlois. *A quantum bit commitment scheme provably unbreakable by both parties*. In: *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*. 1993, pp. 362–371.

[129] M. Blum. *Coin flipping by telephone a protocol for solving impossible problems*. *ACM SIGACT News* 15(1): pp. 23–27. 1983.

[130] A. S. Arora, J. Roland, and C. Vlachou. *Analytic quantum weak coin flipping protocols with arbitrarily small bias*, 2019.

[131] S. Goldwasser, Y. T. Kalai, and G. N. Rothblum. *One-time programs*. In: *Annual International Cryptology Conference*. Springer. 2008, pp. 39–56.

[132] C. H. Bennett, G. Brassard, C. Crépeau, and M.-H. Skubiszewska. *Practical Quantum Oblivious Transfer*. In: *Advances in Cryptology — CRYPTO '91*. Ed. by J. Feigenbaum. Berlin, Heidelberg: Springer Berlin Heidelberg, 1992, pp. 351–366. ISBN: 978-3-540-46766-3.

[133] J. Bartusek, A. Coladangelo, D. Khurana, and F. Ma. *One-Way Functions Imply Secure Computation in a Quantum World*. In: *Advances in Cryptology – CRYPTO 2021*. Ed. by T. Malkin and C. Peikert. Cham: Springer International Publishing, 2021, pp. 467–496. ISBN: 978-3-030-84242-0.

[134] H.-K. Lo and H. F. Chau. *Unconditional Security of Quantum Key Distribution over Arbitrarily Long Distances*. *Science* 283(5410): pp. 2050–2056. 1999.

[135] M. Ben-Or, M. Horodecki, D. Leung, D. Mayers, and J. Oppenheim. *The Universal Composable Security of Quantum Key Distribution*. In: *Proceedings of the Second International Conference on Theory of Cryptography(TCC '05)'*. Ed. by J. Kilian. Springer. 2005, pp. 386–406.

[136] D. Gottesman, H.-K. Lo, N. Lütkenhaus, and J. Preskill. *Security of Quantum Key Distribution with Imperfect Devices*. *Quantum Info. Comput.* 4(5): pp. 325–360. 2004.

[137] G. Brassard and L. Salvail. *Advances in Cryptology—EUROCRYPT'93*. *Lecture Notes in Computer Science* 765: pp. 410–423. 1994.

[138] J. Ha, J. Kim, D. Klinc, and S. McLaughlin. *Rate-compatible punctured low-density parity-check codes with short block lengths*. *IEEE Transactions on Information Theory* 52(2): pp. 728–738. 2006.

[139] M. Tomamichel, C. Schaffner, A. Smith, and R. Renner. *Leftover Hashing Against Quantum Side Information*. *IEEE Transactions on Information Theory* 57(8): pp. 5524–5535. 2011.

[140] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi. *Fundamental limits of repeaterless quantum communications*. *Nature communications* 8: p. 15043. 2017.

[141] L. Salvail, M. Peev, E. Diamanti, R. Alléaume, N. Lütkenhaus, and T. Länger. *Security of trusted repeater quantum key distribution networks*. *Journal of Computer Security* 18(1): pp. 61–87. 2010.

[142] S. Wang, Z.-Q. Yin, D.-Y. He, W. Chen, R.-Q. Wang, P. Ye, Y. Zhou, G.-J. Fan-Yuan, F.-X. Wang, Y.-G. Zhu, P. V. Morozov, A. V. Divochiy, Z. Zhou, G.-C. Guo, and Z.-F. Han. *Twin-field quantum key distribution over 830-km fibre*. *Nature Photonics* 16(2): pp. 154–161. 2022.

[143] M. Blum. *Coin flipping by telephone*. In: *Advances in Cryptology: A Report on CRYPTO 81*. 1981, pp. 11–15.

[144] S. Even. *Protocol for signing contracts*. In: *Advances in Cryptology: A Report on CRYPTO 81*. 1981, pp. 148–153.

[145] L. Hardy and A. Kent. *Cheat Sensitive Quantum Bit Commitment*. *Phys. Rev. Lett.* 92: p. 157901. 2004.

[146] H. Buhrman, M. Christandl, P. Hayden, H.-K. Lo, and S. Wehner. *Possibility, impossibility, and cheat sensitivity of quantum-bit string commitment*. *Phys. Rev. A* 78: p. 022316. 2008.

[147] I. B. DamgÅrd, S. Fehr, L. Salvail, and C. Schaffner. *Cryptography in the Bounded-Quantum-Storage Model*. *SIAM Journal on Computing* 37(6): pp. 1865–1890. 2008.

[148] S. Wehner, C. Schaffner, and B. M. Terhal. *Cryptography from Noisy Storage*. *Phys. Rev. Lett.* 100: p. 220502. 2008.

[149] E. Verbanis, A. Martin, R. Houlmann, G. Boso, F. Bussières, and H. Zbinden. *24-Hour Relativistic Bit Commitment*. *Phys. Rev. Lett.* 117: p. 140506. 2016.

[150] S. Wehner, M. Curty, C. Schaffner, and H.-K. Lo. *Implementation of two-party protocols in the noisy-storage model*. *Phys. Rev. A* 81: p. 052336. 2010.

[151] F. Furrer, T. Gehring, C. Schaffner, C. Pacher, R. Schnabel, and S. Wehner. *Continuous-variable protocol for oblivious transfer in the noisy-storage model*. *Nature communications* 9(1): p. 1450. 2018.

[152] R. König and S. Wehner. *A Strong Converse for Classical Channel Coding Using Entangled Inputs*. *Phys. Rev. Lett.* 103: p. 070504. 2009.

[153] N. Jain, B. Stiller, I. Khan, D. Elser, C. Marquardt, and G. Leuchs. *Attacks on practical quantum key distribution systems (and how to prevent them)*. *Contemporary Physics* 57(3): pp. 366–387. 2016.

[154] M. Bozzio, A. Cavaillès, E. Diamanti, A. Kent, and D. Pitalúa-García. *Multiphoton and Side-Channel Attacks in Mistrustful Quantum Cryptography*. *PRX Quantum* 2: p. 030338. 2021.

[155] C. W. Helstrom. *Quantum detection and estimation theory*. *Journal of Statistical Physics* 1(2): 1969.

[156] S. M. Barnett and S. Croke. *Quantum state discrimination*. *Advances in Optics and Photonics* 1(2): p. 238. 2009.

[157] A. Vakhitov, V. Makarov, and D. R. Hjelme. *Large pulse attack as a method of conventional optical eavesdropping in quantum cryptography*. *Journal of Modern Optics* 48(13): pp. 2023–2038. 2001.

[158] L. Lydersen, C. Wiechers, C. Wittmann, D. Elser, J. Skaar, and V. Makarov. *Hacking commercial quantum cryptography systems by tailored bright illumination*. *Nature Photonics* 4(10): pp. 686–689. 2010.

[159] Y.-J. Qian, D.-Y. He, S. Wang, W. Chen, Z.-Q. Yin, G.-C. Guo, and Z.-F. Han. *Robust countermeasure against detector control attack in a practical quantum key distribution system*. *Optica* 6(9): pp. 1178–1184. 2019.

[160] N. Aharon, S. Massar, S. Pironio, and J. Silman. *Device-independent bit commitment based on the CHSH inequality*. *New Journal of Physics* 18(2): p. 025014. 2016.

[161] I. V. Puthoor, R. Amiri, P. Wallden, M. Curty, and E. Andersson. *Measurement-device-independent quantum digital signatures*. *Phys. Rev. A* 94: p. 022328. 2016.

[162] M. Bozzio, E. Diamanti, and F. Grosshans. *Semi-device-independent quantum money with coherent states*. *Phys. Rev. A* 99: p. 022336. 2019.

[163] C.-L. Li, Y. Fu, W.-B. Liu, Y.-M. Xie, B.-H. Li, M.-G. Zhou, H.-L. Yin, and Z.-B. Chen. *Breaking the rate-distance limitation of measurement-device-independent quantum secret sharing*. *Phys. Rev. Res.* 5: p. 033077. 2023.

[164] J. Kim, O. Benson, H. Kan, and Y. Yamamoto. *A single-photon turnstile device*. *Nature* 397(6719): pp. 500–503. 1999.

[165] D. Bimberg, M. Grundmann, and N. Ledentsov. *Quantum Dot Heterostructures*. Wiley, 1998.

[166] M. Gurioli, Z. Wang, A. Rastelli, T. Kuroda, and S. Sanguinetti. *Droplet epitaxy of semiconductor nanostructures for quantum photonic devices*. *Nature materials* 18(8): pp. 799–810. 2019.

[167] E. Goldmann, M. Paul, F. F. Krause, K. Müller, J. Kettler, T. Mehrtens, A. Rosenauer, M. Jetter, P. Michler, and F. Jahnke. *Structural and emission properties of InGaAs/GaAs quantum dots emitting at 1.3um*. *Applied Physics Letters* 105(15): p. 152102. 2014.

[168] R. Sittig, C. Nawrath, S. Kolatschek, S. Bauer, R. Schaber, J. Huang, P. Vijayan, P. Pruy, S. L. Portalupi, M. Jetter, and P. Michler. *Thin-film InGaAs metamorphic buffer for telecom C-band InAs quantum dots and optical resonators on GaAs platform*. *Nanophotonics* 11(6): pp. 1109–1116. 2022.

[169] P. Hawrylak and A. Wojs. *Electronic structure and optical properties of self-assembled quantum dots*. *Semiconductor Science and Technology* 11(11S): p. 1516. 1996.

[170] *Band Structure*. In: *The Physics of Semiconductors: An Introduction Including Devices and Nanophysics*. Springer Berlin Heidelberg, Berlin, Heidelberg. Pp. 111–147. 2006.

[171] T. Takagahara. *Biexciton states in semiconductor quantum dots and their nonlinear optical properties*. *Phys. Rev. B* 39: pp. 10206–10231. 1989.

[172] M. Bayer, G. Ortner, O. Stern, A. Kuther, A. A. Gorbunov, A. Forchel, P. Hawrylak, S. Fafard, K. Hinzer, T. L. Reinecke, S. N. Walck, J. P. Reithmaier, F. Klopf, and F. Schäfer. *Fine structure of neutral and charged excitons in self-assembled In(Ga)As/(Al)GaAs quantum dots*. *Phys. Rev. B* 65: p. 195315. 2002.

[173] J. Liu, R. Su, Y. Wei, B. Yao, S. F. C. d. Silva, Y. Yu, J. Iles-Smith, K. Srinivasan, A. Rastelli, J. Li, and X. Wang. *A solid-state source of strongly entangled photon pairs with high brightness and indistinguishability*. *Nature nanotechnology* 14(6): pp. 586–593. 2019.

[174] T. Seidelmann, C. Schimpf, T. K. Bracht, M. Cosacchi, A. Vagov, A. Rastelli, D. E. Reiter, and V. M. Axt. *Two-Photon Excitation Sets Limit to Entangled Photon Pair Generation from Quantum Emitters*. *Physical Review Letters* 129(19): p. 193604. 2022.

[175] F. Basso Basset, M. B. Rota, M. Beccaceci, T. M. Krieger, Q. Buchinger, J. Neuwirth, H. Huet, S. Stroj, S. F. Da Covre Silva, G. Ronco, C. Schimpf, S. Höfling, T. Huber-Loyola, A. Rastelli, and R. Trotta. *Signatures of the Optical Stark Effect on Entangled Photon Pairs from Resonantly Pumped Quantum Dots*. *Physical review letters* 131(16): p. 166901. 2023.

[176] A. Laneve, G. Ronco, M. Beccaceci, P. Barigelli, F. Salusti, N. Claro-Rodriguez, G. D. Pascalis, A. Suprano, L. Chiaudano, E. Schöll, L. Hanschke, T. M. Krieger, Q. Buchinger, S. F. C. da Silva, J. Neuwirth, S. Stroj, S. Höfling, T. Huber-Loyola, M. A. U. Castaneda, G. Carvacho, N. Spagnolo, M. B. Rota, F. B. Basset, A. Rastelli, F. Sciarrino, K. Jöns, and R. Trotta. *Quantum teleportation with dissimilar quantum dots over a hybrid quantum network*, 2024.

[177] D. Huber, M. Reindl, S. F. Da Covre Silva, C. Schimpf, J. Martín-Sánchez, H. Huang, G. Piredda, J. Edlinger, A. Rastelli, and R. Trotta. *Strain-Tunable GaAs Quantum Dot: A Nearly Dephasing-Free Source of Entangled Photon Pairs on Demand*. *Physical review letters* 121(3): p. 033902. 2018.

[178] P.-L. Ardelt, L. Hanschke, K. A. Fischer, K. Müller, A. Kleinkauf, M. Koller, A. Bechtold, T. Simmet, J. Wierzbowski, H. Riedl, G. Abstreiter, and J. J. Finley. *Dissipative preparation of the exciton and biexciton in self-assembled quantum dots on picosecond time scales*. *Phys. Rev. B* 90: p. 241404. 2014.

[179] S. Bounouar, M. Müller, A. M. Barth, M. Glässl, V. M. Axt, and P. Michler. *Phonon-assisted robust and deterministic two-photon biexciton preparation in a quantum dot*. *Phys. Rev. B* 91: p. 161302. 2015.

[180] T. K. Bracht, M. Cygorek, T. Seidelmann, V. M. Axt, and D. E. Reiter. *Temperature-independent almost perfect photon entanglement from quantum dots via the SUPER scheme*. *Optica Quantum* 1(2): p. 103. 2023.

[181] J. Bienfang, T. Gerrits, P. Kuo, A. Migdall, S. Polyakov, and O. Slattery. *Single-photon Sources and Detectors Dictionary*, 2023.

[182] A. V. Kuhlmann, J. Houel, A. Ludwig, L. Greuter, D. Reuter, A. D. Wieck, M. Poggio, and R. J. Warburton. *Charge noise and spin noise in a semiconductor quantum device*. *Nature Physics* 9(9): pp. 570–575. 2013.

[183] J. Kettler. *Telecom-wavelength nonclassical light from single In(As)GaAs quantum dots*. PhD thesis. Universität Stuttgart, Stuttgart. 2017.

[184] R. J. Warburton, C. Schäflein, D. Haft, F. Bickel, A. Lorke, K. Karrai, J. M. Garcia, W. Schoenfeld, and P. M. Petroff. *Optical emission from a charge-tunable quantum ring*. *Nature* 405(6789): pp. 926–929. 2000.

[185] J. J. Finley, P. W. Fry, A. D. Ashmore, A. Lemaıtre, A. I. Tartakovskii, R. Oulton, D. J. Mowbray, M. S. Skolnick, M. Hopkinson, P. D. Buckle, and P. A. Maksym. *Observation of multicharged excitons and biexcitons in a single InGaAs quantum dot*. *Phys. Rev. B* 63: p. 161305. 2001.

[186] J. H. Prechtel, P. A. Dalgarno, R. H. Hadfield, J. McFarlane, A. Badolato, P. M. Petroff, and R. J. Warburton. *Fast electro-optics of a single self-assembled quantum dot in a charge-tunable device*. *Journal of Applied Physics* 111(4): 2012.

[187] L. Schweickert, K. D. Jöns, K. D. Zeuner, S. F. Covre da Silva, H. Huang, T. Lettner, M. Reindl, J. Zichi, R. Trotta, A. Rastelli, and V. Zwiller. *On-demand generation of background-free single photons from a solid-state source*. *Applied Physics Letters* 112(9): p. 093106. 2018.

[188] L. Seravalli, G. Trevisi, P. Frigeri, S. Franchi, M. Geddo, and G. Guizzetti. *The role of wetting layer states on the emission efficiency of InAs/InGaAs metamorphic quantum dot nanostructures*. *Nanotechnology* 20(27): p. 275703. 2009.

[189] S. Kolatschek, C. Nawrath, S. Bauer, J. Huang, J. Fischer, R. Sittig, M. Jetter, S. L. Portalupi, and P. Michler. *Bright Purcell enhanced single-photon source in the telecom O-band based on a quantum dot in a circular Bragg grating*. *Nano Letters* 21(18): pp. 7740–7745. 2021.

[190] F. Olbrich, J. Kettler, M. Bayerbach, M. Paul, J. Höschele, S. L. Portalupi, M. Jetter, and P. Michler. *Temperature-dependent properties of single long-wavelength InGaAs quantum dots embedded in a strain reducing layer*. *Journal of Applied Physics* 121(18): p. 184302. 2017.

[191] K. A. Fischer, K. Müller, K. G. Lagoudakis, and J. Vučković. *Dynamical modeling of pulsed two-photon interference*. *New Journal of Physics* 18(11): p. 113053. 2016.

[192] F. Sbresny, C. Calcagno, S. K. Kim, K. Boos, W. Rauhaus, F. Bopp, H. Riedl, J. J. Finley, E. Z. Casalengua, and K. Müller. *Selective filtering of multi-photon events from a single-photon emitter*, 2025.

[193] V. Zwiller, T. Aichele, and O. Benson. *Single-photon Fourier spectroscopy of excitons and biexcitons in single quantum dots*. *Phys. Rev. B* 69: p. 165307. 2004.

[194] M. Anderson, T. Müller, J. Skiba-Szymanska, A. B. Krysa, J. Huwer, R. M. Stevenson, J. Heffernan, D. A. Ritchie, and A. J. Shields. *Coherence in single photon emission from droplet epitaxy and Stranski–Krastanov quantum dots in the telecom C-band*. *Applied Physics Letters* 118(1): p. 014003. 2021.

[195] B. Krummheuer, V. M. Axt, and T. Kuhn. *Theory of pure dephasing and the resulting absorption line shape in semiconductor quantum dots*. *Physical Review B* 65(19): 2002.

[196] A. Berthelot, I. Favero, G. Cassabois, C. Voisin, C. Delalande, P. Roussignol, R. Ferreira, and J. M. Gérard. *Unconventional motional narrowing in the optical spectrum of a semiconductor quantum dot*. *Nature Physics* 2(11): pp. 759–764. 2006.

[197] J. Houel, A. V. Kuhlmann, L. Greuter, F. Xue, M. Poggio, B. D. Gerardot, P. A. Dalgarno, A. Badolato, P. M. Petroff, A. Ludwig, D. Reuter, A. D. Wieck, and R. J. Warburton. *Probing Single-Charge Fluctuations at a* GaAs/AlAs *Interface Using Laser Spectroscopy on a Nearby InGaAs Quantum Dot*. *Phys. Rev. Lett.* 108: p. 107401. 2012.

[198] A. V. Kuhlmann, J. H. Prechtel, J. Houel, A. Ludwig, D. Reuter, A. D. Wieck, and R. J. Warburton. *Transform-limited single photons from a single quantum dot*. *Nature communications* 6: p. 8204. 2015.

[199] M. E. Reimer, G. Bulgarini, A. Fognini, R. W. Heeres, B. J. Witek, M. A. M. Versteegh, A. Rubino, T. Braun, M. Kamp, S. Höfling, D. Dalacu, J. Lapointe, P. J. Poole, and V. Zwiller. *Overcoming power broadening of the quantum dot emission in a pure wurtzite nanowire*. *Phys. Rev. B* 93: p. 195316. 2016.

[200] H. Vural, J. Maisch, I. Gerhardt, M. Jetter, S. L. Portalupi, and P. Michler. *Characterization of spectral diffusion by slow-light photon-correlation spectroscopy*. *Phys. Rev. B* 101: p. 161401. 2020.

[201] H. Vural, S. L. Portalupi, and P. Michler. *Perspective of self-assembled InGaAs quantum-dots for multi-source quantum implementations*. *Applied Physics Letters* 117(3): p. 030501. 2020.

[202] D. Chen, G. R. Lander, K. S. Krowpman, G. S. Solomon, and E. B. Flagg. *Characterization of the local charge environment of a single quantum dot via resonance fluorescence*. *Phys. Rev. B* 93: p. 115307. 2016.

[203] A. Reigue, R. Hostein, and V. Voliotis. *Resonance fluorescence of a single semiconductor quantum dot: the impact of a fluctuating electrostatic environment*. *Semiconductor Science and Technology* 34(11): p. 113001. 2019.

[204] S. M. Ulrich, S. Ates, S. Reitzenstein, A. Löffler, A. Forchel, and P. Michler. *Dephasing of Triplet-Sideband Optical Emission of a Resonantly Driven* InAs/GaAs *Quantum Dot inside a Microcavity*. *Phys. Rev. Lett.* 106: p. 247402. 2011.

[205] A. Thoma, P. Schnauber, M. Gschrey, M. Seifried, J. Wolters, J.-H. Schulze, A. Strittmatter, S. Rodt, A. Carmele, A. Knorr, T. Heindel, and S. Reitzenstein. *Exploring Dephasing of a Solid-State Quantum Emitter via Time- and Temperature-Dependent Hong-Ou-Mandel Experiments*. *Phys. Rev. Lett.* 116: p. 033601. 2016.

[206] T. Takagahara. *Electron-phonon interactions and excitonic dephasing in semiconductor nanocrystals*. *Phys. Rev. Lett.* 71: pp. 3577–3580. 1993.

[207] T. E. Hodgson, L. Viola, and I. D'Amico. *Decoherence-protected storage of exciton qubits through ultrafast multipulse control*. *Physical Review B* 78(16): 2008.

[208] J. Iles-Smith, D. P. S. McCutcheon, A. Nazir, and J. Mørk. *Phonon scattering inhibits simultaneous near-unity efficiency and indistinguishability in semiconductor single-photon sources*. *Nature Photonics* 11(8): pp. 521–526. 2017.

[209] C. K. Hong, Z. Y. Ou, and L. Mandel. *Measurement of subpicosecond time intervals between two photons by interference*. *Phys. Rev. Lett.* 59: pp. 2044–2046. 1987.

[210] E. Knill, R. Laflamme, and G. J. Milburn. *A scheme for efficient quantum computation with linear optics*. *Nature* 409(6816): pp. 46–52. 2001.

[211] A.-N. Zhang, C.-Y. Lu, X.-Q. Zhou, Y.-A. Chen, Z. Zhao, T. Yang, and J.-W. Pan. *Experimental construction of optical multiqubit cluster states from Bell states*. *Phys. Rev. A* 73: p. 022330. 2006.

[212] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn. *Linear optical quantum computing with photonic qubits*. *Rev. Mod. Phys.* 79(1): pp. 135–174. 2007.

[213] M. C. Löbl, L. A. Pettersson, S. Paesani, and A. S. Sørensen. *Transforming graph states via Bell state measurements*. *Quantum* 9: p. 1795. 2025.

[214] J. C. Loredo, C. Antón, B. Reznychenko, P. Hilaire, A. Harouri, C. Millet, H. Ollivier, N. Somaschi, L. de Santis, A. Lemaître, I. Sagnes, L. Lanco, A. Auffèves, O. Krebs, and P. Senellart. *Generation of non-classical light in a photon-number superposition*. *Nature Photonics* 13(11): pp. 803–808. 2019.

[215] Y. Karli, D. A. Vajner, F. Kappe, P. C. A. Hagen, L. M. Hansen, R. Schwarz, T. K. Bracht, C. Schimpf, S. F. Da Covre Silva, P. Walther, A. Rastelli, V. M. Axt, J. C. Loredo, V. Remesh, T. Heindel, D. E. Reiter, and G. Weihs. *Controlling the photon number coherence of solid-state quantum light sources for quantum cryptography*. *npj Quantum Information* 10(1): 2024.

[216] Y.-J. Wei, Y.-M. He, M.-C. Chen, Y.-N. Hu, Y. He, D. Wu, C. Schneider, M. Kamp, S. Höfling, C.-Y. Lu, and J.-W. Pan. *Deterministic and robust generation of single photons from a single quantum dot with 99.5% indistinguishability using adiabatic rapid passage*. *Nano letters* 14(11): pp. 6515–6519. 2014.

[217] M. Reindl, J. H. Weber, D. Huber, C. Schimpf, S. F. Da Covre Silva, S. L. Portalupi, R. Trotta, P. Michler, and A. Rastelli. *Highly indistinguishable single photons from incoherently excited quantum dots*. *Physical Review B* 100(15): 2019.

[218] N. Hauser, M. Bayerbach, J. Kaupp, Y. Reum, G. Peniakov, J. Michl, M. Kamp, T. Huber-Loyola, A. T. Pfenning, S. Höfling, and S. Barz. *Deterministic and highly indistinguishable single photons in the telecom C-band*, 2025.

[219] C. Schimpf. *Quantum communication with semiconductor quantum dots*. PhD thesis. Universität Linz, Linz. 2022.

[220] F. Giorgino, P. Zahálka, L. Jehle, L. Carosini, L. M. Hansen, J. C. Loredo, and P. Walther. *Multi-photon emission from a resonantly pumped quantum dot*. *Optica Quantum* 3(4): pp. 402–407. 2025.

[221] P. Grünwald. *Effective second-order correlation function and single-photon detection*. *New Journal of Physics* 21(9): p. 093003. 2019.

[222] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg. *The Dressed Atom Approach*. In: *Atom—Photon Interactions*. John Wiley & Sons, Ltd, pp. 407–514. 1998.

[223] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg. *The Optical Bloch Equations*. In: *Atom—Photon Interactions*. John Wiley & Sons, Ltd, pp. 353–405. 1998.

[224] B. R. Mollow. *Power Spectrum of Light Scattered by Two-Level Systems*. *Physical Review Letters* 188(5): 1969.

[225] L. Besombes, K. Kheng, L. Marsal, and H. Mariette. *Acoustic phonon broadening mechanism in single quantum dot emission*. *Phys. Rev. B* 63: p. 155307. 2001.

[226] *Many-Particle Physics*. Springer Berlin, 2000.

[227] R. A. DeCrescent, Z. Wang, J. T. Bush, P. Imany, A. Kwiatkowski, D. V. Reddy, S. W. Nam, R. P. Mirin, and K. L. Silverman. *Coherent Control of an Optical Quantum Dot Using Phonons and Photons*, 2024.

[228] A. J. Ramsay, T. M. Godden, S. J. Boyle, E. M. Gauger, A. Nazir, B. W. Lovett, A. M. Fox, and M. S. Skolnick. *Phonon-Induced Rabi-Frequency Renormalization of Optically Driven Single* InGaAs/GaAs *Quantum Dots*. *Phys. Rev. Lett.* 105: p. 177402. 2010.

[229] A. J. Ramsay, A. V. Gopal, E. M. Gauger, A. Nazir, B. W. Lovett, A. M. Fox, and M. S. Skolnick. *Damping of Exciton Rabi Rotations by Acoustic Phonons in Optically Excited* InGaAs/GaAs *Quantum Dots*. *Phys. Rev. Lett.* 104: p. 017402. 2010.

[230] Y. Toda, O. Moriwaki, M. Nishioka, and Y. Arakawa. *Efficient Carrier Relaxation Mechanism in InGaAs/GaAs Self-Assembled Quantum Dots Based on the Existence of Continuum States*. *Phys. Rev. Lett.* 82: pp. 4114–4117. 1999.

[231] F. Findeis, A. Zrenner, G. Böhm, and G. Abstreiter. *Phonon-assisted biexciton generation in a single quantum dot*. *Phys. Rev. B* 61: R10579–R10582. 2000.

[232] R. Oulton, J. J. Finley, A. I. Tartakovskii, D. J. Mowbray, M. S. Skolnick, M. Hopkinson, A. Vasanelli, R. Ferreira, and G. Bastard. *Continuum transitions and phonon coupling in single self-assembled Stranski-Krastanow quantum dots*. *Phys. Rev. B* 68: p. 235301. 2003.

[233] T. Takagahara. *Theory of exciton dephasing in semiconductor quantum dots*. *Phys. Rev. B* 60: pp. 2638–2652. 1999.

[234] M. Glässl, A. M. Barth, and V. M. Axt. *Proposed robust and high-fidelity preparation of excitons and biexcitons in semiconductor quantum dots making active use of phonons*. *Physical review letters* 110(14): p. 147401. 2013.

[235] J. H. Quilter, A. J. Brash, F. Liu, M. Glässl, A. M. Barth, V. M. Axt, A. J. Ramsay, M. S. Skolncik, and A. M. Fox. *Phonon-assisted Population Inversion of a Single InGaAs/GaAs Quantum Dot by Pulsed Laser Excitation*. *Physical Review Letters* 114(13): 2015.

[236] M. Glässl, A. Vagov, S. Lüker, D. E. Reiter, M. D. Croitoru, P. Machnikowski, V. M. Axt, and T. Kuhn. *Long-time dynamics and stationary nonequilibrium of an optically driven strongly confined quantum dot coupled to phonons*. *Phys. Rev. B* 84: p. 195311. 2011.

[237] S. Frick, R. Keil, V. Remesh, and G. Weihs. *Single-Photon Sources for Multi-Photon Applications*. In: *Photonic Quantum Technologies*. John Wiley & Sons, Ltd, pp. 53–84. 2023.

[238] Y.-M. He, H. Wang, C. Wang, M.-C. Chen, X. Ding, J. Qin, Z.-C. Duan, S. Chen, J.-P. Li, R.-Z. Liu, C. Schneider, M. Atatüre, S. Höfling, C.-Y. Lu, and J.-W. Pan. *Coherently driving a single quantum two-level system with dichromatic laser pulses*. *Nature Physics* 15(9): pp. 941–946. 2019.

[239] T. K. Bracht, M. Cosacchi, T. Seidelmann, M. Cygorek, A. Vagov, V. M. Axt, T. Heindel, and D. E. Reiter. *Swing-Up of Quantum Emitter Population Using Detuned Pulses*. *PRX Quantum* 2(4): 2021.

[240] Y. Karli, F. Kappe, V. Remesh, T. K. Bracht, J. Münzberg, S. Da Covre Silva, T. Seidelmann, V. M. Axt, A. Rastelli, D. E. Reiter, and G. Weihs. *SUPER Scheme in Action: Experimental Demonstration of Red-Detuned Excitation of a Quantum Emitter*. *Nano letters* 22(16): pp. 6567–6572. 2022.

[241] H. Wang, J. Qin, X. Ding, M.-C. Chen, S. Chen, X. You, Y.-M. He, X. Jiang, L. You, Z. Wang, C. Schneider, J. J. Renema, S. Höfling, C.-Y. Lu, and J.-W. Pan. *Boson Sampling with 20 Input Photons and a 60-Mode Interferometer in a* 10^{14}*-Dimensional Hilbert Space*. *Phys. Rev. Lett.* 123: p. 250503. 2019.

[242] L. Zhai, G. N. Nguyen, C. Spinnler, J. Ritzmann, M. C. Löbl, A. D. Wieck, A. Ludwig, A. Javadi, and R. J. Warburton. *Quantum interference of identical photons from remote GaAs quantum dots*. *Nature nanotechnology* 17(8): pp. 829–833. 2022.

[243] C. Matthiesen, A. N. Vamivakas, and M. Atatüre. *Subnatural Linewidth Single Photons from a Quantum Dot*. *Phys. Rev. Lett.* 108: p. 093602. 2012.

[244] J. C. López Carreño, E. Z. Casalengua, F. P. Laussy, and E. del Valle. *Joint subnatural-linewidth and single-photon emission from resonance fluorescence*. *Quantum Science and Technology* 3(4): p. 045001. 2018.

[245] L. Hanschke, L. Schweickert, J. C. L. Carreño, E. Schöll, K. D. Zeuner, T. Lettner, E. Z. Casalengua, M. Reindl, S. F. C. da Silva, R. Trotta, J. J. Finley, A. Rastelli, E. del Valle, F. P. Laussy, V. Zwiller, K. Müller, and K. D. Jöns. *Origin of Antibunching in Resonance Fluorescence*. *Phys. Rev. Lett.* 125: p. 170402. 2020.

[246] A. Aspect, G. Roger, S. Reynaud, J. Dalibard, and and C. Cohen-Tannoudji. *Time Correlations between the Two Sidebands of the Resonance Fluorescence Triplet*. *Phys. Rev. Lett.*: 1980.

[247] A. Ulhaq, S. Weiler, S. M. Ulrich, R. Roßbach, M. Jetter, and P. Michler. *Cascaded single-photon emission from the Mollow triplet sidebands of a quantum dot*. *Nature Photonics* 6(4): pp. 238–242. 2012.

[248] J. Yang, I. Strandberg, A. Vivas-Viaña, A. Gaikwad, C. Castillo-Moreno, A. F. Kockum, M. A. Ullah, C. S. Muñoz, A. M. Eriksson, and S. Gasparinetti. *Entanglement of photonic modes from a continuously driven two-level system*. *npj Quantum Information* 11(1): p. 69. 2025.

[249] T. H. Stievater, X. Li, D. G. Steel, D. Gammon, D. S. Katzer, D. Park, C. Piermarocchi, and L. J. Sham. *Rabi Oscillations of Excitons in Single Quantum Dots*. *Phys. Rev. Lett.* 87: p. 133603. 2001.

[250] J. Förstner, C. Weber, J. Danckwerts, and A. Knorr. *Phonon-Assisted Damping of Rabi Oscillations in Semiconductor Quantum Dots*. *Phys. Rev. Lett.* 91: p. 127401. 2003.

[251] D. E. Reiter, T. Kuhn, M. Glässl, and V. M. Axt. *The role of phonons for exciton and biexciton generation in an optically driven quantum dot*. *Journal of physics. Condensed matter : an Institute of Physics journal* 26(42): p. 423203. 2014.

[252] A. Vagov, M. D. Croitoru, V. M. Axt, T. Kuhn, and F. M. Peeters. *Nonmonotonic Field Dependence of Damping and Reappearance of Rabi Oscillations in Quantum Dots*. *Phys. Rev. Lett.* 98: p. 227403. 2007.

[253] L. Hanschke, T. K. Bracht, E. Schöll, D. Bauch, E. Berger, P. Kallert, M. Peter, A. J. Garcia, S. F. Covre da Silva, S. Manna, A. Rastelli, S. Schumacher, D. E. Reiter, and K. D. Jöns. *Experimental Measurement of the Reappearance of Rabi Rotations in Semiconductor Quantum Dots*. *Physical Review Letters* 135(26): 2025.

[254] N. Somaschi, V. Giesz, L. de Santis, J. C. Loredo, M. P. Almeida, G. Hornecker, S. L. Portalupi, T. Grange, C. Antón, J. Demory, C. Gómez, I. Sagnes, N. D. Lanzillotti-Kimura, A. Lemaítre, A. Auffeves, A. G. White, L. Lanco, and P. Senellart. *Near-optimal single-photon sources in the solid state*. *Nature Photonics* 10(5): pp. 340–345. 2016.

[255] A. Javadi, N. Tomm, N. O. Antoniadis, A. J. Brash, R. Schott, S. R. Valentin, A. D. Wieck, A. Ludwig, and R. J. Warburton. *Cavity-enhanced excitation of a quantum dot in the picosecond regime*. *New Journal of Physics* 25(9): p. 093027. 2023.

[256] P. Laccotripes, J. Huang, G. Shooter, A. Barbiero, M. S. Winnel, D. A. Ritchie, A. J. Shields, T. Muller, and R. M. Stevenson. *An entangled photon source for the telecom C-band based on a semiconductor-confined spin*, 2025.

[257] L. Rickert, D. A. Vajner, M. von Helversen, J. Schall, S. Rodt, S. Reitzenstein, H. Liu, S. Li, H. Ni, Z. Niu, and T. Heindel. *High Purcell Enhancement in Quantum-Dot Hybrid Circular Bragg Grating Cavities for GHz Clock Rate Generation of Indistinguishable Photons*. *ACS Photonics* 12(1): pp. 464–475. 2025.

[258] P. Lodahl, S. Mahmoodian, and S. Stobbe. *Interfacing single photons and single quantum dots with photonic nanostructures*. *Rev. Mod. Phys.* 87: pp. 347–400. 2015.

[259] M. Arcari, I. Söllner, A. Javadi, S. Lindskov Hansen, S. Mahmoodian, J. Liu, H. Thyrrestrup, E. H. Lee, J. D. Song, S. Stobbe, and P. Lodahl. *Near-Unity Coupling Efficiency of a Quantum Emitter to a Photonic Crystal Waveguide*. *Phys. Rev. Lett.* 113: p. 093603. 2014.

[260] M. Reindl, K. D. Jöns, D. Huber, C. Schimpf, Y. Huo, V. Zwiller, A. Rastelli, and R. Trotta. *Phonon-Assisted Two-Photon Interference from Remote Quantum Emitters*. *Nano letters* 17(7): pp. 4090–4095. 2017.

[261] S. Weiler, A. Ulhaq, S. M. Ulrich, D. Richter, M. Jetter, P. Michler, C. Roy, and S. Hughes. *Phonon-assisted incoherent excitation of a quantum dot and its emission properties*. *Phys. Rev. B* 86: p. 241304. 2012.

[262] C. Gustin and S. Hughes. *Efficient Pulse–Excitation Techniques for Single Photon Sources from Quantum Dots in Optical Cavities*. *Advanced Quantum Technologies* 3(2): 2020.

[263] A. M. Barth, S. Lüker, A. Vagov, D. E. Reiter, T. Kuhn, and V. M. Axt. *Fast and selective phonon-assisted state preparation of a quantum dot by adiabatic undressing*. *Phys. Rev. B* 94: p. 045306. 2016.

[264] S. E. Thomas, M. Billard, N. Coste, S. C. Wein, Priya, H. Ollivier, O. Krebs, L. Tazaïrt, A. Harouri, A. Lemaitre, I. Sagnes, C. Anton, L. Lanco, N. Somaschi, J. C. Loredo, and P. Senellart. *Bright Polarized Single-Photon Source Based on a Linear Dipole*. *Physical review letters* 126(23): p. 233601. 2021.

[265] R. Behrends, M. v. Helversen, P. K. Saha, L. Rickert, K. Kaymazlar, M. Lach, N. D. Kewitz, J. Kaupp, Y. Reum, T. Huber-Loyola, S. Höfling, A. Pfenning, and T. Heindel. *Highly-Indistinguishable Single-Photons at 1550 nm from a Two-photon Resonantly Excited Purcell-enhanced Quantum Dot*, 2026.

[266] A. C. Dada, T. S. Santana, R. N. E. Malein, A. Koutroumanis, Y. Ma, J. M. Zajac, J. Y. Lim, J. D. Song, and B. D. Gerardot. *Indistinguishable single photons with flexible electronic triggering*. *Optica* 3(5): pp. 493–498. 2016.

[267] K. A. Fischer, L. Hanschke, M. Kremser, J. J. Finley, K. Müller, and J. Vučković. *Pulsed Rabi oscillations in quantum two-level systems: beyond the area theorem*. *Quantum Science and Technology* 3(1): p. 014006. 2018.

[268] C. Gustin and S. Hughes. *Pulsed excitation dynamics in quantum-dot–cavity systems: Limits to optimizing the fidelity of on-demand single-photon sources*. *Physical Review B* 98(4): 2018.

[269] K. Müller, K. A. Fischer, A. Rundquist, C. Dory, K. G. Lagoudakis, T. Sarmiento, Y. A. Kelaita, V. Borish, and J. Vu čkovi ć. *Ultrafast Polariton-Phonon Dynamics of Strongly Coupled Quantum Dot-Nanocavity Systems*. *Phys. Rev. X* 5: p. 031006. 2015.

[270] L. Rickert, K. Żołnacz, D. A. Vajner, M. von Helversen, S. Rodt, S. Reitzenstein, H. Liu, S. Li, H. Ni, P. Wyborski, G. Sek, A. Musiał, Z. Niu, and T. Heindel. *A fiber-pigtailed quantum dot device generating indistinguishable photons at GHz clock-rates*. *Nanophotonics* 14(11): pp. 1795–1808. 2025.

[271] H. Y. Hui and R. B. Liu. *Proposal for geometric generation of a biexciton in a quantum dot using a chirped pulse*. *Phys. Rev. B* 78: p. 155315. 2008.

[272] K. Schuh, F. Jahnke, and M. Lorke. *Rapid adiabatic passage in quantum dots: Influence of scattering and dephasing*. *Applied Physics Letters* 99(1): p. 011105. 2011.

[273] K. Bergmann, H. Theuer, and B. W. Shore. *Coherent population transfer among quantum states of atoms and molecules*. *Rev. Mod. Phys.* 70: pp. 1003–1025. 1998.

[274] F. Kappe, Y. Karli, T. K. Bracht, S. F. Covre da Silva, T. Seidelmann, V. M. Axt, A. Rastelli, G. Weihs, D. E. Reiter, and V. Remesh. *Collective excitation of spatio-spectrally distinct quantum dots enabled by chirped pulses*. *Materials for Quantum Technology* 3(2): p. 025006. 2023.

[275] C. Piccinini, A. Rodek, A. A. Madigawa, A. G. Jr., S. F. C. da Silva, V. Remesh, L. Vannucci, G. Weihs, A. Rastelli, N. Gregersen, and B. Munkhbat. *Exciton and biexciton preparation via coherent swing-up excitation in a GaAs quantum dot embedded in micropillar cavity*, 2025.

[276] K. Boos, F. Sbresny, S. K. Kim, M. Kremser, H. Riedl, F. W. Bopp, W. Rauhaus, B. Scaparra, K. D. Jöns, J. J. Finley, K. Müller, and L. Hanschke. *Coherent Swing-Up Excitation for Semiconductor Quantum Dots*. *Advanced Quantum Technologies* 7(4): p. 2300359. 2024.

[277] T. K. Bracht, T. Seidelmann, T. Kuhn, V. M. Axt, and D. E. Reiter. *Phonon Wave Packet Emission during State Preparation of a Semiconductor Quantum Dot using Different Schemes*. *physica status solidi (b)* 259(6): p. 2100649. 2022.

[278] T. Flissikowski, A. Betke, I. A. Akimov, and F. Henneberger. *Two-Photon Coherent Control of a Single Quantum Dot*. *Phys. Rev. Lett.* 92: p. 227401. 2004.

[279] S. Stufler, P. Machnikowski, P. Ester, M. Bichler, V. M. Axt, T. Kuhn, and A. Zrenner. *Two-photon Rabi oscillations in a single* $\mathrm{In}_x\mathrm{Ga}_{1-x}\mathrm{AsGaAs}$ *quantum dot*. *Phys. Rev. B* 73: p. 125304. 2006.

[280] E. Schöll, L. Schweickert, L. Hanschke, K. D. Zeuner, F. Sbresny, T. Lettner, R. Trivedi, M. Reindl, S. F. Da Covre Silva, R. Trotta, J. J. Finley, J. Vučković, K. Müller, A. Rastelli, V. Zwiller, and K. D. Jöns. *Crux of Using the Cascaded Emission of a Three-Level Quantum Ladder System to Generate Indistinguishable Photons*. *Physical review letters* 125(23): p. 233605. 2020.

[281] C. Schimpf, M. Reindl, F. Basso Basset, K. D. Jöns, R. Trotta, and A. Rastelli. *Quantum dots as potential sources of strongly entangled photons: Perspectives and challenges for applications in quantum networks*. *Applied Physics Letters* 118(10): p. 100502. 2021.

[282] F. Sbresny, L. Hanschke, E. Schöll, W. Rauhaus, B. Scaparra, K. Boos, E. Zubizarreta Casalengua, H. Riedl, E. Del Valle, J. J. Finley, K. D. Jöns, and K. Müller. *Stimulated Generation of Indistinguishable Single Photons from a Quantum Ladder System*. *Physical review letters* 128(9): p. 093603. 2022.

[283] T. Kaldewey, S. Lüker, A. V. Kuhlmann, S. R. Valentin, A. Ludwig, A. D. Wieck, D. E. Reiter, T. Kuhn, and R. J. Warburton. *Coherent and robust high-fidelity generation of a biexciton in a quantum dot by rapid adiabatic passage*. *Phys. Rev. B* 95: p. 161302. 2017.

[284] V. Zwiller and G. Björk. *Improved light extraction from emitters in high refractive index materials using solid immersion lenses*. *Journal of Applied Physics* 92(2): pp. 660–665. 2002.

[285] S. Moehl, H. Zhao, B. D. Don, S. Wachter, and H. Kalt. *Solid immersion lens-enhanced nano-photoluminescence: Principle and applications*. *Journal of Applied Physics* 93(10): pp. 6265–6272. 2003.

[286] M. Gschrey, A. Thoma, P. Schnauber, M. Seifried, R. Schmidt, B. Wohlfeil, L. Krüger, J.-H. Schulze, T. Heindel, S. Burger, F. Schmidt, A. Strittmatter, S. Rodt, and S. Reitzenstein. *Highly indistinguishable photons from deterministic quantum-dot microlenses utilizing three-dimensional in situ electron-beam lithography*. *Nature communications* 6: p. 7662. 2015.

[287] M. T. Borgström, V. Zwiller, E. Müller, and A. Imamoglu. *Optically Bright Quantum Dots in Single Nanowires*. *Nano Letters* 5(7): pp. 1439–1443. 2005.

[288] J. Claudon, J. Bleuse, N. S. Malik, M. Bazin, P. Jaffrennou, N. Gregersen, C. Sauvan, P. Lalanne, and J.-M. Gérard. *A highly efficient single-photon source based on a quantum dot in a photonic nanowire*. *Nature Photonics*: 2010.

[289] M. Munsch, N. S. Malik, E. Dupuy, A. Delga, J. Bleuse, J.-M. Gérard, J. Claudon, N. Gregersen, and J. Mørk. *Dielectric GaAs Antenna Ensuring an Efficient Broadband Coupling between an InAs Quantum Dot and a Gaussian Optical Beam*. *Phys. Rev. Lett.* 110: p. 177402. 2013.

[290] M. Heiss, Y. Fontana, A. Gustafsson, G. Wüst, C. Magen, D. D. O'Regan, J. W. Luo, B. Ketterer, S. Conesa-Boj, A. V. Kuhlmann, J. Houel, E. Russo-Averchi, J. R. Morante, M. Cantoni, N. Marzari, J. Arbiol, A. Zunger, R. J. Warburton, and A. Fontcuberta i Morral. *Self-assembled quantum dots in a nanowire system for quantum photonics*. *Nature materials* 12(5): pp. 439–444. 2013.

[291] A. Laucht, S. Pütz, T. Günthner, N. Hauke, R. Saive, S. Frédérick, M. Bichler, M.-C. Amann, A. W. Holleitner, M. Kaniber, and J. J. Finley. *A Waveguide-Coupled On-Chip Single-Photon Source*. *Phys. Rev. X* 2: p. 011014. 2012.

[292] J. M. Gérard, B. Sermage, B. Gayral, B. Legrand, E. Costard, and V. Thierry-Mieg. *Enhanced Spontaneous Emission by Quantum Boxes in a Monolithic Optical Microcavity*. *Phys. Rev. Lett.* 81: pp. 1110–1113. 1998.

[293] G. S. Solomon, M. Pelton, and Y. Yamamoto. *Single-mode Spontaneous Emission from a Single Quantum Dot in a Three-Dimensional Microcavity*. *Phys. Rev. Lett.* 86: pp. 3903–3906. 2001.

[294] J. P. Reithmaier, G. Sek, A. Löffler, C. Hofmann, S. Kuhn, S. Reitzenstein, L. V. Keldysh, V. D. Kulakovskii, T. L. Reinecke, and A. Forchel. *Strong coupling in a single quantum dot-semiconductor microcavity system*. *Nature* 432(7014): pp. 197–200. 2004.

[295] M. Davanço, M. T. Rakher, D. Schuh, A. Badolato, and K. Srinivasan. *A circular dielectric grating for vertical extraction of single quantum dot emission*. *Applied Physics Letters* 99(4): p. 041102. 2011.

[296] F. Liu, A. J. Brash, J. O'Hara, L. M. P. P. Martins, C. L. Phillips, R. J. Coles, B. Royall, E. Clarke, C. Bentham, N. Prtljaga, I. E. Itskevich, L. R. Wilson, M. S. Skolnick, and A. M. Fox. *High Purcell factor generation of indistinguishable on-chip single photons*. *Nature Nanotechnology* 13(9): pp. 835–840. 2018.

[297] T. Yoshie, A. Scherer, J. Hendrickson, G. Khitrova, H. M. Gibbs, G. Rupper, C. Ell, O. B. Shchekin, and D. G. Deppe. *Vacuum Rabi splitting with a single quantum dot in a photonic crystal nanocavity*. *Nature* 432(7014): pp. 200–203. 2004.

[298] D. Englund, D. Fattal, E. Waks, G. Solomon, B. Zhang, T. Nakaoka, Y. Arakawa, Y. Yamamoto, and J. Vu čkovi ć. *Controlling the Spontaneous Emission Rate of Single Quantum Dots in a Two-Dimensional Photonic Crystal*. *Phys. Rev. Lett.* 95: p. 013904. 2005.

[299] T. Lund-Hansen, S. Stobbe, B. Julsgaard, H. Thyrrestrup, T. Sünner, M. Kamp, A. Forchel, and P. Lodahl. *Experimental Realization of Highly Efficient Broadband Coupling of Single Quantum Dots to a Photonic Crystal Waveguide*. *Phys. Rev. Lett.* 101: p. 113903. 2008.

[300] D. H. Ahn, Y. D. Jang, J. S. Baek, C. Schneider, S. Höfling, and D. Lee. *A broad-band planar-microcavity quantum-dot single-photon source with a solid immersion lens*. *Applied Physics Letters* 118(17): p. 174001. 2021.

[301] M. E. Reimer, G. Bulgarini, N. Akopian, M. Hocevar, M. B. Bavinck, M. A. Verheijen, E. P. A. M. Bakkers, L. P. Kouwenhoven, and V. Zwiller. *Bright single-photon sources in bottom-up tailored nanowires*. *Nature communications* 3: p. 737. 2012.

[302] K. H. Nielsen, Y. Wang, E. C. R. Deacon, P. I. Sund, Z. Liu, S. Scholz, A. D. Wieck, A. Ludwig, L. Midolo, A. S. Sørensen, S. Paesani, and P. Lodahl. *Programmable nonlinear quantum photonic circuits*. *Nature Communications* 16(1): 2025.

[303] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade. *Photonic Crystals: Molding the Flow of Light*. Princeton University Press, Princeton, 2008.

[304] C. Sauvan, J. P. Hugonin, I. S. Maksymov, and P. Lalanne. *Theory of the Spontaneous Optical Emission of Nanosize Photonic and Plasmon Resonators*. *Phys. Rev. Lett.* 110: p. 237401. 2013.

[305] K. Hennessy, A. Badolato, M. Winger, D. Gerace, M. Atatüre, S. Gulde, S. Fält, E. L. Hu, and A. Imamoğlu. *Quantum nature of a strongly coupled single quantum dot-cavity system*. *Nature* 445(7130): pp. 896–899. 2007.

[306] K. Srinivasan and O. Painter. *Linear and nonlinear optical spectroscopy of a strongly coupled microdisk-quantum dot system*. *Nature* 450(7171): pp. 862–865. 2007.

[307] D. Najer, I. Söllner, P. Sekatski, V. Dolique, M. C. Löbl, D. Riedel, R. Schott, S. Starosielec, S. R. Valentin, A. D. Wieck, N. Sangouard, A. Ludwig, and R. J. Warburton. *A gated quantum dot strongly coupled to an optical microcavity*. *Nature* 575(7784): pp. 622–627. 2019.

[308] O. Gazzano, S. Michaelis de Vasconcellos, C. Arnold, A. Nowak, E. Galopin, I. Sagnes, L. Lanco, A. Lemaître, and P. Senellart. *Bright solid-state sources of indistinguishable single photons*. *Nature communications* 4: p. 1425. 2013.

[309] B. Gayral, J. M. Gérard, B. Legrand, E. Costard, and V. Thierry-Mieg. *Optical study of GaAs/AlAs pillar microcavities with elliptical cross section*. *Applied Physics Letters* 72(12): pp. 1421–1423. 1998.

[310] E. Moreau, I. Robert, J. M. Gérard, I. Abram, L. Manin, and V. Thierry-Mieg. *Single-mode solid-state single photon source based on isolated quantum dots in pillar microcavities*. *Applied Physics Letters* 79(18): pp. 2865–2867. 2001.

[311] B. Yao, R. Su, Y. Wei, Z. Liu, T. Zhao, and J. Liu. *Design for Hybrid Circular Bragg Gratings for a Highly Efficient Quantum-Dot Single-Photon Source*. *Journal of the Korean Physical Society* 73(10): pp. 1502–1505. 2018.

[312] M. B. Rota, T. M. Krieger, Q. Buchinger, M. Beccaceci, J. Neuwirth, H. Huet, N. Horová, G. Lovicu, G. Ronco, S. F. Da Covre Silva, G. Pettinari, M. Moczała-Dusanowska, C. Kohlberger, S. Manna, S. Stroj, J. Freund, X. Yuan, C. Schneider, M. Ježek, S. Höfling, F. Basso Basset, T. Huber-Loyola, A. Rastelli, and R. Trotta. *A source of entangled photons based on a cavity-enhanced and strain-tuned GaAs quantum dot*. *eLight* 4(1): p. 13. 2024.

[313] G. Peniakov, Q. Buchinger, M. Helal, S. Betzold, Y. Reum, M. B. Rota, G. Ronco, M. Beccaceci, T. M. Krieger, S. F. C. da Silva, A. Rastelli, R. Trotta, A. Pfenning, S. Höfling, and T. Huber-Loyola. *Polarized and Unpolarized Emission from a Single Emitter in a Bullseye Resonator*. *Laser & Photonics Reviews* 18(4): p. 2300835. 2024.

[314] A. Laneve, M. B. Rota, F. Basso Basset, M. Beccaceci, V. Villari, T. Oberleitner, Y. Reum, T. M. Krieger, Q. Buchinger, R. Prasad, S. F. C. Da Silva, A. Pfenning, S. Stroj, S. Höfling, A. Rastelli, T. Huber-Loyola, and R. Trotta. *Wavevector-resolved polarization entanglement from radiative cascades*. *Nature communications* 16(1): p. 6209. 2025.

[315] A. Muller, E. B. Flagg, M. Metcalfe, J. Lawall, and G. S. Solomon. *Coupling an epitaxial quantum dot to a fiber-based external-mirror microcavity*. *Applied Physics Letters* 95(17): p. 173101. 2009.

[316] R. J. Barbour, P. A. Dalgarno, A. Curran, K. M. Nowak, H. J. Baker, D. R. Hall, N. G. Stoltz, P. M. Petroff, and R. J. Warburton. *A tunable microcavity*. *Journal of Applied Physics* 110(5): 2011.

[317] Y. Colombe, T. Steinmetz, G. Dubois, F. Linke, D. Hunger, and J. Reichel. *Strong atom-field coupling for Bose-Einstein condensates in an optical cavity on a chip*. *Nature* 450(7167): pp. 272–276. 2007.

[318] L. Bremer, S. Rodt, and S. Reitzenstein. *Fiber-coupled quantum light sources based on solid-state quantum emitters*. *Materials for Quantum Technology* 2(4): p. 042002. 2022.

[319] G. Shambat, J. Provine, K. Rivoire, T. Sarmiento, J. Harris, and J. Vučković. *Optical fiber tips functionalized with semiconductor photonic crystal cavities*. *Applied Physics Letters* 99(19): p. 191102. 2011.

[320] H. Sasakura, X. Liu, S. Odashima, H. Kumano, S. Muto, and I. Suemune. *Fiber-Based Bidirectional Solid-State Single-Photon Emitter Based on Semiconductor Quantum Dot*. *Applied Physics Express* 6(6): p. 065203. 2013.

[321] O. Ortiz, F. Pastier, A. Rodriguez, Priya, A. Lemaitre, C. Gomez-Carbonell, I. Sagnes, A. Harouri, P. Senellart, V. Giesz, M. Esmann, and N. D. Lanzillotti-Kimura. *Fiber-integrated microcavities for efficient generation of coherent acoustic phonons*. *Applied Physics Letters* 117(18): p. 183102. 2020.

[322] H. Kumano, T. Harada, I. Suemune, H. Nakajima, T. Kuroda, T. Mano, K. Sakoda, S. Odashima, and H. Sasakura. *Stable and efficient collection of single photons emitted from a semiconductor quantum dot into a single-mode optical fiber*. *Applied Physics Express* 9(3): p. 032801. 2016.

[323] L. Bremer, K. Weber, S. Fischbach, S. Thiele, M. Schmidt, A. Kaganskiy, S. Rodt, A. Herkommer, M. Sartison, S. L. Portalupi, P. Michler, H. Giessen, and S. Reitzenstein. *Quantum dot single-photon emission coupled into single-mode fibers with 3D printed micro-objectives*. *APL Photonics* 5(10): p. 106101. 2020.

[324] B.-H. Ahn, C.-M. Lee, H.-J. Lim, T. W. Schlereth, M. Kamp, S. Höfling, and Y.-H. Lee. *Direct fiber-coupled single photon source based on a photonic crystal waveguide*. *Applied Physics Letters* 107(8): p. 081113. 2015.

[325] R. S. Daveau, K. C. Balram, T. Pregnolato, J. Liu, E. H. Lee, J. D. Song, V. Verma, R. Mirin, S. W. Nam, L. Midolo, S. Stobbe, K. Srinivasan, and P. Lodahl. *Efficient fiber-coupled single-photon source based on quantum dots in a photonic-crystal waveguide*. *Optica* 4(2): pp. 178–184. 2017.

[326] C.-M. Lee, M. A. Buyukkaya, S. Aghaeimeibodi, A. Karasahin, C. J. K. Richardson, and E. Waks. *A fiber-integrated nanobeam single photon source emitting at telecom wavelengths*. *Applied Physics Letters* 114(17): p. 171101. 2019.

[327] B. Zeng, C. De-Eknamkul, D. Assumpcao, D. Renaud, Z. Wang, D. Riedel, J. Ha, C. Robens, D. Levonian, M. Lukin, R. Riedinger, M. Bhaskar, D. Sukachev, M. Loncar, and B. Machielse. *Cryogenic packaging of nanophotonic devices with a low coupling loss <1 dB*. *Applied Physics Letters* 123(16): p. 161106. 2023.

[328] W. B. Jeon, J. S. Moon, K.-Y. Kim, Y.-H. Ko, C. J. K. Richardson, E. Waks, and J.-H. Kim. *Plug-and-Play Single-Photon Devices with Efficient Fiber-Quantum Dot Interface*. *Advanced Quantum Technologies* 5(10): p. 2200022. 2022.

[329] Y. Chen, S. Li, X. Shang, X. Su, H. Hao, J. Shen, Y. Zhang, H. Ni, Y. Ding, and Z. Niu. *Fiber coupled high count-rate single-photon generated from InAs quantum dots*. *Journal of Semiconductors* 42(7): p. 072901. 2021.

[330] N. Margaria, F. Pastier, T. Bennour, M. Billard, E. Ivanov, W. Hease, P. Stepanov, A. F. Adiyatullin, R. Singla, M. Pont, M. Descampeaux, A. Bernard, A. Pishchagin, M. Morassi, A. Lemaître, T. Volz, V. Giesz, N. Somaschi, N. Maring, S. Boissier, T. H. Au, and P. Senellart. *Efficient fibre-pigtailed source of indistinguishable single photons*. *Nature communications* 16(1): p. 7553. 2025.

[331] M. Paul, F. Olbrich, J. Höschele, S. Schreier, J. Kettler, S. L. Portalupi, M. Jetter, and P. Michler. *Single-photon emission at 1.55 μ m from MOVPE-grown InAs quantum dots on InGaAs/GaAs metamorphic buffers*. *Applied Physics Letters* 111(3): p. 033102. 2017.

[332] C. Carmesin, F. Olbrich, T. Mehrtens, M. Florian, S. Michael, S. Schreier, C. Nawrath, M. Paul, J. Höschele, B. Gerken, J. Kettler, S. L. Portalupi, M. Jetter, P. Michler, A. Rosenauer, and F. Jahnke. *Structural and optical properties of InAs/(In)GaAs/GaAs quantum dots with single-photon emission in the telecom C-band up to 77 K*. *Physical Review B* 98(12): p. 125407. 2018.

[333] Ł. Dusanowski, C. Nawrath, S. L. Portalupi, M. Jetter, T. Huber, S. Klembt, P. Michler, and S. Höfling. *Optical charge injection and coherent control of a quantum-dot spin-qubit emitting at telecom wavelengths*. *Nat. Commun.* 13(1): p. 748. 2022.

[334] M. Scheibner, T. Schmidt, L. Worschech, A. Forchel, G. Bacher, T. Passow, and D. Hommel. *Superradiance of quantum dots*. *Nature Physics* 3(2): pp. 106–110. 2007.

[335] N. Hauser, M. J. Bayerbach, S. E. D'Aurelio, R. Weber, M. Santandrea, S. P. Kumar, I. Dhand, and S. Barz. *Boosted Bell-state measurements for photonic quantum computation*. *npj Quantum Information* 11(1): 2025.

[336] S. Liu, C. Gustin, H. Liu, X. Li, Y. Yu, H. Ni, Z. Niu, S. Hughes, X. Wang, and J. Liu. *Dynamic resonance fluorescence in solid-state cavity quantum electrodynamics*. *Nature Photonics* 18(4): pp. 318–324. 2024.

[337] K. Boos, S. K. Kim, T. Bracht, F. Sbresny, J. M. Kaspari, M. Cygorek, H. Riedl, F. W. Bopp, W. Rauhaus, C. Calcagno, J. J. Finley, D. E. Reiter, and K. Müller. *Signatures of Dynamically Dressed States*. *Phys. Rev. Lett.* 132: p. 053602. 2024.

[338] L. M. Hansen, F. Giorgino, L. Jehle, L. Carosini, J. C. L. Carreño, I. Arrazola, P. Walther, and J. C. Loredo. *Non-classical excitation of a solid-state quantum emitter*, 2024.

[339] P. Michler. *Single Semiconductor Quantum Dots*. Springer Berlin Heidelberg, Berlin, Heidelberg. 2009.

[340] K. A. Fischer, K. Müller, A. Rundquist, T. Sarmiento, A. Y. Piggott, Y. Kelaita, C. Dory, K. G. Lagoudakis, and J. Vučković. *Self-homodyne measurement of a dynamic Mollow triplet in the solid state*. *Nature Photonics* 10(3): pp. 163–166. 2016.

[341] E. Martín-López, A. Laing, T. Lawson, R. Alvarez, X.-Q. Zhou, and J. L. O'Brien. *Experimental realization of Shor's quantum factoring algorithm using qubit recycling*. *Nature Photonics* 6(11): pp. 773–776. 2012.

[342] V. Saggio, B. E. Asenbeck, A. Hamann, T. Strömberg, P. Schiansky, V. Dunjko, N. Friis, N. C. Harris, M. Hochberg, D. Englund, S. Wölk, H. J. Briegel, and P. Walther. *Experimental quantum speed-up in reinforcement learning agents*. *Nature* 591(7849): pp. 229–233. 2021.

[343] F. Centrone, N. Kumar, E. Diamanti, and I. Kerenidis. *Experimental demonstration of quantum advantage for NP verification with limited information*. *Nat. Commun.* 12(1): 2021.

[344] R. Bedington, J.-M. Arrazola, and A. Ling. *Progress in satellite quantum key distribution*. *npj Quantum Inf.* 3: p. 30. 2017.

[345] A. Boaron, G. Boso, D. Rusca, C. Vulliez, C. Autebert, M. Caloz, M. Perrenoud, G. Gras, F. Bussières, M.-J. Li, D. Nolan, A. Martin, and H. Zbinden. *Secure Quantum Key Distribution over 421 km of Optical Fiber*. *Phys. Rev. Lett.* 121: p. 190502. 2018.

[346] S. Neves, V. Yacoub, U. Chabaud, M. Bozzio, I. Kerenidis, and E. Diamanti. *Experimental cheat-sensitive quantum weak coin flipping*. *Nat. Commun.* 14: p. 1855. 2023.

[347] G. Berlín, G. Brassard, F. Bussières, N. Godbout, J. A. Slater, and W. Tittel. *Experimental loss-tolerant quantum coin flipping*. *Nat. Commun.* 2: p. 561. 2011.

[348] A. Pappa, P. Jouguet, T. Lawson, A. Chailloux, M. Legré, P. Trinkler, I. Kerenidis, and E. Diamanti. *Experimental plug and play quantum coin flipping*. *Nat. Commun.* 5: p. 3717. 2014.

[349] P. Schiansky, J. Kalb, E. Sztatecsny, M.-C. Roehsner, T. Guggemos, A. Trenti, M. Bozzio, and P. Walther. *Demonstration of quantum-digital payments*. *Nat. Commun.* 14(1): 2023.

[350] A. Kent, D. Lowndes, D. Pitalúa-García, and J. Rarity. *Practical quantum tokens without quantum memories and experimental tests*. *npj Quantum Inf.* 8(1): 2022.

[351] J.-Y. Guan, J.-M. Arrazola, R. Amiri, W. Zhang, H. Li, L. You, Z. Wang, Q. Zhang, and J.-W. Pan. *Experimental preparation and verification of quantum money*. *Phys. Rev. A* 97: p. 032338. 2018.

[352] S. K. Ng N. Huei Y.and Joshi, C. Chen Ming, C. Kurtsiefer, and S. Wehner. *Experimental implementation of bit commitment in the noisy-storage model*. *Nat. Commun.* 3(1): 2012.

[353] T. Lunghi, J. Kaniewski, F. Bussières, R. Houlmann, M. Tomamichel, A. Kent, N. Gisin, S. Wehner, and H. Zbinden. *Experimental Bit Commitment Based on Quantum Communication and Special Relativity*. *Phys. Rev. Lett.* 111: p. 180504. 2013.

[354] H. J. Kimble. *The quantum internet*. *Nature* 453(7198): pp. 1023–1030. 2008.

[355] J.-G. Ren, P. Xu, H.-L. Yong, L. Zhang, S.-K. Liao, J. Yin, W.-Y. Liu, W.-Q. Cai, M. Yang, L. Li, K.-X. Yang, X. Han, Y.-Q. Yao, J. Li, H.-Y. Wu, S. Wan, L. Liu, D.-Q. Liu, Y.-W. Kuang, Z.-P. He, P. Shang, C. Guo, R.-H. Zheng, K. Tian, Z.-C. Zhu, N.-L. Liu, C.-Y. Lu, R. Shu, Y.-A. Chen, C.-Z. Peng, J.-Y. Wang, and J.-W. Pan. *Ground-to-satellite quantum teleportation*. *Nature* 549(7670): pp. 70–73. 2017.

[356] X.-S. Ma, T. Herbst, T. Scheidl, D. Wang, S. Kropatschek, W. Naylor, B. Wittmann, A. Mech, J. Kofler, E. Anisimova, V. Makarov, T. Jennewein, R. Ursin, and A. Zeilinger. *Quantum teleportation over 143 kilometres using active feed-forward*. *Nature* 489(7415): pp. 269–273. 2012.

[357] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson. *Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres*. *Nature* 526(7575): pp. 682–686. 2015.

[358] S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe. *An elementary quantum network of single atoms in optical cavities*. *Nature* 484(7393): pp. 195–200. 2012.

[359] F. Basso Basset, M. B. Rota, C. Schimpf, D. Tedeschi, K. D. Zeuner, S. F. Covre da Silva, M. Reindl, V. Zwiller, K. D. Jöns, A. Rastelli, and R. Trotta. *Entanglement Swapping with Photons Generated on Demand by a Quantum Dot*. *Phys. Rev. Lett.* 123(16): p. 160501. 2019.

[360] M. Zopf, R. Keil, Y. Chen, J. Yang, D. Chen, F. Ding, and O. G. Schmidt. *Entanglement Swapping with Semiconductor-Generated Photons Violates Bell's Inequality*. *Phys. Rev. Lett.* 123(16): p. 160502. 2019.

[361] M. Anderson, T. Müller, J. Huwer, J. Skiba-Szymanska, A. B. Krysa, R. M. Stevenson, J. Heffernan, D. A. Ritchie, and A. J. Shields. *Quantum teleportation using highly coherent emission from telecom C-band quantum dots*. *npj Quantum Inf.* 6: p. 14. 2020.

[362] F. B. Basset, F. Salusti, L. Schweickert, M. B. Rota, D. Tedeschi, S. F. C. da Silva, E. Roccia, V. Zwiller, K. D. Jöns, A. Rastelli, and R. Trotta. *Quantum teleportation with imperfect quantum dots*. *npj Quantum Inf.* 7: p. 7. 2021.

[363] S. K. Liao, H. L. Yong, C. Liu, G. L. Shentu, D. D. Li, J. Lin, H. Dai, S. Q. Zhao, B. Li, J. Y. Guan, W. Chen, Y. H. Gong, Y. Li, Z. H. Lin, G. S. Pan, J. S. Pelc, M. M. Fejer, W. Z. Zhang, W. Y. Liu, J. Yin, J. G. Ren, X. B. Wang, Q. Zhang, C. Z. Peng, and J. W. Pan. *Long-distance free-space quantum key distribution in daylight towards inter-satellite communication*. *Nat. Photon.* 11(8): pp. 509–513. 2017.

[364] J. Wang, F. Sciarrino, A. Laing, and M. G. Thompson. *Integrated photonic quantum technologies*. *Nat. Photon.* 14(5): pp. 273–284. 2020.

[365] T. Miyazawa, K. Takemoto, Y. Nambu, S. Miki, T. Yamashita, H. Terai, M. Fujiwara, M. Sasaki, Y. Sakuma, M. Takatsu, T. Yamamoto, and Y. Arakawa. *Single-photon emission at 1.5m from an InAs/InP quantum dot with highly suppressed multi-photon emission probabilities*. *Appl. Phys. Lett.* 109(13): p. 132106. 2016.

[366] M. Anderson, T. Müller, J. Skiba-Szymanska, A. B. Krysa, J. Huwer, R. M. Stevenson, J. Heffernan, D. A. Ritchie, and A. J. Shields. *Gigahertz-Clocked Teleportation of Time-Bin Qubits with a Quantum Dot in the Telecommunication C Band*. *Phys. Rev. Applied* 13: p. 054052. 2020.

[367] G. Shooter, Z.-H. Xiang, J. R. A. Müller, J. Skiba-Szymanska, J. Huwer, J. Griffiths, T. Mitchell, M. Anderson, T. Müller, A. B. Krysa, R. Mark Stevenson, J. Heffernan, D. A. Ritchie, and A. J. Shields. *1GHz clocked distribution of electrically generated entangled photon pairs*. *Optics Express* 28(24): p. 36838. 2020.

[368] T. Lettner, S. Gyger, K. D. Zeuner, L. Schweickert, S. Steinhauer, C. Reuterskiöld Hedlund, S. Stroj, A. Rastelli, M. Hammar, R. Trotta, K. D. Jöns, and V. Zwiller. *Strain-Controlled Quantum Dot Fine Structure for Entangled Photon Generation at 1550 nm*. *Nano Letters* 21(24): pp. 10501–10506. 2021.

[369] T. van Leent, M. Bock, R. Garthoff, K. Redeker, W. Zhang, T. Bauer, W. Rosenfeld, C. Becher, and H. Weinfurter. *Long-Distance Distribution of Atom-Photon Entanglement at Telecom Wavelength*. *Phys. Rev. Lett.* 124: p. 010510. 2020.

[370] X. Ma. *Quantum cryptography: theory and practice*, 2008.

[371] M. Bozzio, A. Orieux, L. T. Vidarte, I. Zaquine, I. Kerenidis, and E. Diamanti. *Experimental investigation of practical unforgeable quantum money*. *npj Quantum Inf.* 4: p. 5. 2018.

[372] X.-B. Wang. *Beating the Photon-Number-Splitting Attack in Practical Quantum Cryptography*. *Phys. Rev. Lett.* 94: p. 230503. 2005.

[373] S. Ates, I. Agha, A. Gulinatti, I. Rech, A. Badolato, and K. Srinivasan. *Improving the performance of bright quantum dot single photon sources using temporal filtering via amplitude modulation*. *Scientific Reports* 3(1): 2013.

[374] T. Kupko, M. von Helversen, L. Rickert, J.-H. Schulze, A. Strittmatter, M. Gschrey, S. Rodt, S. Reitzenstein, and T. Heindel. *Tools for the performance optimization of single-photon quantum key distribution*. *npj Quantum Information* 6(1): 2020.

[375] D. Cogan, Z.-E. Su, O. Kenneth, and D. Gershoni. *Deterministic generation of indistinguishable photons in a cluster state*. *Nature Photonics* 17(4): pp. 324–329. 2023.

[376] A. K. Nowak, S. L. Portalupi, V. Giesz, O. Gazzano, C. Dal Savio, P.-F. Braun, K. Karrai, C. Arnold, L. Lanco, I. Sagnes, A. Lemaître, and P. Senellart. *Deterministic and electrically tunable bright single-photon source*. *Nature Communications* 5: p. 3240. 2014.

[377] A. J. Bennett, R. B. Patel, J. Skiba-Szymanska, C. A. Nicoll, I. Farrer, D. A. Ritchie, and A. J. Shields. *Giant Stark effect in the emission of single semiconductor quantum dots*. *Applied Physics Letters* 97(3): 2010.

[378] P. Tighineanu, C. L. Dreeßen, C. Flindt, P. Lodahl, and A. S. Sørensen. *Phonon Decoherence of Quantum Dots in Photonic Structures: Broadening of the Zero-Phonon Line and the Role of Dimensionality*. *Phys. Rev. Lett.* 120: p. 257401. 2018.

[379] J. Liu, K. Konthasinghe, M. Davan ço, J. Lawall, V. Anant, V. Verma, R. Mirin, S. W. Nam, J. D. Song, B. Ma, Z. S. Chen, H. Q. Ni, Z. C. Niu, and K. Srinivasan. *Single Self-Assembled* InAs/GaAs *Quantum Dots in Photonic Nanostructures: The Role of Nanofabrication*. *Phys. Rev. Appl.* 9: p. 064019. 2018.

[380] R. Stockill, C. Le Gall, C. Matthiesen, L. Huthmacher, E. Clarke, M. Hugues, and M. Atatüre. *Quantum dot spin coherence governed by a strained nuclear environment*. *Nature Communications* 7(1): p. 12745. 2016.

[381] J. Yoneda, K. Takeda, T. Otsuka, T. Nakajima, M. R. Delbecq, G. Allison, T. Honda, T. Kodera, S. Oda, Y. Hoshi, N. Usami, K. M. Itoh, and S. Tarucha. *A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%*. *Nature Nanotechnology* 13(2): pp. 102–106. 2018.

[382] C. Cohen-Tannoudji and S. Reynaud. *Dressed-atom description of resonance fluorescence and absorption spectra of a multi-level atom in an intense laser beam*. *Journal of Physics B: Atomic and Molecular Physics* 10(3): pp. 345–363. 1977.

[383] T. Unold, K. Mueller, C. Lienau, T. Elsaesser, and A. D. Wieck. *Optical Stark Effect in a Quantum Dot: Ultrafast Control of Single Exciton Polarizations*. *Phys. Rev. Lett.* 92: p. 157401. 2004.

[384] A. Muller, W. Fang, J. Lawal, and G. S. Solomon. *Emission Spectrum of a Dressed Exciton-Biexciton Complex in a Semiconductor Quantum Dot*. *Phys. Rev. Lett.* 101: p. 027401. 2008.

[385] A. Muller, W. Fang, J. Lawall, and G. S. Solomon. *Creating Polarization-Entangled Photon Pairs from a Semiconductor Quantum Dot Using the Optical Stark Effect*. *Phys. Rev. Lett.* 103: p. 217402. 2009.

[386] A. Ripin, R. Peng, X. Zhang, S. Chakravarthi, M. He, X. Xu, K.-M. Fu, T. Cao, and M. Li. *Tunable phononic coupling in excitonic quantum emitters*. *Nature Nanotechnology* 18(9): pp. 1020–1026. 2023.

[387] A. Vagov, M. D. Croitoru, M. Glässl, V. M. Axt, and T. Kuhn. *Real-time path integrals for quantum dots: Quantum dissipative dynamics with superohmic environment coupling*. *Phys. Rev. B* 83(9): 2011.

[388] R. Impagliazzo and M. Luby. *One-way functions are essential for complexity based cryptography*. In: *30th Annual Symposium on Foundations of Computer Science*. 1989, pp. 230–235.

[389] M. Naor. *Bit commitment using pseudorandomness*. *Journal of Cryptology* 4(2): pp. 151–158. 1991.

[390] I. Haitner, M.-H. Nguyen, S. J. Ong, O. Reingold, and S. Vadhan. *Statistically Hiding Commitments and Statistical Zero-Knowledge Arguments from Any One-Way Function*. *SIAM Journal on Computing* 39(3): pp. 1153–1218. 2009.

[391] A. Kent. *Unconditionally Secure Bit Commitment*. *Phys. Rev. Lett.* 83: pp. 1447–1450. 1999.

[392] T. Lunghi, J. Kaniewski, F. Bussières, R. Houlmann, M. Tomamichel, S. Wehner, and H. Zbinden. *Practical Relativistic Bit Commitment*. *Phys. Rev. Lett.* 115: p. 030502. 2015.

[393] M. Ben-Or, S. Goldwasser, J. Kilian, and A. Wigderson. *Multi-prover interactive proofs: how to remove intractability assumptions*. In: *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*. STOC '88. Chicago, Illinois, USA: Association for Computing Machinery, 1988, pp. 113–131.

[394] P. Alikhani, N. Brunner, C. Crépeau, S. Designolle, R. Houlmann, W. Shi, N. Yang, and H. Zbinden. *Experimental relativistic zero-knowledge proofs*. *Nature* 599(7883): pp. 47–50. 2021.

[395] T. Strömberg, P. Schiansky, and P. Walther. *Prescriptive method for fiber polarization compensation in two bases*. *Appl. Opt.* 63(7): pp. 1822–1827. 2024.

[396] R. M. Karp. *Reducibility Among Combinatorial Problems*. In: *50 Years of Integer Programming 1958-2008*. Springer Berlin Heidelberg, pp. 219–241. 2009.

[397] M. R. Garey and D. S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., USA. 1990.

[398] S. Arora and S. Safra. *Probabilistic checking of proofs: a new characterization of NP*. *J. ACM* 45(1): pp. 70–122. 1998.

[399] I. Dinur. *The PCP theorem by gap amplification*. *J. ACM* 54(3): 12–es. 2007.

[400] J. Håstad. *Some optimal inapproximability results*. *J. ACM* 48(4): pp. 798–859. 2001.

[401] V. V. Vazirani. *Approximation Algorithms*. *J. Artif. Int. Res.*: 2003.

[402] U. Feige. *A threshold of ln n for approximating set cover*. *J. ACM* 45(4): pp. 634–652. 1998.

[403] A. Flamini, G. Sciarretta, M. Scuro, A. Sharif, A. Tomasi, and S. Ranise. *On cryptographic mechanisms for the selective disclosure of verifiable credentials*. *Journal of Information Security and Applications* 83: p. 103789. 2024.

[404] Š. B. Ramić, E. Cogo, I. Prazina, E. Cogo, M. Turkanović, R. T. Mulahasanović, and S. Mrdović. *Selective disclosure in digital credentials: A review*. *ICT Express* 10(4): pp. 916–934. 2024.

[405] J. E. Holt and K. E. Seamons. *Selective disclosure credential sets*, 2002.

[406] N. O. Tippenhauer, C. Pöpper, K. B. Rasmussen, and S. Capkun. *On the Requirements for Successful GPS Spoofing Attacks*. *Proc. of the 18th ACM Conference on Computer and Communications Security*: pp. 75–86. 2011.

[407] C. Schaffner, B. Terhal, and S. Wehner. *Robust Cryptography in the Noisy-Quantum-Storage Model*. *Quantum Info. Comput.* 9(11): pp. 963–996. 2009.

[408] C. Schaffner. *Simple protocols for oblivious transfer and secure identification in the noisy-quantum-storage model*. *Phys. Rev. A* 82: p. 032308. 2010.

[409] C. Erven, N. Ng, N. Gigov, R. Laflamme, S. Wehner, and G. Weihs. *An experimental implementation of oblivious transfer in the noisy storage model*. *Nat. Commun.* 5(1): 2014.

[410] D.-C. Liu, P.-Y. Li, T.-X. Zhu, L. Zheng, J.-Y. Huang, Z.-Q. Zhou, C.-F. Li, and G.-C. Guo. *On-Demand Storage of Photonic Qubits at Telecom Wavelengths*. *Phys. Rev. Lett.* 129: p. 210501. 2022.

[411] P. Vernaz-Gris, K. Huang, M. Cao, A. Sheremet, and J. Laurat. *Highly-efficient quantum memory for polarization qubits in a spatially-multiplexed cold atomic ensemble*. *Nat. Comm.* 9: p. 363. 2018.

[412] C. Laplane, P. Jobez, J. Etesse, N. Timoney, N. Gisin, and M. Afzelius. *Multiplexed on-demand storage of polarization qubits in a crystal*. *New J. Phys.* 18(1): p. 013006. 2016.

[413] F. Grünenfelder, A. Boaron, D. Rusca, A. Martin, and H. Zbinden. *Performance and security of 5GHz repetition rate polarization-based quantum key distribution*. *Applied Physics Letters* 117(14): p. 144003. 2020.

[414] W. Li, L. Zhang, H. Tan, Y. Lu, S.-K. Liao, J. Huang, H. Li, Z. Wang, H.-K. Mao, B. Yan, Q. Li, Y. Liu, Q. Zhang, C.-Z. Peng, L. You, F. Xu, and J.-W. Pan. *High-rate quantum key distribution exceeding 110 Mb s–1*. *Nature Photonics* 17(5): pp. 416–421. 2023.

[415] A. Javadi, I. Söllner, M. Arcari, S. L. Hansen, L. Midolo, S. Mahmoodian, G. Kiršanskė, T. Pregnolato, E. H. Lee, J. D. Song, S. Stobbe, and P. Lodahl. *Single-photon non-linear optics with a quantum dot in a waveguide*. *Nature Communications* 6(1): 2015.

[416] H. E. Dyte, S. Manna, S. F. Covre da Silva, A. Rastelli, and E. A. Chekhovich. *Storing quantum coherence in a quantum dot nuclear spin ensemble for over 100 milliseconds*. *Nature Communications* 17(1): 2025.

[417] H. Huet, P. R. Ramesh, S. C. Wein, N. Coste, P. Hilaire, N. Somaschi, M. Morassi, A. Lemaître, I. Sagnes, M. F. Doty, O. Krebs, L. Lanco, D. A. Fioretto, and P. Senellart. *Deterministic and reconfigurable graph state generation with a single solid-state quantum emitter*. *Nature Communications* 16(1): 2025.

[418] H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, and M. Van den Nest. *Measurement-based quantum computation*. *Nature Physics* 5(1): pp. 19–26. 2009.

[419] N. Coste, M. Gundin, D. A. Fioretto, S. E. Thomas, C. Millet, E. Mehdi, N. Somaschi, M. Morassi, M. Pont, A. Lemaître, N. Belabas, O. Krebs, L. Lanco, and P. Senellart. *Probing the dynamics and coherence of a semiconductor hole spin via acoustic phonon-assisted excitation*. *Quantum Science and Technology* 8(2): p. 025021. 2023.

[420] C. Piccinini, A. Paralikis, J. F. Neto, A. A. Madigawa, P. Wyborski, V. Remesh, L. Vannucci, N. Gregersen, and B. Munkhbat. *High-purity and stable single-photon emission in bilayer WSe2 via phonon-assisted excitation*. *Communications Physics* 8(1): 2025.

[421] M. Hazra, M. Rieger, A. Kumar, M. N. Mishuk, C. Cholsuk, K. Sripathy, V. Villafañe, K. Müller, J. J. Finley, and T. Vogl. *Temperature-Dependent Emission Spectroscopy of Quantum Emitters in Hexagonal Boron Nitride*. *ACS Photonics*: 2026.

[422] D. Wigger, R. Schmidt, O. Del Pozo-Zamudio, J. A. Preuß, P. Tonndorf, R. Schneider, P. Steeger, J. Kern, Y. Khodaei, J. Sperling, S. M. de Vasconcellos, R. Bratschitsch, and T. Kuhn. *Phonon-assisted emission and absorption of individual color centers in hexagonal boron nitride*. *2D Materials* 6(3): p. 035006. 2019.

[423] C.-Y. Lu and J.-W. Pan. *Quantum-dot single-photon sources for the quantum internet*. *Nature nanotechnology* 16(12): pp. 1294–1296. 2021.

[424] S. Wehner, D. Elkouss, and R. Hanson. *Quantum internet: A vision for the road ahead*. *Science* 362(6412): 2018.

[425] S. Barz, E. Kashefi, A. Broadbent, J. F. Fitzsimons, A. Zeilinger, and P. Walther. *Demonstration of Blind Quantum Computing*. *Science* 335(6066): pp. 303–308. 2012.

[426] D. Barral, F. J. Cardama, G. Díaz-Camacho, D. Faílde, I. F. Llovo, M. Mussa-Juane, J. Vázquez-Pérez, J. Villasuso, C. Piñeiro, N. Costas, J. C. Pichel, T. F. Pena, and A. Gómez. *Review of Distributed Quantum Computing: From single QPU to High Performance Quantum Computing*. *Computer Science Review* 57: p. 100747. 2025.

[427] D. Awschalom, K. K. Berggren, H. Bernien, S. Bhave, L. D. Carr, P. Davids, S. E. Economou, D. Englund, A. Faraon, M. Fejer, S. Guha, M. V. Gustafsson, E. Hu, L. Jiang, J. Kim, B. Korzh, P. Kumar, P. G. Kwiat, M. Lončar, M. D. Lukin, D. A. Miller, C. Monroe, S. W. Nam, P. Narang, J. S. Orcutt, M. G. Raymer, A. H. Safavi-Naeini, M. Spiropulu, K. Srinivasan, S. Sun, J. Vučković, E. Waks, R. Walsworth, A. M. Weiner, and Z. Zhang. *Development of Quantum Interconnects (QuICs) for Next-Generation Information Technologies*. *PRX Quantum* 2(1): 2021.

[428] C. Löschnauer, J. Mosca Toba, A. Hughes, S. King, M. Weber, R. Srinivas, R. Matt, R. Nourshargh, D. Allcock, C. Ballance, C. Matthiesen, M. Malinowski, and T. Harty. *Scalable, High-Fidelity All-Electronic Control of Trapped-Ion Qubits*. *PRX Quantum* 6: p. 040313. 2025.

[429] S. Zaske, A. Lenhard, C. A. Keßler, J. Kettler, C. Hepp, C. Arend, R. Albrecht, W.-M. Schulz, M. Jetter, P. Michler, and C. Becher. *Visible-to-Telecom Quantum Frequency Conversion of Light from a Single Quantum Emitter*. *Phys. Rev. Lett.* 109: p. 147404. 2012.

[430] A. Singh, Q. Li, S. Liu, Y. Yu, X. Lu, C. Schneider, S. Höfling, J. Lawall, V. Verma, R. Mirin, S. W. Nam, J. Liu, and K. Srinivasan. *Quantum frequency conversion of a quantum dot single-photon source on a nanophotonic chip*. *Optica* 6(5): pp. 563–569. 2019.

[431] M. Cohen, A. Martin, R. Dalidet, F. Pastier, M. Billard, A. Lemaitre, V. Giesz, N. Somaschi, S. Thomas, P. Senellart-Mardon, S. Tanzilli, and L. Labonté. *Compact and efficient quantum frequency conversion of a fiber-pigtailed single-photon source*. *APL Quantum* 3(1): p. 016101. 2026.

[432] K. A. G. Fisher, D. G. England, J.-P. W. MacLean, P. J. Bustard, K. J. Resch, and B. J. Sussman. *Frequency and bandwidth conversion of single photons in a room-temperature diamond quantum memory*. *Nature Communications* 7(1): 2016.

[433] M. Allgaier, V. Ansari, L. Sansoni, C. Eigner, V. Quiring, R. Ricken, G. Harder, B. Brecht, and C. Silberhorn. *Highly efficient frequency conversion with bandwidth compression of quantum light*. *Nature Communications* 8(1): 2017.

[434] F. Sośnicki, M. Mikołajczyk, A. Golestani, and M. Karpiński. *Interface between picosecond and nanosecond quantum light pulses*. *Nature Photonics* 17(9): pp. 761–766. 2023.

[435] R. Uppu, F. T. Pedersen, Y. Wang, C. T. Olesen, C. Papon, X. Zhou, L. Midolo, S. Scholz, A. D. Wieck, A. Ludwig, and P. Lodahl. *Scalable integrated single-photon source*. *Sci. Adv.* 6(50): eabc8268. 2020.

[436] N. H. Y. Ng, M. Berta, and S. Wehner. *Min-entropy uncertainty relation for finite-size cryptography*. *Phys. Rev. A* 86: p. 042315. 2012.

[437] I. Damgard, S. Fehr, R. Renner, L. Salvail, and C. Schaffner. *A Tight High-Order Entropic Quantum Uncertainty Relation with Applications*. *Advances in Cryptology* 4622: pp. 360–378. 2007.

A

APPENDIX: Correlation functions

Correlation functions are indispensable for characterizing quantum light. First formulated by Glauber in 1963 [114, 115], they describe the coherence between electric fields (compare Eq. 2.1 -20)

$$\boldsymbol{E}^{+}(t) = \sum_{\mathbf{k},\mu} \sqrt{\frac{\hbar\omega_{\boldsymbol{k}}}{2\epsilon_0 V}} a_{\boldsymbol{k},\mu}\, e^{i(\omega_{\boldsymbol{k}} t - \boldsymbol{k}\cdot\boldsymbol{r}_0)}\, \boldsymbol{e}_{\boldsymbol{k},\mu} \tag{A.0-1}$$

(and $\boldsymbol{E}^{-}$ as the complex conjugate) for arbitrary order n

$$G^{(n)}(t_1, ..., t_{2n}) = \langle \boldsymbol{E}^{-}(t_1) ... \boldsymbol{E}^{-}(t_n) \boldsymbol{E}^{+}(t_{n+1}) ... \boldsymbol{E}^{+}(t_{2n}) \rangle \,. \tag{A.0-2}$$

Since the photonic creation and annihilation operators do not commute, one has to take care of the ordering of electric fields, where we adopted the so-called normal ordering. Generally, the fields analyzed by means of correlation functions can stem from different light sources, but in this work, we focus on fields emitted by the same source but at different times, to which one might refer as autocorrelation functions. Limiting ourselves to the description of single-mode fields and absorbing the time dependence in the photon operator, we simplify the notation of the following by using $E^{+}(t) \propto a(t)$ and $E^{-}(t) \propto a^{\dagger}(t)$. Then, the first-order autocorrelation function is

$$G^{(1)}(t_1, t_2) = \langle a^{\dagger}(t_1)\, a(t_2) \rangle \,, \tag{A.0-3}$$

which reduces to the photon number operator for $t_2 = t_1$

$$G^{(1)}(t_1, t_1) = \langle n(t_1) \rangle \,. \tag{A.0-4}$$

Next, we define the *normalized* first-order autocorrelation function

$$g^{(1)}(t_1, t_2) = \frac{G^{(1)}(t_1, t_2)}{G^{(1)}(t_1, t_1)} = \frac{G^{(1)}(t_1, t_2)}{\langle n(t_1) \rangle} \,. \tag{A.0-5}$$

For electric fields with stationary statistics (that is, field fluctuations do not change on long time scales), it is convenient to change from absolute times t_1 and t_2 to a reference time $t = t_1$ and a relative delay from this reference time $\tau = t_2 - t_1$ such that

$$G^{(1)}(t, \tau) = \langle a^{\dagger}(t)\, a(t + \tau) \rangle \,. \tag{A.0-6}$$

Since the field fluctuations are assumed time independent we can set $t = 0$ without loss of generality and apply statistical averaging $\langle n(t) \rangle = \bar{n}$ to find

$$G^{(1)}(\tau) = \langle a^{\dagger}(0)\, a(\tau) \rangle \tag{A.0-7}$$

$$g^{(1)}(\tau) = \frac{G^{(1)}(\tau)}{\bar{n}} \,, \tag{A.0-8}$$

The temporal coherence of a photon source is inferred by experimentally measuring $g^{(1)}(\tau)$ as described in Sec. 2.2.2.3.

Similarly, we define the second-order autocorrelation functions

$$G^{(2)}(t_1, t_2) = \langle a^\dagger(t_1)\, a^\dagger(t_2)\, a(t_2)\, a(t_1)\rangle \tag{A.0-9}$$

$$g^{(2)}(t_1, t_2) = \frac{G^{(2)}(t_1, t_2)}{G^{(1)}(t_1, t_1)\, G^{(1)}(t_2, t_2)} = \frac{G^{(2)}(t_1, t_2)}{\bar{n}^2}\,. \tag{A.0-10}$$

Analog to the first-order, we also rewrite in terms of relative delay τ

$$g^{(2)}(\tau) = \frac{\langle a^\dagger(0)\, a^\dagger(\tau)\, a(\tau)\, a(0)\rangle}{\bar{n}^2} \tag{A.0-11}$$

and using the commutator relation $a^\dagger(0)\, a(\tau) = a(\tau)\, a^\dagger(0) - \delta(\tau)$ we simplify Eq. A.0-11 for further use

$$\begin{aligned} g^{(2)}(\tau) &= \frac{1}{\bar{n}^2}\langle a^\dagger(\tau)\left[a(\tau)\, a^\dagger(0) - \delta(\tau)\right] a(0)\rangle \\ &= \frac{1}{\bar{n}^2}\langle a^\dagger(\tau) a(\tau)\, a^\dagger(0)\, a(0) - a^\dagger(\tau)\, a(0)\delta(\tau)\rangle \\ &= \frac{1}{\bar{n}^2}\langle n^2 - a^\dagger(\tau)\, a(0)\delta(\tau)\rangle \end{aligned} \tag{A.0-12}$$

and it follows for zero time delay

$$g^{(2)}(0) = \frac{\langle n(n-1)\rangle}{\bar{n}^2}\,. \tag{A.0-13}$$

B

APPENDIX: Security analysis for limited-knowledge proofs

B.1 Technical preliminaries

B.1.1 Polynomial-time reductions

In our protocol, we rely on a specific problem structure to estimate the amount of information leaked in each verification round. Furthermore, we require a soundness gap to prove security against a dishonest P^*. To obtain this, V applies a chain of reductions R to the public 3-SAT instance ϕ at the start of the protocol, yielding a MAX-E3SAT(5) instance ψ with a guaranteed soundness gap. Here, MAX-E3SAT(5) denotes a restricted version of 3-SAT in which every clause has exactly three literals and each variable appears in exactly five distinct clauses. We now outline this chain of reductions:

Theorem B.1.1 *[400]. Given a 3-CNF Boolean formula ϕ, it is NP-hard to distinguish between the following two cases for any constant $\delta \in (0, \frac{1}{8})$:*

- *(Completeness) ϕ is satisfiable.*
- *(Soundness) No truth assignment can satisfy at least a $(\frac{7}{8}+\delta)$-fraction of the clauses of ϕ.*

This theorem implies that there is a polynomial-time mapping R taking any 3-CNF Boolean formula ϕ to a MAX-3-SAT instance $R(\phi) = \psi_{MaxSAT}$ such that

- $\phi \in SAT \quad \Rightarrow \quad \mathrm{OPT}(\psi) = 1,$
- $\phi \notin SAT \quad \Rightarrow \quad \mathrm{OPT}(\psi) \leq \frac{7}{8} + \delta,$

where $\mathrm{OPT}(\psi)$ is the maximum fraction of clauses of ψ that any assignment can satisfy, i.e. $\mathrm{OPT}(\psi) = \max_{\tilde{s}} f_{\tilde{s}}(\psi)$ where $f_{\tilde{s}}(\psi)$ denotes the fraction of clauses satisfied by assignment $\tilde{s}$.

We then use Theorem 29.11 in [401] to transform the MAX-3-SAT to MAX-3SAT(29), which is a special case of MAX-3-SAT where every variable appears in at most 29 clauses of a 3-CNF Boolean formula:

Theorem B.1.2 *[401]. There is a polynomial-time gap-preserving reduction that transforms an instance ψ_{MaxSAT} of MAX-3-SAT to an instance ψ_{Max29} of MAX-3SAT(29) such that the following is satisfied:*

- $OPT(\psi_{MaxSAT}) = 1 \quad \Rightarrow OPT(\psi_{Max29}) = 1.$
- $OPT(\psi_{MaxSAT}) \leq 1 - \delta' \quad \Rightarrow OPT(\psi_{Max29}) \leq 1 - \frac{\delta'}{43}.$

For our reduction this means that $\delta' = \frac{1}{8} - \delta \Rightarrow 1 - \frac{\delta'}{43} = 1 - \frac{\frac{1}{8}-\delta}{43}$, hence

- $\mathrm{OPT}(\psi_{MaxSAT}) = 1 \quad \Rightarrow \quad \mathrm{OPT}(\psi_{Max29}) = 1.$
- $\mathrm{OPT}(\psi_{MaxSAT}) \leq \frac{7}{8} + \delta \quad \Rightarrow \quad \mathrm{OPT}(\psi_{Max29}) \leq 1 - \frac{\frac{1}{8}-\delta}{43}.$

Next, one can reduce from MAX-3SAT(29) to MAX-E3SAT(5), which is a special case of MAX-3SAT where an input 3-CNF formula contains L variables and $M = 5L/3$ clauses for some positive integer divisible by 3, each clause contains exactly 3 literals, and each variable appears in exactly 5 clauses but never appears twice in the same clause:

Theorem B.1.3 *[402], **Proposition 2.1.2** There is a polynomial-time gap-preserving reduction that transforms an instance ψ_{Max29} of MAX-3SAT(29) to an instance ψ_{Max5} of MAX-E3SAT(5) such that the following is satisfied:*

- $OPT(\psi_{Max29}) = 1 \quad \Rightarrow \quad OPT(\psi_{Max5}) = 1.$
- $OPT(\psi_{Max29}) \leq 1 - \delta'' \quad \Rightarrow \quad OPT(\psi_{Max5}) \leq 1 - \frac{\delta''}{200}.$

For our reduction this means $\delta'' = \frac{\frac{1}{8}-\delta}{43} \Rightarrow 1 - \frac{\delta''}{200} = 1 - \frac{\frac{1}{8}-\delta}{8600}$:

- $\mathrm{OPT}(\psi_{Max29}) = 1 \quad \Rightarrow \quad \mathrm{OPT}(\psi_{Max5}) = 1.$
- $\mathrm{OPT}(\psi_{Max29}) \leq 1 - \frac{\frac{1}{8}-\delta}{43} \quad \Rightarrow \quad \mathrm{OPT}(\psi_{Max5}) \leq 1 - \frac{\frac{1}{8}-\delta}{8600}.$

Now we arrived at a 3-SAT instance with a gap promise whose variables are uniformly distributed. We now implement a last step to create a balanced instance ψ_{Max5} whose literal polarities are also uniform:

For each variable y_l, pick a random bit $r_{y_l} \in \{0, 1\}$.

- If $r_{y_l} = 1$, flip the sign of every occurrence of y_l:

  $$y_l \longleftrightarrow \neg y_l,$$

  and change the assignment:

  $$s'_l = s_l \oplus 1$$

- If $r_{y_l} = 0$, do nothing.

We can check that the value of the literal doesn't change by the new assignment:

- If $l = y_l$: call the value under the new assignment s'_l, and the literal after the flip becomes $l' = \neg y_l$, when $r_{y_l} = 1$, and $l' = y_l$ when $r_{y_l} = 0$.

  $$l'(s'_l) = \begin{cases} \neg(s_l \oplus 1) = s_l & (r_{y_l} = 1) \\ s_l & (r_{y_l} = 0) \end{cases} \tag{B.1-1}$$

  which in both cases is the same: $l'(s'_l) = s_l = l(s_l)$.

- If $l = \neg y_l$: the literal after the flip becomes $l' = y_l$, when $r_{y_l} = 1$, and $l' = \neg y_l$ when $r_{y_l} = 0$.

  $$l'(s'_l) = \begin{cases} s_l \oplus 1 = \neg s_l & (r_{y_l} = 1) \\ \neg s_l & (r_{y_l} = 0) \end{cases} \tag{B.1-2}$$

  which in both cases is the same: $l'(s'_l) = \neg s_l = l(s_l)$.

To check why this routine makes literal polarities uniform, let t_l be the number of occurrences of variable y_l in the formula, and let p_l be the number of positive occurrences of y_l (so $n_l = t_l - p_l$ are the negative occurrences $\neg y_l$) before the random flips and let p'_l, n'_l be the occurrences afterwards. We have

- With probability $\frac{1}{2}$, $p'_l = p_l$
- With probability $\frac{1}{2}$, $p'_l = n_l$

Then

$$\mathbb{E}[p_l'] = \frac{1}{2}p_l + \frac{1}{2}n_l = \frac{1}{2}p_l + \frac{1}{2}(t_l - p_l) = \frac{t_l}{2}. \tag{B.1-3}$$

Obviously, the soundness gap is preserved in this routine, and our final gap problem is

$$\begin{aligned} &\phi \in \text{SAT} \Rightarrow \text{OPT}(\psi) = 1, \\ &\phi \notin \text{SAT} \Rightarrow \text{OPT}(\psi) \leq 1 - \frac{\frac{1}{8} - \delta}{8600}, \end{aligned} \tag{B.1-4}$$

where ψ denotes a balanced MAX-E3SAT(5) instance.

B.1.2 Calculation of min-entropies

In many quantum-cryptographic protocols, security is established by lower-bounding the min-entropy $H_\infty(X|Q)$ associated with an adversary guessing a variable X, conditioned on some quantum side-information Q. In this context, the joint state shared between an honest party and the adversary can be described by a classical-quantum state living in $\mathcal{H}_X \otimes \mathcal{H}_Q$ of the form

$$\rho_{XQ} = \sum_x P(X = x)|x\rangle\langle x|_X \otimes \rho_Q^x, \tag{B.1-5}$$

where $\{|x\rangle\}$ is the orthonormal set of possible classical strings in $\mathcal{H}_X$ weighted by $P(X = x)$ and ρ_Q^x is the quantum state in $\mathcal{H}_Q$ held by the adversary. The conditional min-entropy of an adversary holding the Q register then is defined as

$$H_\infty(X|Q)_{\rho_{XQ}} := -\log p_{\text{guess}}(X|Q), \tag{B.1-6}$$

where

$$p_{\text{guess}}(X|Q) := \max_{\{D_Q^x\}} \sum_x P(X = x)\text{Tr}(D_Q^x \rho_Q^x). \tag{B.1-7}$$

Here, the maximization is over all POVMs $\{D_Q^x\}$ on $\mathcal{H}_Q$. This captures the maximum probability of correctly guessing X given Q, i.e. how correlated the adversary's quantum state is to the classical string. Note that the min-entropy takes a maximum value of $H_\infty(X|Q)_{\rho_{XQ}} = \log|\mathcal{X}|$ when the classical register is completely uncorrelated from the adversary's quantum register, i.e. when $\rho_{XQ} = \frac{\mathbb{1}}{|\mathcal{X}|} \otimes \rho_Q$ (and where $|\mathcal{X}|$ denotes the alphabet size of X).

The smooth version of the conditional min-entropy is defined as

$$H_\infty^\varepsilon(X|Q)_{\rho_{XQ}} := \sup_{\substack{\bar{\rho}_{XQ} \geq 0: \ \frac{1}{2}||\bar{\rho}_{XQ} - \rho_{XQ}||_1 \leq \text{Tr}(\rho_{XQ})\cdot\varepsilon \\ \text{Tr}(\bar{\rho}_{XQ}) \leq \text{Tr}(\rho_{XQ})}} H_\infty^\varepsilon(X|Q)_{\bar{\rho}_{XQ}}, \tag{B.1-8}$$

which can be understood as an optimization over the min-entropies corresponding to states $\{\bar{\rho}_{XQ}\}$ within an ε-distance of ρ_{XQ}. A very useful identity for the smooth min-entropy of a classical-classical-quantum state ρ_{XYQ} is the following chain rule:

$$H_\infty^\varepsilon(X|YQ)_{\rho_{XYQ}} \geq H_\infty^\varepsilon(XY|Q)_{\rho_{XYQ}} - \log|\mathcal{Y}|. \tag{B.1-9}$$

where $|\mathcal{Y}|$ denotes the alphabet size of Y.

We now summarize the most important techniques for the noisy-storage part of our security proof. The first theorem states that the conditional smooth min-entropy of the measurement outcome $X \in \{0,1\}^N$ of an N-qubit state ρ, measured in a randomly chosen basis $\Theta \in \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowright\}^N$, can be lower-bounded:

Theorem B.1.4 ***Uncertainty relation 1*** *[106] Consider the N-qubit state $\rho \in (\mathbf{C}^2)^{\otimes N}$. Let $\Theta \in_R \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowright\}^N$ be independent of ρ and let $X \in \{0,1\}^N$ be the random variable for the measurement outcome of ρ measured in basis Θ. Then, the conditional smooth min-entropy is lower-bounded by*

$$H_\infty^\varepsilon(X|\Theta) \geq \left(\frac{1}{2} - 2\delta\right)N \tag{B.1-10}$$

for any $0 < \delta < \frac{1}{2}$ *satisfying*

$$\varepsilon = \exp\Big(-\frac{\delta^2 N}{32(2-\log(\delta))^2}\Big), \tag{B.1-11}$$

where ε should ideally be very low for a secure protocol. To make the implementation feasible, we need a tighter uncertainty relation:

Theorem B.1.5 ***Uncertainty relation 2*** *[436] Consider the N-qubit state $\rho \in (\mathbf{C}^2)^{\otimes N}$. Let $\Theta \in_R \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowright\}^N$ be independent of ρ and let $X \in \{0,1\}^N$ be the random variable for the measurement outcome of ρ measured in basis Θ. Then, the conditional smooth min-entropy is lower-bounded by*

$$H_\infty^\varepsilon(X|\Theta) \geq N \cdot c_{BB84} \tag{B.1-12}$$

where

$$c_{BB84} := \max_{s\in(0,1]} \frac{1}{s}[1+s-\log(1+2^s)] - \frac{1}{sN}\log\frac{2}{\varepsilon^2} \tag{B.1-13}$$

for any $\varepsilon > 0$.

The following lemma then shows that if the joint entropy of two random variables X_0 and X_1 is lower-bounded by some α, then the min-entropy of X_D where $D \in_R \{0,1\}$ is a *random* variable, is at least $\alpha/2 - 1$:

Lemma B.1.6 ***Min-entropy splitting*** *[408, 437], Let $\varepsilon \geq 0$, and X_0, X_1, Θ be random variables with $H_\infty^\varepsilon(X_0, X_1|\Theta) \geq \alpha$. Then, there exists a random variable $D \in_R \{0,1\}$ such that*

$$H_\infty^\varepsilon(X_D|D\,\Theta) \geq \alpha/2 - 1 \tag{B.1-14}$$

The above bound on the min-entropy quantifies the number of bits of uncertainty when guessing a substring $X_D \subset X$ given Θ. The uncertainty about X given classical side-information determines the effective rate an adversary must send through their storage device $\mathcal{F}$. We now state two useful lemmas:

Lemma B.1.7 *([106], Lemma II.1.) Consider an arbitrary cq-state ρ_{XQ} and a CPTPM $\mathcal{F}$: $\mathcal{H}_{Q_{in}} \to \mathcal{H}_{Q_{out}}$. Then*

$$H_\infty(X|\mathcal{F}(Q)) \geq -\log P_{succ}^{\mathcal{F}}(\lfloor H_\infty(X)\rfloor). \tag{B.1-15}$$

Lemma B.1.8 *([106], Lemma II.2.) Consider an arbitrary ccq-state ρ_{XTQ}, and let $\varepsilon, \varepsilon' \geq 0$ be arbitrary. Let $\mathcal{F} : \mathcal{H}_{Q_{in}} \to \mathcal{H}_{Q_{out}}$ be an arbitrary CPTP map. Then*

$$H_\infty^{\varepsilon+\varepsilon'}(X|T\mathcal{F}(Q)) \geq -\log P_{succ}^{\mathcal{F}}\Big(\lfloor H_\infty^\varepsilon(X|T) - \log\frac{1}{\varepsilon'}\rfloor\Big). \tag{B.1-16}$$

B.2 Security against P*

B.2.1 Length notations and finite-size effects

Considering a channel transmission of $\tilde{\eta}_{ch}$ and V's photon-number generation probabilities $\{p_0, p_1, p_{\mathrm{m}}\}$, P will detect vacuum, single-photon and multi-photon states with probabilities

$$\begin{aligned}
\tilde{p}_0 &= p_0 + p_1\cdot(1-\tilde{\eta}_{ch}) + p_{\mathrm{m}}\cdot(1-\tilde{\eta}_{ch})^2,\\
\tilde{p}_1 &= p_1\cdot\tilde{\eta}_{ch} + p_{\mathrm{m}}[1-\tilde{\eta}_{ch}^2 - (1-\tilde{\eta}_{ch})^2],\\
\tilde{p}_{\mathrm{m}} &= 1-\tilde{p}_1-\tilde{p}_0.
\end{aligned} \tag{B.2-1}$$

For the security proof, we must identify which probabilities are controlled by which party. The set $\{p_0, p_1, p_\mathrm{m}\}$ is trusted by V. However, once the states leave $V's$ lab, V has no control over the probabilities anymore, hence the untrusted probabilities $\{\tilde{p}_0, \tilde{p}_1, \tilde{p}_\mathrm{m}\}$ denoted with a tilde.

Let N be the total number of states sent per variable, and $\{Q_1^{(n)}, ..., Q_N^{(n)}\}$ be a set of N independent random variables. Furthermore, let $Q^{(n)} = \sum_{i=1}^{N} Q_i^{(n)}$, where $Q_i^{(n)} = 1$ with probability p_n for $n \in \{0, 1, \mathrm{m}\}$ and $Q_i^{(n)} = 0$ otherwise. The expectation value is $\mu_n = p_n \cdot N$. We may use a Chernoff-type argument to upper-bound the probability that emitted and detected photon numbers exceed their expectation value μ_n by some $\sigma_{n+} \in [0, 1]$:

$$\mathbf{P}[Q^{(n)} \geq (1 + \sigma_n) \cdot \mu_n] \leq \exp\left(-\frac{\sigma_{n+}^2}{3} \cdot \mu_n\right). \tag{B.2-2}$$

Similarly, we can bound the probability that $Q^{(n)} = \sum_{i=1}^{N} Q_i^{(n)}$ is smaller than the expectation value μ_n by some $\sigma_{n-} \in [0, 1]$:

$$\mathbf{P}[Q^{(n)} \leq (1 - \sigma_{n-}) \cdot \mu_n] \leq \exp\left(-\frac{\sigma_{n-}^2}{2} \cdot \mu_n\right). \tag{B.2-3}$$

Setting the probability that $Q^{(n)}$ lies outside the interval $[(1 - \sigma_{n-}) \cdot p_n \cdot N, (1 + \sigma_{n+}) \cdot p_n \cdot N]$ to be smaller than some $\varepsilon_{p_n} > 0$ yields the following relations:

$$\sigma_{n+} \leq \sqrt{\frac{3 \ln\left(\frac{1}{\varepsilon_{p_n}}\right)}{p_n \cdot N}}. \tag{B.2-4}$$

$$\sigma_{n-} \leq \sqrt{\frac{2 \ln\left(\frac{1}{\varepsilon_{p_n}}\right)}{p_n \cdot N}}. \tag{B.2-5}$$

In the same way, we can bound the probability $\tilde{p}_n$ that P receives n-photon states after considering additional channel losses $\tilde{\eta}_{ch}$. If $\tilde{Q}^{(n)}$ is the number of n-photon states that P detects, the probability that $\tilde{Q}^{(n)}$ lies outside the interval $[(1-\tilde{\sigma}_{n-}) \cdot \tilde{p}_n \cdot N, (1+\tilde{\sigma}_{n+}) \cdot \tilde{p}_n \cdot N]$ is upper-bounded by $\varepsilon_{\tilde{p}_n}$.

Our security proof holds under the assumption that V's source produces phase-randomized states. This is well-justified by the fact that our quantum dot is excited with an LA phonon-assisted pumping scheme, which generates mixed states in the photon-number basis [45]. In this context, P^* only attempts to store single-photon states in their quantum memory, since multi-photon states allow them to perfectly decommit by measuring one photon in each basis. Based on these considerations, we define and calculate the quantities N_detect and $N_\mathrm{detect}^{(1)}$.

When N states are generated by V, N_detect is the number of states we expect an honest P to detect *and* measure in the basis they wish to commit to. This is the set of states that can later be challenged. Given photon-number generation probabilities $\{p_0, p_1, p_\mathrm{m}\}$, channel transmission $\tilde{\eta}_{ch}$, the probability $\tilde{p}_\mathrm{dc}$ that any of the four detector registers a dark count, and assuming a balanced measurement BS, we write

$$N_\mathrm{detect} = \left(\frac{\tilde{p}_1}{2} + \frac{\tilde{p}_\mathrm{m}}{2} + \frac{\tilde{p}_\mathrm{dc}}{2}\right) \cdot N. \tag{B.2-6}$$

We upper-bound the probability that N_detect is not in the interval

$$\begin{aligned}\Bigg[&\left(\frac{(1 - \tilde{\sigma}_{1-}) \cdot \tilde{p}_1}{2} + \frac{(1 - \tilde{\sigma}_{\mathrm{m}-}) \cdot \tilde{p}_\mathrm{m}}{2} + \frac{(1 - \tilde{\sigma}_{\mathrm{dc}-}) \cdot \tilde{p}_\mathrm{dc}}{2}\right) \cdot N, \\ &\left(\frac{(1 + \tilde{\sigma}_{1+}) \cdot \tilde{p}_1}{2} + \frac{(1 + \tilde{\sigma}_{\mathrm{m}+}) \cdot \tilde{p}_\mathrm{m}}{2} + \frac{(1 + \tilde{\sigma}_{\mathrm{dc}+}) \cdot \tilde{p}_\mathrm{dc}}{2}\right) \cdot N\Bigg]\end{aligned} \tag{B.2-7}$$

by $\varepsilon_\mathrm{det} = \varepsilon_{\tilde{p}_1} + \varepsilon_{\tilde{p}_\mathrm{m}} + \varepsilon_{\tilde{p}_\mathrm{dc}} = 3\varepsilon_{\tilde{p}_1}$, where we calculate the various $\tilde{\sigma}_{n+}$ and $\tilde{\sigma}_{n-}$ parameters using Eq. B.2-4 and B.2-5 by setting $\varepsilon_{\tilde{p}_1} = \varepsilon_{\tilde{p}_\mathrm{m}} = \varepsilon_{\tilde{p}_\mathrm{dc}}$. If the number of reported states does not lie within this interval, we allow V to abort the protocol. In the following, we use the lower bound on N_detect, i.e.

$$N_{\text{detect,min}} = \left(\frac{(1-\tilde{\sigma}_{1-}) \cdot \tilde{p}_1}{2} + \frac{(1-\tilde{\sigma}_{\text{m}-}) \cdot \tilde{p}_{\text{m}}}{2} + \frac{(1-\tilde{\sigma}_{\text{dc}-}) \cdot \tilde{p}_{\text{dc}}}{2} \right) \cdot N \tag{B.2-8}$$

$N^{(1)}_{\text{detect}}$ is the number of expected single-photon states within N_{detect} that P^* cannot report as lost. Since, in the worst case, a dishonest P^* will never report multi-photon states as lost, we can write

$$N^{(1)}_{\text{detect}} = N_{\text{detect}} - p_{\text{m}} \cdot N = \left(\frac{\tilde{p}_1}{2} + \frac{\tilde{p}_{\text{m}}}{2} + \frac{\tilde{p}_{\text{dc}}}{2} - p_{\text{m}} \right) \cdot N \tag{B.2-9}$$

Again, we can upper-bound the probability that $N^{(1)}_{\text{detect}}$ is not within the interval

$$\left[\left(\frac{(1-\tilde{\sigma}_{1-}) \cdot \tilde{p}_1}{2} + \frac{(1-\tilde{\sigma}_{\text{m}-}) \cdot \tilde{p}_{\text{m}}}{2} + \frac{(1-\tilde{\sigma}_{\text{dc}-}) \cdot \tilde{p}_{\text{dc}}}{2} - (1+\sigma_{m+}) \cdot p_{\text{m}} \right) \cdot N, \right.$$
$$\left. \left(\frac{(1+\tilde{\sigma}_{1-}) \cdot \tilde{p}_1}{2} + \frac{(1+\tilde{\sigma}_{\text{m}-}) \cdot \tilde{p}_{\text{m}}}{2} + \frac{(1+\tilde{\sigma}_{\text{dc}-}) \cdot \tilde{p}_{\text{dc}}}{2} - (1-\sigma_{\text{m}-}) \cdot p_{\text{m}} \right) \cdot N \right]$$

by $\varepsilon_{\text{det}} = \varepsilon_{\tilde{p}_1} + \varepsilon_{\tilde{p}_{\text{m}}} + \varepsilon_{\tilde{p}_{\text{dc}}} + \varepsilon_{p_{\text{m}}} = 4\varepsilon_{\tilde{p}_1}$. Assuming the worst case, we take $N^{(1)}_{\text{detect, min}}$ as the lower bound on $N^{(1)}_{\text{detect}}$ for calculating the security bounds, i.e.

$$N^{(1)}_{\text{detect, min}} = \left(\left(\frac{(1-\tilde{\sigma}_{1-}) \cdot \tilde{p}_1}{2} + \frac{(1-\tilde{\sigma}_{\text{m}-}) \cdot \tilde{p}_{\text{m}}}{2} + \frac{(1-\tilde{\sigma}_{\text{dc}-}) \cdot \tilde{p}_{\text{dc}}}{2} \right) - \left((1+\sigma_{m+}) \cdot p_{\text{m}} + (1+\tilde{\sigma}_{\text{dc}+}) \cdot \tilde{p}_{\text{dc}} \right) \right) \cdot N. \tag{B.2-10}$$

Finally, we focus on the treatment of quantum bit error rates. In the honest setting, the expected total error that needs to be corrected is a combination of the channel error $p_{\text{error,ch}}$ and dark count probability $\tilde{p}_{\text{dc}}$. Normalizing to the total expected detection events, we have:

$$\tilde{p}_{\text{error}_+} = \frac{d}{2} + (1-d) \cdot p_{\text{error, ch}} \tag{B.2-11}$$

where, in the worst-case,

$$d = \frac{\tilde{p}_{\text{dc}} \cdot (1+\tilde{\sigma}_{\text{dc},+})}{\tilde{p}_{\text{dc}} \cdot (1+\tilde{\sigma}_{\text{dc},+}) + \tilde{p}_1 \cdot (1-\tilde{\sigma}_{1,-}) + \tilde{p}_{\text{m}} \cdot (1-\tilde{\sigma}_{\text{m},-})}.$$

Again we set $\varepsilon_{\text{err}} = \varepsilon_{\tilde{p}_{\text{dc}}} + \varepsilon_{\tilde{p}_1} + \varepsilon_{\tilde{p}_{\text{m}}} = 3\varepsilon_{\tilde{p}_1}$, i.e. the probability that the error is larger than $\tilde{p}_{\text{error},+}$ is upper-bounded by ε_{err}.

B.2.2 Security bounds

In our analysis, we derive two bounds on an adversary's success probability when trying to break the protocol. The first bound $\varepsilon_{\text{s,sat}}$ dictates the minimum number of variables (or their corresponding quantum states) that an adversary must store to exceed a cheating probability of $\varepsilon_{\text{s,sat}}$ when k clauses are challenged. The second bound $\varepsilon_{\text{s,q}}$ shows that, by sending a sufficiently large number N of quantum states per variable, the probability that an adversary with storage size S guesses the string X_{v} is upper bounded. Both bounds relate to one another in the following way: whenever P^* tries to increase their cheating probability above $\varepsilon_{\text{s,sat}}$ by storing the quantum states associated with more problem variables, our protocol ensures that the verification string X_{v} achieves at least 7-bit security, except with probability $\varepsilon_{\text{min}} + 2\varepsilon_{\text{det}}$ which comes from statistical errors (ε_{det}) and smoothing the min-entropy (ε_{min}). Conversely, if P^* stores fewer quantum states for which the storage capacity suffices to reduce the uncertainty of X_{v} below 7 bits, the probability of detecting an error in k challenges becomes larger than $(1-\varepsilon_{\text{s,sat}})$. With this, the total soundness, i.e. the probability that V accepts a dishonest proof, is bounded by

$$\begin{aligned} \varepsilon_{\text{s}} &= \varepsilon_{\text{s,sat}} + \varepsilon_{\text{s,q}} \\ &= \varepsilon_{\text{s,sat}} + 2\varepsilon_{\text{det}} + \varepsilon_{\text{min}} + \varepsilon_{\text{key}}, \end{aligned} \tag{B.2-12}$$

where the factor of 2 for ε_{det} comes from $N_{\text{detect, min}}$ and $N^{(1)}_{\text{detect, min}}$ (defined in Eqs. B.2-8 and B.2-10, respectively), and $\varepsilon_{\text{key}} = 2^{-7}$. From B.1.1 we know that if $\phi \notin$ SAT, then at most $1 - \frac{\frac{1}{8}-\delta}{8600}$ of all clauses of ψ can be satisfied. Recalling that each variable appears in five distinct clauses, the probability that V challenges only satisfied or stored clauses over k rounds is given by

$$\left(1 - \frac{\frac{1}{8}-\delta}{8600} + 5 \cdot \frac{y_{\text{store}}}{M}\right)^k = \left(1 - \frac{\frac{1}{8}-\delta}{8600} + 3 \cdot \frac{y_{\text{store}}}{L}\right)^k < \varepsilon_{\text{s,sat}}, \tag{B.2-13}$$

where y_{store} is the number of variables stored by P^* and M is the number of clauses in ψ. If P^* can store a variable in a clause reliably, then they can alter the commitment of this variable, making an initially unsatisfied clause evaluate to *true*. Note that Eq. B.2-13 holds in the worst case, i.e. when the set of satisfied clauses does not overlap with the set of clauses for which a variable is stored. From Eq. B.2-13, we calculate the threshold $y_{\min}$ for which P^* can successfully cheat with probability $\varepsilon_{\text{s,sat}}$:

$$y_{\text{store}} > y_{\min} = \frac{L}{3} \cdot \left(\varepsilon_{\text{s,sat}}^{1/k} - (1 - \frac{\frac{1}{8}-\delta}{8600})\right). \tag{B.2-14}$$

By increasing the number of stored variables - or equivalently, the number of quantum states - P^* decreases their storage rate ν, given by the ratio of the actual storage size S to the number of states that they store. For a given $\varepsilon_{\text{s,sat}}$, we define the maximum storage rate $\nu_{\varepsilon_{\text{s,sat}}}$ as the highest rate achievable by P^* while successfully cheating with probability $\varepsilon_{\text{s,sat}}$:

$$\nu_{\varepsilon_{\text{s,sat}}} = \frac{S}{N^{(1)}_{\text{detect,min}} \cdot y_{\min}}. \tag{B.2-15}$$

Here, we assume that for one variable, P^* only has to store single-photon states that were not reported as lost, corresponding to the number $N^{(1)}_{\text{detect,min}}$ expressed in Eq. B.2-10.

When $y_{\text{store}} > y_{\min}$ and *no incorrectly measured clause but a stored variable* is challenged, we choose N (per variable) such that the resulting verification string retains a security of at least 7 bits.

B.2.3 String definitions

First, we define the *total challenged string* X_{K} as the concatenation of all encoded bits corresponding to the challenged variables. Let $\mathcal{L} \subset \{1, \ldots, L\}$ be the set of variables included in the k clauses, with $|\mathcal{L}| \approx 3k$. For each variable $y_{l'}$ with $l' \in \mathcal{L}$, let $X_{y_{l'}}$ denote the string that V encoded (regardless of their chosen basis) which agrees with the committed set of timings $\tilde{t}_{s_{l'}}$ sent by P in Step 3 of the commit phase, satisfying $|X_{y_{l'}}| \geq N_{\text{detect,min}}$. Then X_{K} may be expressed as

$$X_{\text{K}} = X_{y_{1'}} || \ldots || X_{y_{|\mathcal{L}|'}}, \tag{B.2-16}$$

where $|X_{\text{K}}| \geq |\mathcal{L}| \cdot N_{\text{detect,min}}$.

A *valid verification string* X_{v} is then defined as the concatenation of substrings which P and V both measured/encoded in the basis corresponding to a *satisfying assignment* of the k clauses. For a given challenged variable $y_{l'}$, these substrings satisfy either

$$X_{s_{l'}} = X_{\text{true}} = \mathcal{C}_{\left[(\Theta_j = \tilde{\Theta}_{j=} \leftrightarrow\updownarrow) \wedge (j \in \tilde{t}_{\text{true}})\right]}(X_j) \quad \text{for } s_{l'} = \text{true} \tag{B.2-17}$$

or

$$X_{s_{l'}} = X_{\text{false}} = \mathcal{C}_{\left[(\Theta_j = \tilde{\Theta}_{j=} \leftrightarrow\updownarrow) \wedge (j \in \tilde{t}_{\text{false}})\right]}(X_j) \quad \text{for } s_{l'} = \text{false}, \tag{B.2-18}$$

where $\mathcal{C}_{[A]}(B)$ denotes the concatenation of all B strings such that condition A holds, $\tilde{t}_{\text{true}} = \{t_j \mid \tilde{\Theta}_j = \leftrightarrow\updownarrow\}$, $\tilde{t}_{\text{false}} = \{t_j \mid \tilde{\Theta}_j = \circlearrowleft\circlearrowright\}$ and $|X_{\text{true}}| \approx |X_{\text{false}}| \approx |X_{y_{l'}}|/2$

For $l' \in \mathcal{L}$, we now write $r_{l'} \in \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowright\}$ to denote that only one basis substring of a variable $y_{l'}$ is addressed. Note that we do not use $s_{l'} \in \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowright\}$ to distinguish it from the original honest solution assignment s, because a valid verification string could differ from that as long as the

chosen assignments $\{r_{l'}\}$ satisfy the challenge. We then define the set of *valid verification strings* $\mathbf{X}_\mathrm{v}$ as:

$$\mathbf{X}_\mathrm{v} = \left\{ X_\mathrm{v} = X_{r_{1'}} || X_{r_{2'}} || \ldots || X_{r_{|\mathcal{L}|'}} \;\middle|\; \begin{array}{l} r_{l'} \in \{\leftrightarrow\updownarrow,\ \circlearrowleft\circlearrowleft\} \ , \ l' \in \mathcal{L}, \\ \text{and the assignment satisfies all } k \text{ clauses.} \end{array} \right\}, \quad \text{(B.2-19)}$$

and $|X_\mathrm{v}| \geq |\mathcal{L}| \cdot N_{\mathrm{detect,min}}/2$. To pass the verification, P^* needs to have knowledge about one of the strings in the set $\mathbf{X}_\mathrm{v}$.

Now let $\mathcal{V}$ denote the set of stored, challenged variables and $\mathbf{V}_\mathcal{L}$ be the set of all subsets of $\mathcal{L}$ of size $|\mathcal{V}|$, i.e.

$$\mathbf{V}_\mathcal{L} := \{\ V_\mathcal{L} \subset \mathcal{L} \mid |V_\mathcal{L}| = |\mathcal{V}|\ \}\,. \quad \text{(B.2-20)}$$

Recall that P^* must store at least y_min variables to cheat with probability $\varepsilon_\mathrm{s,sat}$. Storing a variable $y_{l'}$ in this context means storing the quantum states that correspond to the committed set of timings $\tilde{t}_{s_{l'}}$. We define the set of strings corresponding to stored, challenged subsets (i.e. subsets belonging to $\mathbf{V}_\mathcal{L}$) which agree with the committed set of timings and are encoded in *any* basis as:

$$\mathbf{X}^\mathcal{V} = \left\{ X^\mathcal{V} = \ \mathcal{C}_{[v' \in V_\mathcal{L}]}\left(X_{y_{v'}}\right) \quad \mid V_\mathcal{L} \in \mathbf{V}_\mathcal{L} \ \right\}, \quad \text{(B.2-21)}$$

where, once again, $\mathcal{C}_{[A]}\left(B\right)$ denotes the concatenation of all B strings such that condition A holds. Assuming the worst case, namely that all except $|\mathcal{V}|$ challenged variables were measured in a basis satisfying the challenges, the adversary's uncertainty about $X_\mathrm{v} \in \mathbf{X}_\mathrm{v}$ reduces to their uncertainty about $X^\mathcal{V} \in \mathbf{X}^\mathcal{V}$, which, in turn, arises from the limitations of their imperfect storage device.

We now wish to form two subsets of $\mathbf{X}^\mathcal{V}$, respectively labeled $\mathbf{X}^\mathcal{V}_\mathrm{correct}$ and $\mathbf{X}^\mathcal{V}_\mathrm{incorrect}$, such that, in combination with the measurement outcome strings from the challenged and un-stored set $\mathcal{L}\backslash V_\mathcal{L}$, the resulting concatenated string is in the set X_v of valid verification strings (resp. is not in that set). This gives the following definitions:

$$\mathbf{X}^\mathcal{V}_\mathrm{correct} = \left\{ \begin{array}{l} X^\mathcal{V}_\mathrm{correct} = \mathcal{C}_{[v' \in V_\mathcal{L}]}(X_{r_{v'}}) \\ \quad \mathcal{C}_{[i' \in \mathcal{L}\backslash V_\mathcal{L}]}(X_{\mathrm{v},r_{i'}}) \,||\, \mathcal{C}_{[v' \in V_\mathcal{L}]}(X_{r_{v'}}) \in \mathbf{X}_\mathrm{v} \ , \forall\ (X_\mathrm{v}, V_\mathcal{L}) \in \mathbf{X}_\mathrm{v} \times \mathbf{V}_\mathcal{L} \\ \text{and } (r_{i'}, r_{v'}) \in \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowleft\}^2. \end{array} \right\} \quad \text{(B.2-22)}$$

and

$$\mathbf{X}^\mathcal{V}_\mathrm{incorrect} = \left\{ \begin{array}{l} X^\mathcal{V}_\mathrm{incorrect} = \mathcal{C}_{[v' \in V_\mathcal{L}]}(X_{r_{v'}}) \\ \quad \mathcal{C}_{[i' \in \mathcal{L}\backslash V_\mathcal{L}]}(X_{\mathrm{v},r_{i'}}) \,||\, \mathcal{C}_{[v' \in V_\mathcal{L}]}(X_{r_{v'}}) \notin \mathbf{X}_\mathrm{v} \ , \forall\ (X_\mathrm{v}, V_\mathcal{L}) \in \mathbf{X}_\mathrm{v} \times \mathbf{V}_\mathcal{L} \\ \text{and } (r_{i'}, r_{v'}) \in \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowleft\}^2. \end{array} \right\}, \quad \text{(B.2-23)}$$

where $X_{\mathrm{v},\,i'}$ denotes the i^{th} substring of a valid verification key X_v, and $X_{\mathrm{v},\,r_{i'}}$ denotes the basis choice for the i^{th} substring of a valid verification string X_v.

B.2.4 Uncertainty about X_v

We work with a *purified* version of the protocol, in which V prepares N EPR pairs for each variable:

1. For each y_l, V prepares N EPR pairs and sends one half to P.
2. For each pair, P measures their half in a random basis $\tilde{\Theta} \in_R \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowleft\}^{2N_\mathrm{detect}}$, obtaining two outcome strings $\tilde{X}_{\leftrightarrow\updownarrow}$ and $\tilde{X}_{\circlearrowleft\circlearrowleft}$. Each of theses strings has length $\approx N_\mathrm{detect}$.
3. For each pair, V measures their half in a random basis $\Theta \in_R \{\leftrightarrow\updownarrow, \circlearrowleft\circlearrowleft\}^N$. This yields outcome strings $X_{\leftrightarrow\updownarrow}$ or $X_{\circlearrowleft\circlearrowleft}$ for each variable y_l, where $|X_{\leftrightarrow\updownarrow}| \approx |X_{\circlearrowleft\circlearrowleft}| \approx N/2$.

In this context, P^*' view of both protocols is identical: one can transform the purified version into the original one by letting V perform their measurement before Step 2. Thus, security against

P^* in the purified protocol implies security in the original protocol. P^*'s uncertainty about X_v effectively reduces to their uncertainty about the basis strings associated with the set $\mathcal{V}$ of stored and challenged variables. Focusing on the worst case, we set $|\mathcal{V}| = 1$, implying that $|\mathbf{V}_\mathcal{L}| = |\mathcal{L}|$. Then, for each $(X_\mathrm{v}, V_\mathcal{L}) \in \mathbf{X}_\mathrm{v} \times \mathbf{V}_\mathcal{L}$, there exists only one string $X^\mathcal{V}_\mathrm{correct}$ and one string $X^\mathcal{V}_\mathrm{incorrect}$.

Note that, once V projects their classical registers, the state received by P^* has a tensor-product form. Given that the min-entropy of a tensor-product state is additive [148], the min-entropy of X_{N_enc} given classical side-information K_{N_enc} and Θ_{N_enc} is additive. This allows to compute the min-entropy for the relevant subsets of rounds independently of the rounds outside the set. In particular, we can derive the remaining min-entropy of X_K (defined in Eq. B.2-16), the min-entropy of $X^\mathcal{V}$ (defined in Eq. B.2-21) and that of $X^{(1),\mathcal{V}}$ which covers single-photon states that P^* did not declare missing.

We are now ready to derive the security of the verification bit string X_v when $y_\mathrm{store} > y_\mathrm{min}$ and all challenged clauses are satisfied but one stored variable is challenged, i.e. $|\mathcal{V}| = 1$ (we keep the index $\mathcal{V}$ to make clear that we refer to the stored variable). Assuming that neither multi-photon events nor rounds corresponding to the variables $\mathcal{L}\backslash\mathcal{V}$ contribute to P^*'s uncertainty, we find that

$$H_\infty^{\varepsilon_\mathrm{min}/2}(X_K|K_K\Theta_K) = H_\infty^{\varepsilon_\mathrm{min}/2}(X^{(1),\mathcal{V}}|K^{(1),\mathcal{V}}\Theta^{(1),\mathcal{V}}) \tag{B.2-24}$$

where $K^{(1),\mathcal{V}}$ and $\Theta^{(1),\mathcal{V}}$ denote the classical side-information strings corresponding only to single-photon events. By the uncertainty relation from Eq. (B.1-12), we have

$$H_\infty^{\varepsilon_\mathrm{min}/2}(X^{(1),\mathcal{V}}|K^{(1),\mathcal{V}}\Theta^{(1),\mathcal{V}}) \geq |\mathcal{V}| \cdot N^{(1)}_{\mathrm{detect,min}} \cdot c_{BB84}, \tag{B.2-25}$$

where $N^{(1)}_{\mathrm{detect,min}}$ is the minimum number of single-photon events that are not reported missing, defined in Eq. B.2-10. From Eq. B.2-25, we can calculate the rate at which a dishonest P^* must send information through their storage. To derive a lower bound on the min-entropy associated with the substrings $\{X^{(1),\mathcal{V}}_\mathrm{D}\}_{D\in_R\{\mathrm{correct,incorrect}\}}$, we apply the min-entropy splitting theorem from Eq. B.1-14 and get

$$H_\infty^{\varepsilon_\mathrm{min}/2}(X_D^{(1),\mathcal{V}}|D\ K^{(1),\mathcal{V}}\ \Theta^{(1),\mathcal{V}}) \geq \frac{|\mathcal{V}| \cdot N^{(1)}_{\mathrm{detect,min}} \cdot c_{BB84}}{2} - 1. \tag{B.2-26}$$

Since $X_D^{(1),\mathcal{V}}$ is a subset of the secret bit-string X_v, and by assumption the only subset for which P^* has some uncertainty, we can relate the lower bound on $X_D^{(1),\mathcal{V}}$ conditioned on classical information to the lower bound on X_v:

$$H_\infty^{\varepsilon_\mathrm{min}/2}(X_\mathrm{v}|K_\mathrm{v}\Theta_\mathrm{v}) = H_\infty^{\varepsilon_\mathrm{min}/2}(X_D^{(1),\mathcal{V}}|D\ K_{(1),\mathcal{V}}\ \Theta_{(1),\mathcal{V}}). \tag{B.2-27}$$

Again, we used the fact that conditioning on independent information does not decrease the uncertainty, and the additivity of min-entropy for tensor-product states. Given Eq. B.2-26, we can bound the uncertainty about X_v given classical and quantum side-information via B.1.8:

$$H_\infty^{\varepsilon_\mathrm{min}}(X_\mathrm{v}|K_\mathrm{v}\ \Theta_\mathrm{v}\ \mathcal{F}(Q_{in})) \geq -\log P^\mathcal{F}_\mathrm{succ}\left(\frac{|\mathcal{V}| \cdot N^{(1)}_{\mathrm{detect,min}} \cdot c_{BB84}}{2} - 1 - \log\left(\frac{2}{\varepsilon_\mathrm{min}}\right)\right). \tag{B.2-28}$$

The rate at which a dishonest P^* needs to send information through their strorage channel is then given by:

$$R = \frac{1}{y_\mathrm{store} \cdot N^{(1)}_{\mathrm{detect,min}}}\left(\frac{|\mathcal{V}| \cdot N^{(1)}_{\mathrm{detect,min}} \cdot c_{BB84}}{2} - 1 - \log\left(\frac{2}{\varepsilon_\mathrm{min}}\right)\right). \tag{B.2-29}$$

We now have to subtract the additional syndrome information that V sends in Step 3 of the challenge phase. Here, we can disregard the syndrome information for the outcome strings corresponding to variables outside of the set $\mathcal{V}$, as we assume P^* has zero uncertainty about them,

and they do not provide any information about $X^{\mathcal{V}}$, i.e. we only need to consider $Syn(X_D^{\mathcal{V}})$. We take the error-correction efficiency as 1.2. The resulting min-entropy bound on X_{v} is

$$\begin{aligned} H_\infty^{\varepsilon_{\min}}(X_{\mathrm{v}}|K_{\mathrm{v}}\,\Theta_{\mathrm{v}}\,\mathcal{F}(Q_{in})Syn(X_D^{\mathcal{V}})) &\geq H_\infty^{\varepsilon_{\min}}(X_{\mathrm{v}}|K_{\mathrm{v}}\,\Theta_{\mathrm{v}}\,\mathcal{F}(Q_{in})) \\ &\quad - 1.2\cdot\frac{|\mathcal{V}|\cdot N_{\mathrm{detect,min}}}{2}\cdot h(\tilde{p}_{\mathrm{error}_+}) \\ &\geq -\log P_{\mathrm{succ}}^{\mathcal{F}}\left(\frac{|\mathcal{V}|\cdot N_{\mathrm{detect,min}}^{(1)}\cdot c_{BB84}}{2} - 1 - \log\left(\frac{2}{\varepsilon_{\min}}\right)\right) \\ &\quad - 1.2\cdot\frac{|\mathcal{V}|\cdot N_{\mathrm{detect,min}}}{2}\cdot h(\tilde{p}_{\mathrm{error}_+}), \end{aligned} \tag{B.2-30}$$

where $h(\tilde{p}_{\mathrm{error}_+})$ is the binary entropy function $h(p) := -p\cdot\log(p) - (1-p)\cdot\log(1-p)$ restricted to $0 < p \leq 0.5$. Requiring B.2-30 to be ≥ 7, we get a verification string that has at least 7-bit security.

We also briefly discuss the security of the verification bit string X_{v} when $y_{\mathrm{store}} > y_{\min}$ and *no stored variable is challenged*, and *all except* $|\mathcal{W}|$ *variables are measured in a basis that satisfies the challenged clauses*. Let $X^{\mathcal{W}}$ denote the bit-string that corresponds to the wrongly measured variables and that both V and P^* declare as measured in the same basis. In this case, the uncertainty of the verification string is independent of the adversary's storage. We again look at the worst case where $|\mathcal{W}| = 1$. With the same argumentation as above, namely that the uncertainty of the verification bit string effectively reduces to that of the substrings corresponding to the variables in $\mathcal{W}$, from the uncertainty relation and the min-entropy splitting theorem, we get that

$$H_\infty^{\varepsilon_{\min}}(X_{\mathrm{v}}|D\,\Theta_{\mathrm{v}}) \geq \frac{|\mathcal{W}|\cdot N_{\mathrm{detect,min}}^{(1)}\cdot c_{BB84}}{2} - 1. \tag{B.2-31}$$

After V sends the syndrome information we get

$$H_\infty^{\varepsilon_{\min}}(X_{\mathrm{v}}|D\,\Theta_{\mathrm{v}}\,Syn(X_D^{\mathcal{W}})) \geq \frac{|\mathcal{W}|\cdot N_{\mathrm{detect,min}}^{(1)}\cdot c_{BB84}}{2} - 1 - 1.2\cdot\frac{|\mathcal{W}|\cdot N_{\mathrm{detect,min}}}{2}\cdot h(\tilde{p}_{\mathrm{error}_+}), \tag{B.2-32}$$

which we again require to be ≥ 7.

B.2.5 Depolarizing channel

We now analyze security for the depolarizing channel $\mathcal{F} = \mathcal{N}^{\otimes\nu_{\varepsilon_{\mathrm{s,sat}}}y_{\mathrm{store}}N_{\mathrm{detect,min}}^{(1)}} = \mathcal{N}^{\otimes S}$ with $\mathcal{N}(\rho) = r\cdot\rho + (1-r)\frac{\mathbb{1}}{d}$, and storage rate $\nu_{\varepsilon_{\mathrm{s,sat}}}$ (note that the storage rate is a function of the soundness threshold $\varepsilon_{\mathrm{s,sat}}$, see B.2-15). We model the number of qubits that determine the storage size as the fraction of states P can store to pass the verification with probability $\varepsilon_{\mathrm{s,sat}}$, i.e. $S = \nu_{\varepsilon_{\mathrm{s,sat}}}\cdot y_{\mathrm{store}}\cdot N_{\mathrm{detect,min}}^{(1)}$. From the strong converse property we know that the success probability for rates above the capacity decays as

$$-\log P_{succ}^{\mathcal{N}^{\otimes\nu_{\varepsilon_{\mathrm{s,sat}}}y_{\mathrm{store}}N_{\mathrm{detect,min}}^{(1)}}}(R\,y_{\mathrm{store}}N_{\mathrm{detect,min}}^{(1)}) \geq \nu_{\varepsilon_{\mathrm{s,sat}}}\cdot y_{\mathrm{store}}\cdot N_{\mathrm{detect,min}}^{(1)}\cdot\gamma^{\mathcal{N}}(R/\nu_{\varepsilon_{\mathrm{s,sat}}}). \tag{B.2-33}$$

From B.2-29 we know that the rate is

$$R = \frac{1}{y_{\mathrm{store}}\cdot N_{\mathrm{detect,min}}^{(1)}}\left(\frac{|\mathcal{V}|\cdot N_{\mathrm{detect,min}}^{(1)}\cdot c_{BB84}}{2} - 1 - \log\left(\frac{2}{\varepsilon_{\min}}\right)\right). \tag{B.2-34}$$

We can then insert the bound on the success probability B.2-33 for the depolarizing channel with the rate B.2-29 into equation B.2-30 and get

$$
\begin{aligned}
H_\infty^\varepsilon(X_\mathrm{v}|K_\mathrm{v}\ \Theta_\mathrm{v}\ \mathcal{F}(Q_{in})Syn(X_D^{\mathcal{V}})) \geq & -\log P_{\mathrm{succ}}^{\mathcal{N}^{\otimes \nu_{\varepsilon_{\mathrm{s,sat}}} y_{\mathrm{store}} N^{(1)}_{\mathrm{detect,min}}}} \\
& \cdot \left(\frac{|\mathcal{V}| \cdot N^{(1)}_{\mathrm{detect,min}} \cdot c_{BB84}}{2} - 1 - \log\left(\frac{2}{\varepsilon_{\min}}\right)\right) \\
& - 1.2 \cdot \frac{|\mathcal{V}| \cdot N_{\mathrm{detect,min}}}{2} \cdot h(\tilde{p}_{\mathrm{error}_+}) \\
\geq & \ \nu_{\varepsilon_{\mathrm{s,sat}}} \cdot y_{\mathrm{store}} \cdot N^{(1)}_{\mathrm{detect,min}} \cdot \gamma^{\mathcal{N}} \\
& \cdot \left(\frac{1}{\nu_{\varepsilon_{\mathrm{s,sat}}} \cdot y_{\mathrm{store}} \cdot N^{(1)}_{\mathrm{detect}}} \left(\frac{|\mathcal{V}| \cdot N^{(1)}_{\mathrm{detect,min}} \cdot c_{BB84}}{2} \right.\right. \\
& \left.\left. - 1 - \log\left(\frac{2}{\varepsilon_{\min}}\right)\right)\right) \\
& - 1.2 \cdot \frac{|\mathcal{V}| \cdot N_{\mathrm{detect,min}}}{2} \cdot h(\tilde{p}_{\mathrm{error}_+}).
\end{aligned}
\tag{B.2-35}
$$

To determine the minimum number N that has to be sent by V, we require that the right-hand side of equation B.2-35 is at least 7, which means that the verification string has at least 7-bit security. The soundness error, i.e. the probability that a dishonest P obtains a valid verification string is $\varepsilon_{\mathrm{s,q}} = \varepsilon_{\min} + 2\varepsilon_{\mathrm{det}} + \varepsilon_{\mathrm{key}}$, which consists of the probability $2\varepsilon_{\mathrm{det}}$ that P receives less photons and single photon states than expected, an error $\varepsilon_{\min}$ coming from the smooth min-entropy and the probability $\varepsilon_{\mathrm{key}} = 2^{-7}$ of randomly guessing X_v.

B.3 Security against V*

B.3.1 Commitment phase

Here, we argue that a dishonest V^* cannot learn anything from the satisfying assignment to ψ (or ϕ) other than what is directly revealed by P during the challenge step of the commit phase. Due to the structure of SAT problems, however, this direct leakage may also trigger indirect leakage of variables, which we discuss in Sec. 6.3.4. For the idealized protocol, i.e., in the absence of experimental imperfections, our security argument can be formulated similarly to [437]. Assuming that V^* does not know the satisfying assignment for the input instance ψ, P's measurement basis choice—and hence their committed bit—appears random to V^*, i.e., $P[s_l = true] = P[s_l = false]$. A memory-unbounded P can indeed create a *dummy* commitment that can later be opened as either bit without altering V^*'s view. Specifically, P can declare a random timing set $\tilde{t}_{s_l} \subset t_l$ of size N_{detect} for each variable y_l, store the corresponding quantum states, and await V^*'s challenges. Only after the challenges are revealed does P measure the states and output $\tilde{X}_{s_l}$. The only difference between the honest and hypothetical unbounded P is their measurement basis choice, which V^* cannot learn since they do not have access to P's laboratory. This provides the hiding property of our scheme. As a result, a hypothetically unbounded P can open the commitment as either bit, hence revealing no information about the bit before the reveal phase. This provides the hiding property of our scheme.

We now show that our scheme remains hiding in the presence of experimental imperfections. Before the challenge phase, the only information obtained by V^* is the timing string $\tilde{t}_{s_l}$ provided by P in the commitment phase, revealing the rounds for which the states were measured in the correct basis. If $\tilde{t}_{s_l}$ for $s_l = true$ is distinguishable from $\tilde{t}_{s_l}$ for $s_l = false$ (i.e. the set $\tilde{t}_{s_l}$ cannot be chosen uniformly at random for either bit value), then V^* could distinguish between the two values. Such distinguishability could arise, for instance, if detectors do not behave symmetrically or if P explicitly back-reports the type of click configuration that occurred on their detectors: by sending brighter polarized pulses, V^* can indeed remotely probe the measurement basis chosen by

P [154]. To prevent this leakage, any detection round yielding more than one click is uniformly mapped by P onto one of the four detector outcomes. This is an improvement over previous implementations of noisy-storage schemes [352], in which symmetrization of losses was used to counter some (but not all) loopholes of this type [154]. Finally, we note that if the number of multi-click events exceeds the error threshold that can be corrected through standard syndrome decoding, the protocol will abort in the verification phase when part of the verification string s has already been revealed. To avoid this, P should honestly abort the protocol prior to the challenges.

B.3.2 Verification phase

From Eq. B.1-4, it is NP-hard to distinguish $\mathrm{OPT}(\psi) = 1$ (YES case) from $\mathrm{OPT}(\psi) \leq 1 - \frac{\frac{1}{8}-\delta}{8600}$ (NO case), since this is equivalent to deciding whether $\phi \in \mathrm{SAT}$. We now establish a *hardness under bounded leakage* property: even if an algorithm is given the values of $B = O(1)$ variables of ψ (i.e., $O(1)$ bits of information about a satisfying assignment), this does not yield a polynomial-time procedure that can reliably distinguish the YES from the NO cases on all instances.

Assume there exist constants $B_0, \beta > 0$ and a polynomial-time routine V^* such that for *every* MAX-E3SAT(5) instance ψ, whenever $\mathrm{OPT}(\psi) = 1$, there is a set $S(\psi)$ of at most B_0 distinct variables and values satisfying

$$f_{V^*(\psi,S(\psi))}(\psi) \;\geq\; 1 - \frac{\frac{1}{8}-\delta}{8600} + \beta,$$

where $V^*(\psi, S(\psi))$ is the assignment produced by V^* when given the instance ψ and its leaked set $S(\psi)$, and $f_{V^*(\psi,S(\psi))}(\psi)$ is the fraction of clauses satisfied by $V^*(\psi, S(\psi))$. Intuitively, this assumption says that there is a polynomial-time algorithm V^* which, whenever ψ is a YES instance, can exploit a leakage of at most B_0 distinct variables to produce an assignment that beats the hardness gap by a fixed margin β. We now show that if such a V^* existed, then one could construct a polynomial-time algorithm that *distinguishes* the two cases YES and NO *without any leakage at all*, thus contradicting NP-hardness. We define a polynomial-time algorithm A^* that, on input ψ with L variables, simulates all possible leakages of up to B_0 distinct variables:

1. For each $j = 0, 1, \ldots, B_0$, for each choice of j variables (there are $\binom{L}{j}$ choices), and for each assignment to those positions (2^j choices), run the routine V^* to obtain a candidate assignment $\tilde{s} = V^*(\psi, S)$. Compute $\mathrm{OPT}(\psi)_{\tilde{s}}$ and record the best value
$$V_{\max} = \max_{|S| \leq B_0} f_{V^*(\psi,S)}(\psi).$$
2. Output
$$\begin{aligned} &\phi \in \mathrm{SAT} \quad \text{if} \quad V_{\max} > 1 - \frac{\frac{1}{8}-\delta}{8600} + \beta, \\ &\phi \notin \mathrm{SAT} \quad \text{otherwise.} \end{aligned} \tag{B.3-1}$$

The total number of trials is

$$\sum_{j=0}^{B_0} \binom{L}{j} 2^j \leq \sum_{j=0}^{B_0} (2L)^j = O\big((2L)^{B_0}\big) = \mathrm{poly}(L),$$

since $B_0 = O(1)$. Thus, A^* runs in polynomial-time. By assumption, for every YES instance ψ, there exists a particular leakage $S(\psi)$ of size at most B_0 such that

$$f_{V^*(\psi,S(\psi))}(\psi) > 1 - \frac{\frac{1}{8}-\delta}{8600} + \beta.$$

Among the sets S enumerated by A^*, this particular $S(\psi)$ will appear, and hence $V_{\max}$ will exceed the threshold in the YES case. In contrast, in any NO instance, we have

$$\mathrm{OPT}(\psi) \leq 1 - \frac{\frac{1}{8}-\delta}{8600},$$

so no assignment can achieve a value larger than this bound. Therefore, for all S we have

$$f_{V^*(\psi,S}(\psi) \leq 1 - \frac{\frac{1}{8} - \delta}{8600},$$

and hence $V_{\max}$ never exceeds the threshold in the NO case. Thus A^* correctly distinguishes whether $\phi \in$ SAT or $\phi \notin$ SAT in polynomial time, contradicting the NP-hardness of the underlying gap problem. This implies that, assuming P $\neq$ NP, there is no polynomial-time algorithm that can, for all instances, systematically use only $B_0 = O(1)$ leaked variable values to construct assignments that beat the hardness gap. In particular, as an immediate corollary, assuming P $\neq$ NP, there is no polynomial-time algorithm that, for every satisfiable instance ψ, can use only $B_0 = O(1)$ leaked variable values to reconstruct a full satisfying assignment.