Professorship for Quantum Computing
TUM School of Computation, Information and Technology
Technical University of Munich

Towards More Accurate Numerical Modelling of Qubit Dynamics Controlled by Quantum Light

Addressing the dimensionality problem while maintaining accuracy

Oriol Bertomeu

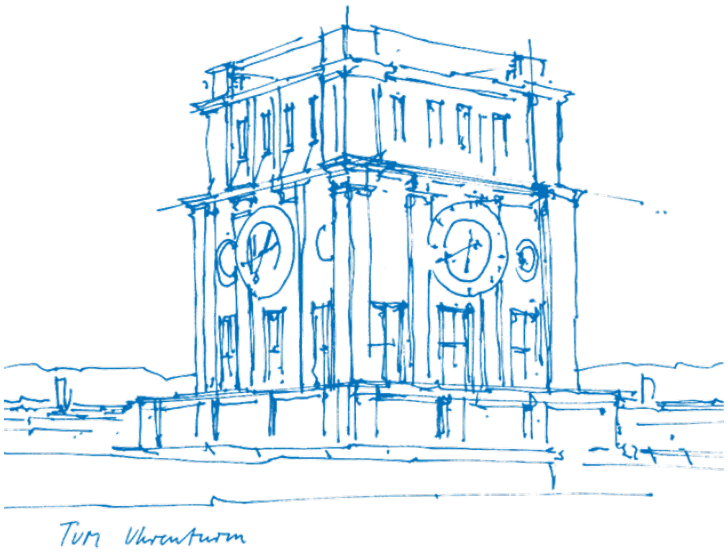

Towards More Accurate Numerical Modelling of Qubit Dynamics Controlled by Quantum Light

Addressing the dimensionality problem while maintaining accuracy

Oriol Bertomeu

Professorship for Quantum Computing
TUM School of Computation, Information and Technology
Technical University of Munich

Towards More Accurate Numerical Modelling of Qubit Dynamics Controlled by Quantum Light

Addressing the dimensionality problem while maintaining accuracy

Oriol Bertomeu

Thesis for the attainment of the academic degree

Master of Science (M.Sc.) in Computational Science and Engineering

at the TUM School of Computation, Information and Technology of the Technical University of Munich.

Examiner:
Prof. Dr. Christian B. Mendl

Supervisor:
Prof. Dr. Akihito Soeda

Submitted:
Munich, July 16, 2025

I hereby declare that this thesis is entirely the result of my own work except where otherwise indicated. I have only used the resources given in the list of references.

Munich, July 16, 2025 Oriol Bertomeu

Abstract

This thesis investigates numerical methods for modeling qubit dynamics controlled by quantum light, focusing on improving accuracy while addressing the curse of dimensionality in light-matter interactions. Qubits, the fundamental units of quantum information, are often implemented using matter states in architectures like superconducting circuits, trapped ions, and neutral atoms. While in some cases the control field can be described well by classical electromagnetism, in other cases, the ability to precisely control qubits via photonic interactions can be crucial. However, simulating these systems is computationally challenging due to the curse of dimensionality present in the Hilbert space of photonic fields.

We explore the Quantum Rabi Hamiltonian to model strong-coupling regimes, relaxing assumptions like the rotating wave approximation to enhance accuracy. Novel simulation techniques are developed, including truncation of the Hilbert space and coherent state basis methods, to balance computational efficiency and accuracy. Results demonstrate high atom state fidelity (up to 99.99% for coupling strengths $g \lesssim 0.1$) in truncated spaces, with significant computational speedups. However, coherent state simulations reveal challenges in numerical stability and fidelity, particularly for photonic states. Two-dimensional wave propagation visualizations highlight interaction phenomena like diffraction and re-emission.

This work advances the understanding of qubit control, supporting the development of fault-tolerant quantum computing by providing accurate, scalable simulation frameworks.

Acknowledgements

First of all, I would like to express my deep gratitude to my supervisor, Prof. Akihito Soeda, for his continued support. He has been the best mentor anybody could ask for, and it would have been impossible for me to complete this thesis without his dedication and patience. He has taught me what it means to be a researcher, and being able to work with him has been an incredible learning experience. His own passion for quantum computing has been key in nurturing and growing my own passion for the field, and I will miss our pen, paper, and hōjicha brainstorming sessions.

I also want to thank the rest of the staff at the National Institute of Informatics, especially Ms. Etsuko Kono and Prof. Emmanuel Planas, who work hard at creating a thriving research environment. Through the NII internship program, they support many young researchers who are starting out in their careers and offer them a unique chance to take part in a top research environment.

I want to give thanks to my advisor, Prof. Christian Mendl, for agreeing to take on the supervision of the project, showing interest in it, and providing guidance in spite of the physical distance.

I also appreciate the role of the Japan Science and Technology Agency, which, through the Moonshot Research and Development program, funded some of the computing resources necessary for the project.

Finally, I want to show my appreciation to my friends and family, especially my parents, for their continued support during my studies. Without having them by my side, it would not have been possible to complete this journey.

Contents

1 Introduction

Quantum computing represents a paradigm shift in computational science, harnessing quantum mechanical principles such as superposition, entanglement, and quantum interference to solve problems intractable for classical computers. Unlike classical bits, which represent either 0 or 1, quantum bits or qubits can exist in superpositions of states, enabling parallel computation and offering potential exponential speedups for specific tasks, such as factoring large numbers or simulating quantum systems.

Many of the currently most promising physical architectures for quantum computers, such as superconducting qubits, trapped ions, and neutral atoms, involve using matter states to describe qubits. For instance, electronic states are used in superconducting circuits, ionic excitation states are used in trapped ions, and hyperfine states are used in neutral atom computing. We investigate how photons interact with atoms, to be able to control such states with precisely designed pulses of quantum light, that is, individual photons.

The interaction between light and matter lies at the heart of quantum control in these architectures. Photons, as carriers of quantum information, can induce transitions in qubit states, enabling operations like state preparation and entanglement generation. However, accurately modeling these light-matter interactions poses significant challenges due to the complex dynamics involved. The photonic field, modeled as a bosonic mode, theoretically supports an infinite number of excitations, corresponding to a countably infinite-dimensional Hilbert space. This infinite dimensionality makes exact numerical simulations computationally infeasible, as no computer can store or process an infinite state space. To address this, truncation techniques are employed, limiting the number of excitations considered in the simulation. However, even with truncation, the state space grows exponentially with the number of modes or qubits, a phenomenon known as the curse of dimensionality. This exponential scaling severely limits the size of systems that can be simulated, necessitating innovative numerical methods to balance computational efficiency with simulation accuracy.

The curse of dimensionality becomes particularly pronounced in architectures where strong coupling between the control field and the qubits is present. In those cases, techniques that had been used previously in the literature, such as limiting the number of excitations present in the system, break down due to the collapse of assumptions present in the theory used to derive the Hamiltonian under which the Hamiltonian evolves. Namely, the rotating wave approximation, which implies that interactions between the field and the atoms that do not conserve the total number of excitations in the system can be neglected, cannot be used anymore, making new methods that take this into account necessary to accurately simulate strongly-coupled systems.

It is also key that the new methods to be developed are computationally efficient on top of being precise. That makes it possible to simulate systems with multiple atoms, which can be arranged in carefully constructed arrays in space to construct quantum optical devices. Quantum optical devices can be understood as elements that impart some kind of transformation on light as it goes through them: examples are polarizers, which selectively transmit light of a certain polarization while blocking the others, mirrors, which reflect waves, and beam splitters, which split an incident beam of light into two beams, among others.

What truly makes pursuing this field of research relevant is the fact that the greatest technical bottleneck for present-day quantum computing is error correction. Besides developing and implementing efficient

error-correcting codes, which is still a fundamental part of achieving fault-tolerant computation, for state-of-the-art error-correcting codes to work, it is necessary to guarantee that error rates stay below a certain threshold. Therefore, it is key to have accurate models that account for how quantum optical devices act on states, so that the error that a physical quantum optical device would generate compared to an ideal quantum gate can be explained away with a more precise description of how the quantum optical device acts on the atom.

The thesis is structured as follows: Chapter 2 provides a foundational overview of quantum computing concepts, including qubits, quantum gates, and error correction. Chapter 3 details the theoretical methods for modeling light-matter interactions, focusing on the Quantum Rabi Hamiltonian and simulation techniques. Chapter 4 outlines the problem setting, emphasizing the dimensionality challenge. Chapter 5 presents the novel contributions of this work, including advanced truncation and basis methods. Chapter 6 discusses numerical methods, such as sparse matrix storage and integration schemes. Chapter 7 evaluates the results, analyzing the fidelity of atom and photon states and visualizing wave propagation. Finally, Chapter 8 concludes with a discussion of limitations, future research directions, and the broader implications for fault-tolerant quantum computing.

By addressing the computational challenges of simulating qubit dynamics, this work aims to contribute to the development of scalable quantum technologies, paving the way for practical quantum computers capable of solving real-world problems efficiently.

2 Fundamental concepts

This section provides a foundational overview of quantum information processing, focusing on the core concepts of qubits, quantum gates, and quantum computer architectures and the critical role of error correction in achieving fault-tolerant quantum computation. Beginning with the definition and properties of qubits—the quantum analog of classical bits—we explore their representation as superpositions of basis states and their visualization on the Bloch sphere. The discussion then extends to quantum gates, which enable reversible transformations on qubit registers, and the physical platforms, such as superconducting circuits and ion traps, that implement these operations. Finally, we address the challenges of quantum error correction, emphasizing the need for precise modeling of qubit dynamics to mitigate errors and meet the stringent thresholds required for fault-tolerant quantum computing. As explored in subsequent sections, these concepts lay the groundwork for understanding the numerical modeling of qubit dynamics controlled by quantum light.

2.1 What is a qubit?

The term "qubit", or quantum bit, was first introduced in [Schumacher, 1995] to refer to spin-$\frac{1}{2}$ systems, to establish a quantum analogue of the noiseless coding theorem in classical computing. They are the fundamental units of quantum information. A formal definition relies on defining them as a linear combination of two base states $|0\rangle$ and $|1\rangle$ [Nielsen and Chuang, 2010], in the form $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. They generalize classical bits: for qubits, α and β can be arbitrary complex numbers, while for bits, one of the parameters must be one and the other zero.

The parameters can be interpreted as the probability of measuring the qubit in each basis state. Therefore, $|\alpha|^2$ corresponds to the probability of measuring $|0\rangle$ and $|\beta|^2$ to the probability of measuring $|1|\rangle$. Making sure this probability interpretation stays valid imposes an additional constraint on the parameters, referred to as normalization: $|\alpha|^2 + |\beta|^2 = 1$.

This normalization condition can be enforced by rewriting the qubit state as

$$|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\phi}\sin(\theta/2)|1\rangle \tag{2.1}$$

States of this form are referred to as pure qubit states. A useful way to visualize them is the Bloch sphere [2.1]. The parameters ϕ, θ in the parameterization can be mapped to spherical coordinates on a shell, thereby providing a handy way to depict pure states.

Nevertheless, states consisting of statistical mixtures of pure states can also arise. These originate in interactive processes such as decoherence [Bacon, 2003] [Schlosshauer, 2019] and quantum noise [Clerk et al., 2010]. These can only be expressed through a density matrix

$$\rho = \sum_i \lambda_i |\psi_i\rangle\langle\psi_i| \tag{2.2}$$

In the Bloch sphere, they correspond to points in the interior of the spherical shell. Density matrices degenerate to pure states in the case $\lambda_0 = 1$, $\lambda_{i>0} = 0$, implying that a pure state also has a density matrix $|\psi\rangle\langle\psi|$.

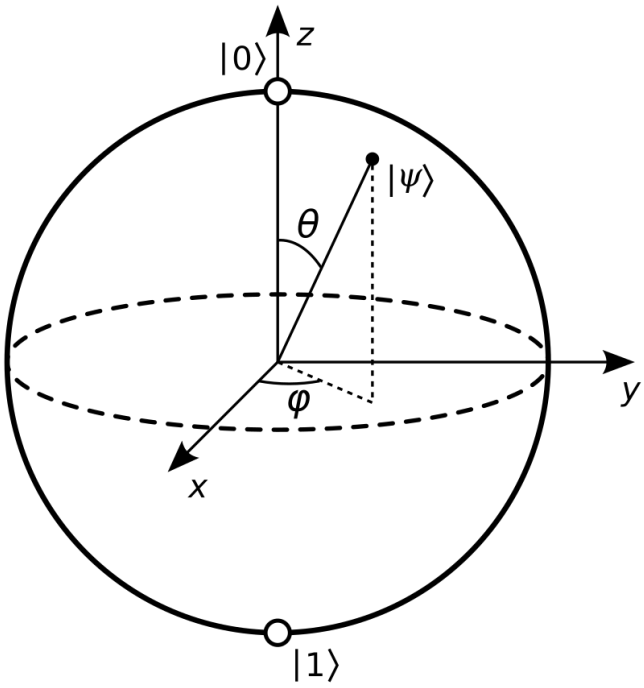

Figure 2.1 Bloch sphere representation of a qubit state

2.2 Quantum gates and computers

A set of qubits is referred to as a register when grouped. An n-qubit register can be written as

$$|\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle \otimes ... \otimes |\psi_n\rangle \equiv |\psi_1\rangle|\psi_2\rangle...|\psi_n\rangle \tag{2.3}$$

Where every $|\psi_i\rangle$ is a state of the form of (2.1). We can transform these analogously to classical logic gates with quantum gates [Barenco et al., 1995], which are basic circuits that apply some operation on a subset of the register. A key property of quantum gates is that they are reversible in contrast to classical gates, which is key to modeling quantum phenomena, since non-reversibility would introduce loss of information [Shah, 2010]. Consequently, quantum gates can be modeled by unitary operators acting on qubits.

A quantum gate operating on qubits i, j of the register can be represented as

$$|\psi'\rangle = (\mathbf{1} \otimes U_{ij} \otimes \mathbf{1})\ |\psi_1\rangle...|\psi_i\rangle|\psi_j\rangle...|\psi_n\rangle \equiv U\ |\psi\rangle \tag{2.4}$$

for a pure state. For a density matrix state

$$\rho' = (\mathbf{1} \otimes U_{ij} \otimes \mathbf{1})|\psi_1\rangle...|\psi_i\rangle|\psi_j\rangle...|\psi_n\rangle\langle\psi_1|...\langle\psi_i|\langle\psi_j|...\langle\psi_n|(\mathbf{1} \otimes U_{ij}^{\dagger} \otimes \mathbf{1}) \equiv U\rho\, U^{\dagger} \tag{2.5}$$

A reduced state, characterizing only a subsystem of the register, can be extracted by taking the partial trace of the density matrix:

$$\rho_A = Tr_{\overline{A}}(\rho) = \sum_k \langle k|\rho|k\rangle \tag{2.6}$$

with the $|k\rangle$ being the basis elements of the subspaces we trace over, $\overline{A}$ (the complement of A, which contains all subspaces but A).

An important fact is that pure states stay pure under unitary action. However, the same cannot be said for subsystems. In these cases, states can become entangled, a key property they must have before we can perform useful computation.

A quantum computer is nothing more than a register of which, through electronics, we can control the state by implementing quantum gates [Wong, 2025].

2.2.1 Quantum computer architectures

To implement a working quantum computer, it is necessary to pick a physical platform to support qubits and operations. DiVincenzo's criteria summarize the requirements that such an architecture must satisfy [DiVincenzo, 2000]:

1. *A scalable physical system with well-characterized qubits*: qubits must be two-level systems, the behavior of which is accurately known. The energy gap between the two levels $|0\rangle$ and $|1\rangle$ must also be big enough to prevent thermal noise from introducing bit flip errors.

2. *The ability to initialize the state of the qubits to a simple fiducial state, such as* $|000...\rangle$: it must be possible to prepare an initial state quickly and with high fidelity.

3. *Long relevant decoherence times, much longer than the gate operation time*: the time scale for interaction with the environment outside the computer must be sufficiently long to perform practical computations before it loses its information.

4. *A "universal" set of quantum gates*: the computer must be able to perform arbitrary operations.

5. *A qubit-specific measurement capability*: measurements must be performed to read out the results of computations.

Several competing architectures aim to fulfill these criteria. Each has its advantages and drawbacks and finds itself at a different stage of development. These are defined by the method they use to construct qubits. They can be classified into three types: (1) intrinsic two-level systems, (2) two-level subset systems, and (3) engineered two-level systems [Chae et al., 2024]. (1) and (2) leverage the quantum mechanical character of nature so that they can be referred to as "natural qubits". However, only (1) has two possible levels, while (2) has more, and care must be taken in designing the qubit so that the probability of entering other states is sufficiently low. Meanwhile, (3), which can be thought of as "synthetic qubits", uses creative Hilbert spaces of, for instance, circuits, to design qubits.

Physical systems must also be able to demonstrate entanglement between qubits to be useful for quantum computation. The main challenge to achieving entanglement is that qubits must be able to interact with each other but not with the environment to avoid decoherence.

Superconducting qubit computers

Superconducting quantum computers use superconducting qubits as physical qubits. These are solid-state circuits built using techniques borrowed from conventional integrated circuits, leveraging existing know-how in the semiconductor industry [Devoret et al., 2004].

Superconducting qubits take advantage of the zero-resistance property of superconductors to provide dissipation-free operation, which is key for quantum mechanical behavior. Superconducting tunnel junctions, also known as Josephson junctions, are used to implement nonlinearity and non-dissipativity. These consist of a sandwich of two superconducting films separated by an insulating layer, through which Cooper pairs can tunnel. The most common fabrication material is aluminum, since its oxidation already produces an insulating layer with good properties.

There are multiple superconducting qubit circuits in use [Büttiker, 1987] [Barone and Paterno, 1982]. They are modeled as non-linear oscillators, and the qubit states are identified as the circuit's fundamental and first excited states. Nonlinearities are exploited to ensure the gap between the second and the first excited state is much larger than that between the fundamental and the first excited state, guaranteeing the qubit has a high probability of remaining well-behaved.

Single-qubit gates are implemented via microwave pulses that are precisely tuned to the transition frequency of the qubit [Krantz et al., 2019]. For two-qubit gates, the qubits are allowed to interact through a resonator, and then microwave pulses are applied.

Their main advantages lie in fast gate speeds, of the order of the tens of nanoseconds, and mature fabrication methods. At the same time, their most considerable drawbacks are short decoherence times of tens of microseconds, onerous cryogenic requirements, and limited connectivity between qubits due to being laid out in a 2D lattice, which leads to scalability challenges.

Ion trap computers

By confining cold ions in an ion trap and making them interact with laser beams, a quantum computer can be implemented [Cirac and Zoller, 1995]. Qubits can be individually manipulated by directing different laser beams to them. CNOT gates can be implemented by exciting the ions' collective quantized motion with lasers, originating from a repulsive Coulomb interaction.

Internal electronic states of the ions are used to represent the qubit states $|0\rangle$ and $|1\rangle$. The magnitude of the energy splitting between states depends on the kind of qubit used, ranging from tens of MHz to hundreds of THz [Bruzewicz et al., 2019]. Initialization is performed via optical pumping into $|1\rangle$, and, for readout, a resonant laser couples $|1\rangle$ to a transition that scatters many photons, which are detected at a detector, while $|0\rangle$ remains dark. Both processes take less than 1 ms to complete.

Entanglement is realized via a Cirac-Zoller gate, which implements a controlled phase shift between two ions trapped in a linear Paul trap by addressing them with a sequence of laser pulses [Schmidt-Kaler et al., 2003].

Their strong points lie in long coherence times, up to several minutes, and high gate fidelity. Qubits can all couple with each other by sharing an ion trap, so full connectivity is easy to implement. Furthermore, the temperature at which they need to be cooled down is not as low as that of superconducting computers. The weak points are slow gate speeds of tens of microseconds, the engineering complexity of precise laser control, and scaling issues from ion trap design limits. Finally, their fabrication processes are less mature, since semiconductor industry processes cannot be leveraged.

2.3 Error correction codes and fault tolerance

A problem common to all quantum computer physical implementations is that, in contrast to classical bits, it is difficult to isolate qubits from external noise, which leads to errors during computation. Furthermore, there are unique challenges due to the quantum nature of qubits. The no-cloning theorem forbids creating copies of arbitrary quantum states [Wootters et al., 1982], and measurement of a quantum state destroys it by collapsing its wavefunction.

An error process can be modeled as a shift of the state in the Bloch sphere. It can be shown [Roffe, 2019] that such an operation can be modeled as a combination of

1. A bit flip error $a|0\rangle + b|1\rangle \rightarrow a|1\rangle + b|0\rangle$.
2. A phase flip error $a|0\rangle + b|1\rangle \rightarrow a|0\rangle - b|1\rangle$.

(1) is the analog of a classical bit flip error, while (2) is a purely quantum effect. Therefore, in contrast to classical computing, two types of errors must be corrected at once, increasing the complexity of the systems required to solve the problem.

Error correction codes work by using extra qubits to encode logical qubits. Projective measurements are then carried out to stabilize the quantum state in the codespace, which refers to the set of physical qubit states where there is no error. For that, ancilla qubits are used: after the projective measurements, they are measured, which is referred to as syndrome extraction. After m stabilizer measurements, a m-bit syndrome is obtained, which informs us of the optimal operation to perform on the qubits to recover the logical state in the codespace.

The error correction code in most state-of-the-art quantum computers is a surface code. Surface codes are characterized by their modularity, which makes scaling straightforward. Furthermore, they can be implemented with only nearest-neighbor interactions, avoiding low-fidelity long-range interactions [Fowler et al., 2012].

A quantum error correction is called fault-tolerant if it can account for errors occurring at any point in the circuit, ensuring small errors do not propagate through the circuit and magnify. Making a given error correction code fault-tolerant adds considerable overhead in the amount of ancilla qubits necessary to recover the codespace state.

A key fact that must be borne in mind when implementing error correction schemes is the presence of code thresholds. The threshold theorem states that adding physical qubits to the code results in lower logical error rates only if the error rate of the individual physical qubits stays below a certain threshold. Thresholds depend on the exact code, but are normally around 10%.

The presence of this threshold underscores that, despite the existence of error correction codes through which we can perform computation with high fidelity, there is still a need for accurate modeling of qubit dynamics, especially of processes like entanglement gates, since they are normally executed with lower fidelity. Having good models of qubit control methods, like laser-qubit interaction, is key to reducing physical error and successfully implementing fault-tolerant error correction.

3 Methods

This section lays the theoretical groundwork for modeling the dynamics of qubits interacting with quantum light, focusing on the state space, Hamiltonians, and simulation techniques critical to understanding light-matter interactions. We begin by defining the composite Hilbert space of the system, comprising the infinite-dimensional Fock space of the photonic field and the two-level qubit subsystem, and extend this to multi-mode scenarios for realistic photon propagation. The discussion then shifts to Hamiltonians, introducing the Jaynes-Cummings and Quantum Rabi models to describe light-matter interactions, emphasizing the latter's applicability in strong-coupling regimes relevant to quantum computing. We explore simulation techniques for temporal evolution across Schrödinger, Heisenberg, and interaction pictures, alongside state representations, including product, entangled, and coherent states, and their characterization via density operators and entanglement measures. Finally, the Wigner picture is introduced as a phase-space representation to bridge classical and quantum descriptions, providing insights into the quantum nature of number and coherent states. These concepts underpin the numerical methods developed for simulating qubit dynamics in subsequent sections.

3.1 State space

Our setup comprises two distinct subsystems. Assuming only one mode is present, the Hilbert space of the whole system can be said to be the tensor product of both subsystems:

1. Field: The electromagnetic field is quantized as a harmonic oscillator with infinitely many eigenstates $|n\rangle = |0\rangle, |1\rangle, |2\rangle, ...$ The Hilbert space for this field is the Fock space $\mathscr{H}_f = \{|0\rangle, |1\rangle, |2\rangle, ...\}$
2. Two-level atom: The atom is a qubit with two energy states, the ground state and the excited state, that we label $|g\rangle$ and $|e\rangle$. It is characterized by a resonant frequency ω_a linked to the energy gap between qubit states by $E = \hbar\omega_a$. We write the Hilbert space spanned by the qubit basis states as $\mathscr{H}_a = \{|g\rangle, |e\rangle\}$.

So, $\mathscr{H} = \mathscr{H}_f \otimes \mathscr{H}_a$, and a state from $\mathscr{H}$ can be written as $|n, g\rangle$ or $|n, e\rangle$.

However, one mode is insufficient to obtain accurate physical descriptions of photon propagation. The propagating photon must be modeled as a wave packet, which is, in the ideal case, a superposition of infinitely many modes. The actual field Hilbert space is, therefore, the tensor product of as many Fock spaces as modes are present:

$$\mathscr{H}_F = \bigotimes_{i=1}^{N_M} \mathscr{H}_f \tag{3.1}$$

3.2 Hamiltonians

We model interactions between light and matter with Hamiltonians. Hamiltonians are operators acting on the Hilbert space where the states of the problem are defined, and they govern the temporal evolution of the system according to the time-dependent Schrödinger equation [Shankar, 1980]:

$$i\hbar\frac{d}{dt}|\Psi(t)\rangle = H|\Psi(t)\rangle \tag{3.2}$$

It is intimately related to the energy of the system, as energy levels correspond to its eigenvalues:

$$H|\Psi_n\rangle = E_n|\Psi_n\rangle \tag{3.3}$$

Choosing a suitable Hamiltonian is crucial when deciding how to model a quantum system, since it contains all the information about how the system will evolve.

Hamiltonians are often written in terms of ladder operators. Ladder operators can be creation or annihilation operators, increasing or decreasing the number of excitations in a number state by one. That way, the following are fulfilled:

$$a|n\rangle = \sqrt{n}|s, n-1\rangle \qquad a^\dagger|n\rangle = \sqrt{n+1}|s, n+1\rangle \tag{3.4}$$

A number operator N can be defined as $N = a^\dagger a$. It also fulfills $N|n\rangle = n|n\rangle$, which is a property that will prove helpful in calculating expectation values. An important property is $a|s, 0\rangle = 0$, since photon numbers cannot be negative.

Similarly, raising and lowering operators can be defined for the atom subsystem. These act on the compound states as $\sigma^+|g, n\rangle = |e, n\rangle$ and $\sigma^-|e, n\rangle = |g, n\rangle$. A key difference with the bosonic ladder operators defined for the photons is that, since the atomic Hilbert space has finite dimension, in contrast to the Fock space, it holds that $\sigma +^+ |e, n\rangle = 0$ as well.

3.2.1 Jaynes-Cummings Hamiltonian

The Jaynes-Cummings Hamiltonian is one of the simplest models for light-matter interaction. It can be obtained by quantizing the electrical and magnetic fields [Mandel and Wolf, 1995] and then writing the Hamiltonian as an energy-density operator. A dipole approximation is then used for the interacting term. Finally, the rotating wave approximation (RWA) is employed, which neglects terms with dependencies on the sum of atom and field frequencies, yielding the following Hamiltonian:

$$H_{JC} = H_F + H_A + H_I \tag{3.5}$$

$$H_F = \sum_k^{N_M} \hbar\omega_k a_k^\dagger a_k \tag{3.6}$$

$$H_A = \sum_j^{N_A} \hbar\Omega_j \sigma_j^+ \sigma_j^- \tag{3.7}$$

$$H_I = \sum_k^{N_M}\sum_j^{N_A} \hbar g(\sigma_j^+ a_k + \sigma_j^- a_k^\dagger) \tag{3.8}$$

The Jaynes-Cummings Hamiltonian conserves the total number of excitations, which can be intuitively understood by observing the interaction Hamiltonian H_I. The two subsystems can interact in two separate ways: either the atom gains an excitation and the photon loses one, or the excitation is lost by the atom and gained by the photon. This property is a direct consequence of the RWA, as discussed in section 3.2.2.

It can be shown [Larson and Mavrogordatos, 2024] that the Jaynes-Cummings Hamiltonian can be rewritten in a way that it only depends on the set of parameters $\Delta_{j,k} = \Omega_j - \omega_k$, which is the detuning between each mode and the atom frequencies, and the coupling parameter g, which determines the strength of the interaction between atom and photon.

Even for only one mode and one atom, the dynamics of the Jaynes-Cummings Hamiltonian are pretty rich. In this case, they can be derived analytically. For instance, Rabi oscillations can be observed in resonance ($\Delta = 0$). These are harmonic oscillations in the probability of detecting the excited state, with frequency $\Omega_R = 2g\sqrt{n+1}$.

3.2.2 Quantum Rabi Hamiltonian

The derivation of the Jaynes-Cummings Hamiltonian is reliant on several assumptions, the most relevant of which are:

1. Dipole approximation: the atoms are assumed to be small compared to the relevant wavelengths of the electromagnetic waves mediating the interaction.

2. Two-level approximation (TLA): excited states from the second on are assumed to be unreachable due to being far detuned.

3. Rotating wave approximation (RWA): interactions that might not conserve the total number of excitations in the system are not considered.

These assumptions are connected, and the key factor that determines their validity is the relative size of the coupling parameter g: as g increases, they progressively break down. However, this does not necessarily happen simultaneously for all of them. As it turns out, the RWA is the first assumption to break down. For this reason, it is helpful to contemplate the case where the other approximations still work, but the RWA is discarded.

Doing this yields the Quantum Rabi model:

$$H_{QR} = H_F + H_A + H_I \tag{3.9}$$

$$H_F = \sum_{k}^{N_M} \hbar\omega_k a_k^\dagger a_k \tag{3.10}$$

$$H_A = \sum_{j}^{N_A} \hbar\Omega_j \sigma_j^+ \sigma_j^- \tag{3.11}$$

$$H_I = \sum_{k}^{N_M}\sum_{j}^{N_A} \hbar g(\sigma_j^+ + \sigma_j^-)(a_k + a_k^\dagger) \tag{3.12}$$

The Quantum Rabi model improves on the performance of the Jaynes-Cummings Hamiltonian for stronger coupling. It is often interesting to have strong coupling in quantum computers, since that reveals entanglement dynamics to a greater degree than weak coupling. For this reason, we have focused our study on Quantum Rabi Hamiltonians.

3.2.3 Coupling parameter

In standard Jaynes-Cummings and Quantum Rabi Hamiltonians, the coupling parameter is assumed to be a real, constant number. However, it is possible to derive its value from the dipole approximation, and it actually has a dependence on frequency:

$$g(j,k) = -\frac{i\hbar}{2\varepsilon_0 L}\sqrt{\omega_k} D_j e^{i\vec{k}\cdot\vec{r}_j} \tag{3.13}$$

However, after a gauge transformation, $g(j,k)$ can be taken to be a positive real number following the procedure outlined in [Larson and Mavrogordatos, 2024], and by grouping all other parameters together and using $\omega_k = c|k|$, the coupling constant can be rewritten as:

$$g(j,k) = g_0\sqrt{|k|} \tag{3.14}$$

This is the value we implement, with the goal of minimizing model assumptions and thereby maximizing fidelity with respect to real behavior.

3.3 Simulation techniques

3.3.1 Time evolution

In quantum mechanics, the time evolution of a state is given by a unitary operator [Sakurai, 1994], which is related to the Hamiltonian governing the system:

$$U(t,t_0) = exp\Big[-\Big(\frac{i}{\hbar}\Big)\int_{t_0}^{t} dt' H(t')\Big] \tag{3.15}$$

If the Hamiltonian under consideration has no time dependence, as is the case for photon-atom interaction Hamiltonians (Jaynes-Cummings, Quantum Rabi), then this can be further simplified as follows:

$$U(t,t_0) = e^{-iH(t-t_0)/\hbar} \tag{3.16}$$

This operator can be understood as a map from the state at time t_0 to the state at time t:

$$|\Psi(t)\rangle = U(t,t_0)|\Psi(t_0)\rangle \tag{3.17}$$

From now on, we will take $t_0 = 0$ without losing generality. The evolved state can then be used to study how expectation values depend on time. For a generic observable M:

$$\langle M(t)\rangle = \langle\Psi(t)|M|\Psi(t)\rangle = \langle\Psi|U^\dagger(t)M|\Psi(t)|U(t)\rangle \tag{3.18}$$

This expression suggests an alternate interpretation of time dependence where, instead of states, what changes with time are observables. By writing $M(t) = U^\dagger(t)M|\Psi(t)|U(t)$, the states can be understood to be constant.

Working with time dependence on the observables is called the Heisenberg picture. Meanwhile, when time dependence is on the states, we are said to be working in the Schrödinger picture.

There is, however, another picture in the middle of the Schrödinger and Heisenberg pictures, called the interaction picture. It leverages the properties of the exponential by rewriting the Hamiltonian as $H = H_0 + H_I$. H_0 is the well-understood, analytically solvable part of the Hamiltonian, while interesting interactions we wish to study are included in H_I. Then, the evolution of an arbitrary observable can be rewritten as

$$\langle M(t)\rangle = \langle\Psi|e^{iH_0t/\hbar}e^{iH_It/\hbar}Me^{-iH_It/\hbar}e^{-iH_0t/\hbar}|\Psi\rangle = \langle\Psi_I(t)|M_I(t)|\Psi_I(t)\rangle \tag{3.19}$$

with

$$|\Psi_I(t)\rangle = e^{-iH_0t/\hbar}|\Psi\rangle \tag{3.20}$$

$$M_I(t) = e^{iH_It/\hbar}Me^{-iH_It/\hbar} \tag{3.21}$$

The interaction is often helpful because it allows us to simulate the interaction term (H_I) exclusively, while the free (H_0) term, which often has been well-studied in previous work, can be accounted for analytically.

A drawback of the interaction picture, however, is that, in the case where the evolution of $|\Psi_I(t)\rangle$ is not known analytically, the evolutions of state and observables must be computed numerically, which can be expensive and even completely negate the advantages of working on the interaction basis.

Furthermore, analogous evolution equations exist for each picture. These describe the evolution of states and operators for the Schrödinger and the Heisenberg picture, respectively.

For the Schrödinger picture, the evolution of a state is given by the Schrödinger equation:

$$i\hbar\frac{d}{dt}|\Psi(t)\rangle = H|\Psi(t)\rangle \tag{3.22}$$

The validity of this equation can be readily verified by inserting $|\Psi(t)\rangle = U(t,t_0)|\Psi(0)\rangle$, upon which one can check that the differential equation is indeed fulfilled.

For the Heisenberg picture, with time dependence on the operators, time evolution is instead governed by the Heisenberg equation:

$$\frac{dA(t)}{dt} = \frac{i}{\hbar}[H(t), A(t)] + \frac{\partial A}{\partial t} \tag{3.23}$$

Where $\frac{\partial A}{\partial t}$ is the derivative with respect to explicit time dependence in the Schrödinger picture.

The Heisenberg equation illustrates that, for operators without explicit time dependence, commutation with the Hamiltonian

$$[H, A] = 0$$

is equivalent to them being constants of motion, since, then, $\frac{dA(t)}{dt} = 0$.

3.4 States

States encode information about the configuration of the quantum system at a given point. As discussed before, they can be understood as elements of a Hilbert space characterizing the system's dynamics.

It is interesting to discuss their characterization, as that has implications for understanding the different ways systems interact with each other and the environment.

In the simplest case, the state of the system composed of many subsystems can be written as a Kronecker product of vectors that are in the Hilbert spaces of the subsystems:

$|\psi\rangle = |\psi_A\rangle \otimes |\psi_B\rangle \otimes ... \otimes |\psi_N\rangle$

This can be rewritten as a vector acting on a different Hilbert space. For instance:

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle_A + |1\rangle_A) \otimes \frac{1}{\sqrt{2}}(|0\rangle_B - |1\rangle_B) \equiv \frac{1}{2}(|00\rangle - |01\rangle + |10\rangle - |11\rangle) \tag{3.24}$$

We have identified the state $|a\rangle_A \otimes |b\rangle_B$, the outer product of vectors of two different states, as a vector in a different Hilbert space, $|ab\rangle$, which is equivalent to relabeling the states $|00\rangle \to |0\rangle, |01\rangle \to |1\rangle, |10\rangle \to |2\rangle, |11\rangle \to |3\rangle$, since they are all orthogonal and, in practice, it is possible to work with them without taking their origin as states resulting of the product of states in subsystems into account.

We can also observe that the dimensionality of the combined Hilbert space is equal to the product of the dimensionalities of the subsystem states.

However, the reverse proposition is not true. Given a vector in the Hilbert space of the combined subsystems, it will not be generally possible to rewrite it as a product of two states in the Hilbert space of the subsystems. For example, the Bell state

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$$

cannot be rewritten as a product of two states in the subsystem spaces.

If a state in a combined Hilbert space can be rewritten as a product of states in the spaces of the subsystems, we will refer to it as a product state. Otherwise, it can be considered an entangled state.

It is interesting to consider the problem of quantifying how much a particular state is entangled. A generic composite state can be written as

$$|\Psi\rangle = \sum_{i=1}^{d_A} \sum_{j=1}^{d_B} c_{ij} |i, j\rangle \tag{3.25}$$

All information about the state is encoded in the coefficients c_{ij}, which can be understood as the matrix elements of C. However, we could have also chosen another basis of the composite system to represent the state. In particular, selecting a basis in which the matrix of coefficients is diagonal:

$$|\Psi\rangle = \sum_{k=1}^{d} d_k |u_k, v_k\rangle \tag{3.26}$$

Where $d = min(d_A, d_B)$ and $\sum_k d_k^2 = 1$ due to normalization.

This expansion is the Schmidt decomposition of a state, equivalent to the singular value decomposition of C.

After calculating the Schmidt decomposition of a state, it is easy to determine if it is a product or an entangled state. In the case where one of the coefficients equals one and the rest are zero, we have a product state. Otherwise, the state is entangled.

Given a state $|\Psi\rangle$ composed of subsystems A and B, it is often desirable to recover information describing only the state of one of the subsystems. The respective reduced states characterize these:

$$\rho_A = tr_B(|\Psi\rangle\langle\Psi|) = \sum_b \langle b|_B \Psi\rangle\langle\Psi|b\rangle_B \tag{3.27}$$

$$\rho_B = tr_A(|\Psi\rangle\langle\Psi|) = \sum_a \langle a|_A \Psi\rangle\langle\Psi|a\rangle_A$$

The sums go over the basis of the subspace being traced over.

It is important to remark that these ρ_A, ρ_B that result from the reduced states are nevertheless not vectors in the Hilbert spaces of A and B. Instead, they are density operators, $\rho \in L(H)$, characterized by being hermitian, positive semi-definite, and having unit trace.

Density operators are a characterization of quantum systems offering an alternative to states. Furthermore, they allow us to characterize systems that are statistical mixtures of states. Such states, which are called mixed states, can be written as

$$\rho = \sum_i p_i \rho_i = \sum_i p_i |\Psi_i\rangle\langle\Psi_i| \tag{3.28}$$

In the case where one of the p_i is equal to one and the rest are zero, this reduces to $\rho = |\Psi\rangle\langle\Psi|$ and the density operator ρ can be identified with the state $|\Psi\rangle$. $|\Psi\rangle$ or ρ are then understood to refer to a pure state.

Mixedness can therefore be understood as the state's rank: pure states have rank one, while mixed states have a greater rank. However, states like $0.999|0\rangle\langle 0| + 0.001|1\rangle\langle 1|$ still have rank two, but are very close to rank one states.

To sidestep this, the most common measures of mixedness used are the following:

- Purity: for a state ρ, its purity is given by $\text{tr}(\rho^2) = \sum_i^d \lambda_i^2$, with λ_i the eigenvalues of ρ. For a pure state, one of the λ_i is one while the others are zero, and therefore $\text{tr}(\rho^2) = 1$. Otherwise, since $0 \leq \lambda_i \leq 1$, it can be guaranteed that $1/d \leq \lambda_i < 1$.
- Von Neumann entropy: for a state ρ, it is given by $S(\rho) = -\text{tr}(\rho \log_2(\rho)) = -\sum_i^d \lambda_i \log_2(\lambda_i)$. It can be understood as the Shannon entropy of the eigenvalues of ρ. For a pure state, $S(\rho) = 0$. Also, it can be shown that $0 \leq S(\rho) \leq \log_2(d)$.

The state $\rho = \mathbf{1}/d$ is referred to as the maximally mixed state, as it corresponds to both the lower bound on purity and the upper bound on von Neumann entropy.

Going back to product and entangled states, we can notice that a composite state $\Psi\rangle$ is a product state if and only if both reduced density operators correspond to pure states. Therefore, entanglement corresponds to the mixedness of the reduced density operators, and we can measure entanglement by measuring the mixedness of the reduced states.

Thus, we introduce the entropy of entanglement of a state $|\Psi\rangle$ as:

$$E(\Psi) = -\text{tr}_A[\rho_A \log_2(\rho_A)] = -\text{tr}_B[\rho_B \log_2(\rho_B)] = -\sum_k d_k^2 \log_2(d_k^2) \tag{3.29}$$

As expected, for a product state, $E(\Psi) = 0$. The maximal entanglement entropy, $E = \log[min(d_A, d_B)]$, is achieved for states with equal Schmidt coefficients.

Entanglement can also be understood as a correlation measurement between two systems after interaction. Suppose two systems are initialized separately and put into contact. In that case, we expect them to evolve from a product state, given that the systems had no knowledge of each other before the interaction, to a state where they have mutually influenced each other. Therefore, we expect to see high entanglement in strong interactions, while we will have low entanglement for weak interactions.

It is interesting to study the temporal evolution of purity:

$$\text{tr}(\rho(t)^2) = \text{tr}(U(t)\rho U^\dagger(t) U(t)\rho U^\dagger(t)) = \text{tr}(U(t)\rho\rho U^\dagger(t)) = \text{tr}(\rho(t)^2) \tag{3.30}$$

For the whole isolated system, purity is a constant of motion. Therefore, if we initialize the composite system as a product of pure states, it will remain pure for all its temporal evolution.

The significance of this fact for simulation lies in the fact that, since the state will never become a statistical mixture, we do not need to model it with a density operator matrix. A vector suffices, saving considerable space in memory.

Nevertheless, it must be underlined that this fact only holds for the state of the composite system. The evolution of the subsystems A and B is generally not unitary, so the purity of the reduced states rho_A and ρ_B cannot be said to be constant.

3.4.1 Coherent states

We have used the so-called number states when defining the Hamiltonians governing the interactions we want to study. These are the eigenstates of the harmonic oscillator: in our notation, they are the eigenstates of the number operator $N = a^\dagger a$. These reflect the number of photons in a state: $|0\rangle$ has zero photons, $|1\rangle$ has one photon, $|2\rangle$ has two photons, and so on.

These states encode phenomena clearly of a quantum nature: light is made up of discrete photons, and through the application of ladder (annihilation and creation) operators, more photons can be subtracted or added to the system.

This raises the question of how to model the classical limit. In classical physics, light is a continuous wave, and there is no discretization. We can think of this as meaning that, in the classical limit, subtracting a photon from the state leaves it unchanged.

This argument naturally leads us to an analogy from classical limit states to the eigenstates of the annihilation operator, which are referred to as coherent states. We write their complex eigenvalue as α and the state itself as $|\alpha\rangle$:

$$a|\alpha\rangle = \alpha|\alpha\rangle \tag{3.31}$$

They can be expressed as a superposition of infinitely many number states.

$$|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle \tag{3.32}$$

Finally, it is often helpful to work in a picture composed of exclusively coherent states, and altogether forgo expansions in terms of number states. In these cases, the coherent state is defined as the result of applying a displacement operator on the vacuum state:

$$|\alpha\rangle = D(\alpha)|0\rangle = e^{\alpha a^\dagger - a^\star a}|0\rangle \tag{3.33}$$

The reason why $D(\alpha)$ is called a displacement operator is that, if we understand the vacuum state as being equivalent to $|\alpha = 0\rangle$, successive applications of displacement operators have the property $D(\alpha)D(\beta)|0\rangle = D(\beta)D(\alpha)|0\rangle = D(\alpha+\beta)|0\rangle = |\alpha+\beta\rangle$. Therefore, the action of a displacement operator $D(\alpha)$ on a coherent state can be understood as "displacing" it by α in the complex plane.

Since we work on a compound Hilbert space $\mathscr{H} = \mathscr{H}_F^M \otimes \mathscr{H}_A$, it is important to bear in mind that the displacement operator only acts on the bosonic part of the state at initialization: the actual state is $|\alpha; s\rangle$, with $s \in \{g, e\}$ being the state of the two-level atom.

Such a state can be initialized from the vacuum state by applying the operator

$$D(\alpha) \otimes |0; s\rangle = |\alpha; s\rangle \tag{3.34}$$

However, for the sake of brevity, we will write $D(\alpha)|0\rangle$ whenever we do not wish to consider the behavior of the atomic part of the system or when it is trivial.

A key fact that must be considered when working with coherent states is that, unlike number states, they are not orthonormal. In fact,

$$\langle\beta|\alpha\rangle = \langle 0|D(\beta)D(\alpha)|0\rangle = e^{\frac{1}{2}(|\alpha|^2+|\beta|^2)-\beta^\star\alpha} = e^{-\frac{1}{2}|\beta-\alpha|^2}e^{-\mathrm{Im}(\alpha\beta^\star)} \tag{3.35}$$

Meaning that, while we can consider that states with $|\beta-\alpha| \gg 1$ (very far apart from each other) almost orthogonal, as their overlap decays exponentially, we cannot do that in general, as for number states.

This fact has profound implications when defining density matrices and calculating expectation values. As long as the basis of the state under consideration is orthonormal, then traces do not depend on the exact choice of the basis, as $\mathrm{tr}(\rho) = \mathrm{tr}(U\rho U^\dagger)$ due to the cyclic property of the trace. However, this property is not fulfilled when working in non-orthonormal bases, and expectation values become ill-defined.

Coherent states can be extended to multi-mode states by writing new multi-mode ladder operators that are a superposition of mode ladder operators:

$$\tilde{a} = \sum_k c_k a_k \tag{3.36}$$

$$\tilde{a}^\dagger = \sum_k c_k^\star a_k^\dagger \tag{3.37}$$

The choice of c_k is in principle free, with the only constraint being normalization:

$$\sum_k |c_k|^2 = 1 \tag{3.38}$$

However, in practice, the constants c_k will be chosen in such a way as to replicate a desired wave packet shape.

3.4.2 Wigner picture

The Wigner picture is a method to bring phase space representations from classical to quantum mechanics. It was developed by Paul Wigner in 1932 [Weinbub and Ferry, 2018], as a way to recover a phase space distribution function from a quantum wave function:

$$W(x,p) = \frac{1}{\pi\hbar}\int_{-\infty}^{\infty} dy \psi^\star(x+y)\psi(x-y)e^{2iyp/\hbar} \tag{3.39}$$

However, the Wigner function cannot be equated to a probability distribution, since, while it is real, it is not positive definite, which is a key difference from classical phase space treatments: we can see negativity of the Wigner function as proof of the exclusively quantum nature of a process, interaction, or phenomenon.

Even though Paul Wigner originally defined his function with a wavefunction as an argument, it is convenient to generalize it so it can also be applied to density matrix states. A position wavefunction can be obtained from a state vector by $\psi(x) = \langle x|\psi\rangle$. Substituting that into the Wigner function expression and identifying $\rho = |\psi\rangle\langle\psi|$, we can rewrite it as

$$W(x,p) = \frac{1}{\pi\hbar}\int_{-\infty}^{\infty} dy \langle x+y|\rho|x-y\rangle e^{2iyp/\hbar} \tag{3.40}$$

Which is often called the Weyl transform. By taking time derivatives, a time evolution equation can be derived for the Wigner distribution. Some interesting mathematical properties are that, even though the Wigner function is not a true probability function, its marginal distributions on x and p, $W(x) = \int_{-\infty}^{\infty} dp W(x,p) = \langle x|\rho|x\rangle$ and $W(p) = \int_{-\infty}^{\infty} dx W(x,p) = \langle p|\rho|p\rangle$, are (from the positive definiteness of the density matrix). Furthermore, for a pure state, the following bound applies:

$$-\frac{2}{h} \leq W(x,p) \leq \frac{2}{h} \tag{3.41}$$

Meaning that, in the classical limit $\hbar \to 0$, the bound disappears.

To perform calculations, it can also be helpful to rewrite the expression for the Wigner function in terms of the Fourier transform of the displacement operator:

$$\delta(a-\beta) = \frac{1}{\pi^2}\int d^2\gamma e^{\gamma^\star\beta-\gamma\beta^\star} D(\gamma) \tag{3.42}$$

$$W(\beta) = \langle\delta(a-\beta)\rangle = \mathsf{Tr}[\rho\delta(a-\beta)] \tag{3.43}$$

The latter equation can be reversed to recover the density matrix from a Wigner distribution:

$$\rho = \pi\int d^2\alpha W(\alpha)\delta(a-\alpha) \tag{3.44}$$

Number states

For the n-th eigenstate of the harmonic oscillator, $|n\rangle$, we can calculate the Wigner function analytically:

$$W_n(\alpha) = \frac{1}{\pi^2}\int d^2\gamma e^{\gamma^\star\beta-\gamma\beta^\star}\langle n|D(\gamma)|n\rangle = \frac{2}{\pi}(-1)^n e^{-2|\alpha|^2} L_n(4|\alpha|^2) \tag{3.45}$$

The $L_n(x)$ refer to Laguerre polynomials of order n. Since $L_n(0) = 1$, the Wigner function will be negative around the origin for all odd n. Furthermore, it can be shown that for all $n > 0$ the Wigner function will be negative at some point in phase space, meaning it cannot be interpreted as a classical probability function. This can be understood as number states encoding uniquely quantum systems with no analog in classical physics. As figure 3.1 demonstrates, increasing the photon number does not make the negativity of the Wigner function go away, meaning number states are still uniquely quantum even for very high values of n.

Coherent states

For a coherent state $|\alpha_0\rangle$, we obtain

$$W(\alpha) = \frac{2}{\pi}e^{-2|\alpha-\alpha_0|^2} \tag{3.46}$$

It is noticeable that the shape of the Wigner function corresponds to a Gaussian centered at α_0. This means that the distribution is positive everywhere. This tracks with the fact that coherent states closely resemble classical light: in this case, it is possible to interpret $W(\alpha)$ as a probability distribution in phase space.

It is interesting to note that, under the action of a harmonic oscillator potential, the Wigner function of a coherent state stays a Gaussian, with its center rotating around the origin, with the period of the rotation being given by the frequency of the harmonic oscillator [Case, 2008].

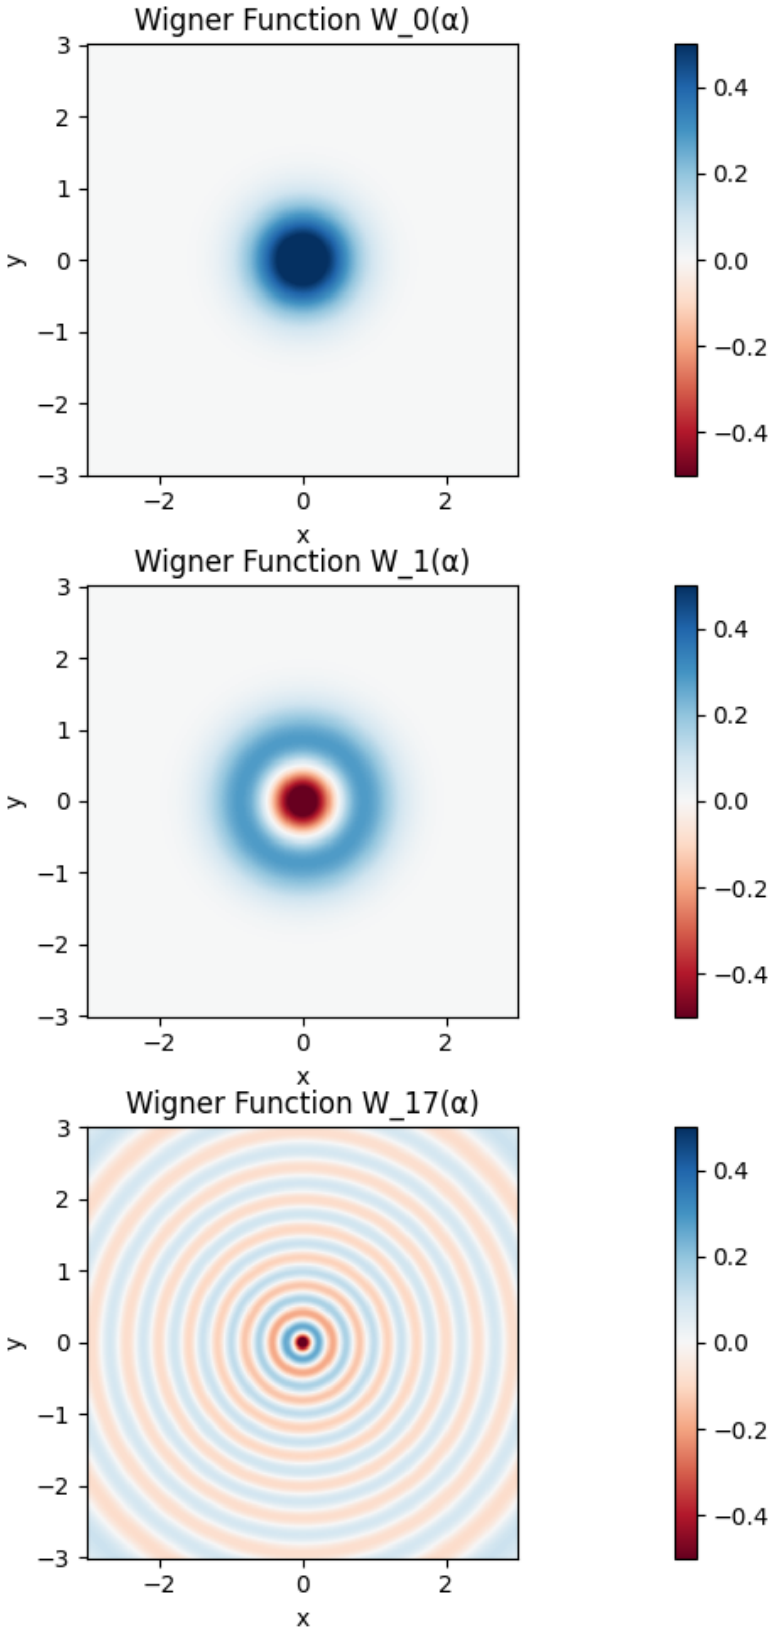

Figure 3.1 Plots of the Wigner function for number states $n = 0$, $n = 1$ and $n = 17$.

4 Problem setting

4.1 Problem statement

As alluded to in section 2.3, the biggest problem faced by quantum computing nowadays, in the quest for computers that can solve practical issues, is fighting error sources. These errors can make whole computations worthless. Even though error correction codes have been developed, fundamental physical error thresholds remain present, making precise handling of qubit operations necessary.

Furthermore, the most commonly used physical implementations of quantum computing, including superconducting circuits and trapped ions, rely on lasers interacting with qubits to implement single and two-qubit gates. Therefore, to keep errors down, it is a vital necessity to be able to exercise precise control of qubits with lasers.

The laser systems industry has made progress on developing packages that allow for high fidelity control and entanglement of qubit systems [Picken et al., 2023]. However, having a good understanding of the physical phenomena governing laser-qubit interaction remains key: better models are a crucial part of minimizing error, since they let us build more accurate descriptions of gates, therefore explaining error away.

We focus on a simplified problem modeling a photon that propagates in space and interacts with a two-level atom that we approximate as a point particle. To properly model space propagation, it is necessary to include many modes in the photon simulation, equivalent to simulating a system of M bosonic modes interacting with a two-level system. Since the eigenspectrum of a bosonic mode is, in theory, infinite, it is necessary to truncate it to a certain maximum amount of excitations that will be considered, N. The way the states fill is depicted in Figure 4.1.

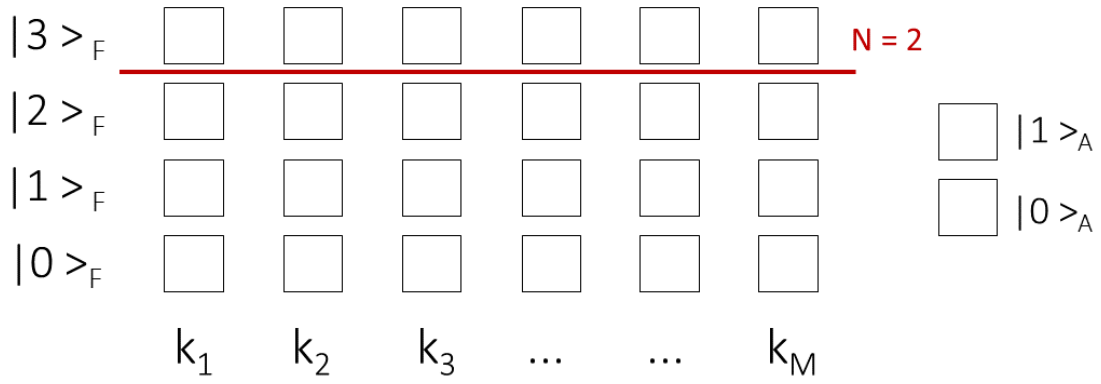

Figure 4.1 Diagram of the bosonic and atom states present in the problem we want to solve. M bosonic modes can be filled up to N excitations, and the two-level atom state, equivalent to a qubit, must also be considered.

4.2 Dimensionality problem

Setting aside the fact that the number of excitations is truncated, we can write a general state as

$$|\psi\rangle = |\psi_F\rangle \otimes |\psi_A\rangle = \sum_{n_1=0}^{\infty} \sum_{n_2=0}^{\infty} ... \sum_{n_M=0}^{\infty} \sum_{s\in\{g,e\}} |n_1\rangle \otimes |n_2\rangle \otimes ... \otimes |n_M\rangle \otimes |s\rangle \tag{4.1}$$

This is impossible to simulate on any computer, since infinite memory would be required to store the untruncated states. Instead, limiting each mode at N excitations, we instead have

$$|\psi\rangle = |\psi_F\rangle \otimes |\psi_A\rangle = \sum_{n_1=0}^{N} \sum_{n_2=0}^{N} ... \sum_{n_M=0}^{N} \sum_{s\in\{g,e\}} |n_1\rangle \otimes |n_2\rangle \otimes ... \otimes |n_M\rangle \otimes |s\rangle \tag{4.2}$$

From now on, for brevity, we will skip the outer products $\otimes$ when writing out states:

$$|\psi\rangle = |\psi_F; \psi_A\rangle = \sum_{n_1=0}^{N} \sum_{n_2=0}^{N} ... \sum_{n_M=0}^{N} \sum_{s\in\{g,e\}} |n_1, n_2, ..., n_M; s\rangle \tag{4.3}$$

Then, the dimensionality of the Hilbert space becomes

$$dim(\mathscr{H}) = 2(N+1)^M \tag{4.4}$$

Table 4.1 expresses the amount of memory needed to save the coefficients of a state like the one in our problem. It becomes evident that memory requirements scale very poorly for this problem, with requirements becoming too onerous even for supercomputers, even at numbers of bosons that are not too high.

N \ M	1	5	10	20	30
1	16 B	256 B	8 kB	8 MB	8 GB
2	24 B	1.9 kB	0.47 MB	27.8 GB	16 TB
5	40 B	25 kB	78 MB	762 TB	7.4×10^6 PB
10	80 B	800 kB	80 GB	8×10^5 PB	8×10^{15} PB

Table 4.1 Memory necessary to store the state of a simulation with M bosonic modes truncated to N excitations, assuming 4-byte floats are used for the coefficients. For reference, a laptop might have 16 GB of memory, a server might have a few TB in its large memory nodes (for example, 6 TB at LRZ), and the supercomputer with the most RAM in the world as of the time of writing (El Capitan at LLNL) has 5.44 PB across all nodes.

Simulating a Gaussian wave packet with acceptable resolution requires more than 50 modes, making the simulation we need unfeasible, even with powerful computers, if the excitation truncation threshold is greater than one.

Therefore, it is clear that it is necessary to find ways to optimize the simulation to be able to model systems of realistic size. To that end, different approaches have been considered.

4.3 Previous work

More than 20 years ago, groundbreaking work was done by [Havukainen et al., 1999] in simulating single-photon wave packets interacting with optical elements. By constraining the number of excitations in the

system, they were able to simulate free propagation and interaction with devices such as mirrors and beam splitters and recreate the two-slit experiment.

However, probably due to the constraints of the computers available at the time, they couldn't simulate more complex phenomena that might involve including more than one photon or recreating more complicated devices. Furthermore, numerical issues were found in their implementation of the solution of the Schrödinger equation via a Runge-Kutta scheme, which became apparent only in more complex scenarios.

For example, when a Mach-Zehnder interferometer is simulated, which involves two beam splitters, two mirrors, and a polarizer, nonphysical probabilities greater than one can be measured for specific detection probabilities. Therefore, a more careful numerical treatment was required.

The work of [Oba et al., 2024] addressed this issue. The authors found that the numerical problems of [Havukainen et al., 1999] could be addressed by solving Schrödinger's equation with a symplectic method based on a Suzuki-Trotter decomposition instead of Runge-Kutta schemes, assuming the spatial intensity of the photons remains low. As in the original paper, the dynamics of quantum optical devices were modeled as composed of many two-level atoms.

A Jaynes-Cummings Hamiltonian in the interaction picture was used. One notices that the Jaynes-Cummings Hamiltonian conserves the number of excitations in the system. Therefore, by initializing a photon in a number state $|1\rangle$, it becomes possible to simulate only the subspace of the combined Hilbert space $\mathscr{H}_F \otimes \mathscr{H}_a$ which has only one combined excitation across photon and atoms. We denote such a subspace as $\mathscr{H}_1$, and a general state in it can be written as

$$|\phi(t)\rangle = \sum_{\vec{k}} c(t,\vec{k})|1_{\vec{k}}\rangle + \sum_{j=1}^{N_A} c_j(t)|1_j\rangle \tag{4.5}$$

Building up on this, we introduce the notation $\mathscr{H}_n$ for the subspace of $\mathscr{H}_F \otimes \mathscr{H}_a$ containing exactly n excitations. Furthermore, since $\mathscr{H}_i \cap \mathscr{H}_j = \emptyset$ for all $i \neq j$, we can also consider the subspace containing states with excitations from 0 to n included and write it as $\mathscr{H}_{:n} = \bigcup_{i=0}^{n} \mathscr{H}_i$.

The state from 4.5 is then evolved with an unitary time evolution operator $|\phi(t)\rangle = U(t)|\phi(0)\rangle = e^{-iHt}|\phi(0)\rangle$. It can be proven that, if the Hamiltonian is of the Jaynes-Cummings type, $H = H_{JC}$, $|\phi(t)\rangle \in \mathscr{H}_1$ for all t.

To enforce this, the bosonic operators that are used to generate excited states have to be slightly modified. Namely,

$$a_j^\dagger a_k^\dagger|0\rangle = a_j^\dagger|1_k\rangle = 0 = a_k^\dagger|1_j\rangle = a_k^\dagger a_j^\dagger|0\rangle \tag{4.6}$$

In their work, they also consider degrees of polarization as states. They achieve this by writing the full state

$$|\phi_p(t)\rangle = \sum_{p\in\{H,V\}} \sum_{\vec{k}} c_p(t,\vec{k})|1_{\vec{k},p}\rangle + \sum_{j=1}^{N_A} c_{j,p}(t)|1_{j,p}\rangle \tag{4.7}$$

This allows them to characterize devices such as polarizers. In this latter case, the birefringent nature of the crystals that make them up is modeled by a coupling parameter $g_p(j,\vec{k})$ that depends on polarization as well. The Jaynes-Cummings type Hamiltonian becomes

$$H = \sum_p \Big(\sum_k^{N_M} \hbar\omega_k a_{k,p}^\dagger a_{k,p} + \sum_k^{N_M}\sum_j^{N_A} \hbar g_p(\sigma_j^+ a_{k,p} + \sigma_j^- a_{k,p}^\dagger)\Big) + \sum_j^{N_A} \hbar\Omega_j \sigma_j^+\sigma_j^- \tag{4.8}$$

A two-dimensional box of dimension (L_x, L_y) is considered, with periodic boundary conditions to model propagation. The box is then discretized with a (N_x, N_y) grid. Then, N_i, $i \in \{x, y\}$, modes are considered for each dimension: modes that are equally spaced in phase space are generated, from $-N_i\pi/L_i$ to $+N_i\pi/L_i$.

It is instructive to consider the size of the Hilbert space $\mathscr{H}_1$. Considering polarization, N_A two-level atoms and N_M photonic modes, $dim(\mathscr{H}_1) = 2(N_M + N_A)$.

This constitutes a clear improvement over the exponential dependence of dimensionality on the number of modes we observe for $\mathscr{H}_{:n}$. Photons can be simulated with many modes, allowing us to design pulses precisely. This fact presents immediate applications in qubit control, where pulses with precise parameters are required to implement single and multiple-qubit gates with high fidelity.

Furthermore, the dependence of the memory requirements being only linear in the number of atoms makes it feasible to simulate systems with a large number of atoms, opening the door to experiments with many quantum optical devices.

With this numerically stable method that avoids the curse of dimensionality, the authors can simulate, besides the previously mentioned Mach-Zehnder interferometer, phenomena such as scattering from an obstacle and Hong-Ou-Mandel interference, as well as demonstrate violations of the Bell-CHSH inequality.

Nevertheless, issues remained that I hope to address with my work. Namely, the fact remains that the Jaynes-Cummings Hamiltonian, which provides the basis for restricting the space of states being considered to $\mathscr{H}_1$, relies on approximations and assumptions, such as the rotating wave approximation, that lose their validity once coupling becomes relevant in the system. In quantum computers, working in a strong coupling regime is often desired, since doing so allows for creating entanglement within qubits mediated by interaction with lasers.

However, dismissing the rotating wave approximation means that states, even if initialized to $\mathscr{H}_1$, cannot be guaranteed to stay there after evolution. In this case, it becomes necessary to simulate a larger Hilbert space, which makes the curse of dimensionality present again.

In this work, we aim to find ways to solve these more complex problems while sidestepping the curse of dimensionality and maintaining a standard of accuracy, allowing us to gain a deeper understanding of the mechanisms of light-matter interaction at higher-coupling regimes.

5 Contributions

In our work, we have aimed to find ways to build accurate models of light-matter interaction that go beyond the Jaynes-Cummings model and can therefore account for phenomena taking place in higher-coupling regimes, where the curse of dimensionality would be reintroduced, as the total number of excitations cannot be said to be conserved anymore.

First of all, we consider the Quantum Rabi Hamiltonian now:

$$H_{QR} = \sum_{k}^{N_M} \hbar\omega_k a_k^\dagger a_k + \sum_{j}^{N_A} \hbar\Omega_j \sigma_j^+ \sigma_j^- + \hbar \sum_{k}^{N_M} \sum_{j}^{N_A} g(j,k)(\sigma_j^+ + \sigma_j^-)(a_k + a_k^\dagger) \tag{5.1}$$

The key difference this Hamiltonian presents when compared to the Jaynes-Cummings Hamiltonian is the presence of counter-rotating terms in the interaction term:

$$\sigma^+ a^\dagger, \qquad \sigma^- a \tag{5.2}$$

These imply that interactions where either two photons are created or two photons are destroyed are possible. The key consequence of this fact is that the total number of excitations in the system is no longer conserved: initializing a state in $\mathscr{H}_1$ does not guarantee that, after time evolution, the evolved state will still be in $\mathscr{H}_1$.

The main implication is that, since the evolved state does not in principle remain restricted to a particular subspace of $\mathscr{H}_F \otimes \mathscr{H}_A$, the whole space must be simulated to replicate the actual dynamics of the system. Nevertheless, as discussed in previous chapters, that is, in any case, unfeasible due to the infinite dimensionality of the Hilbert space. Moreover, even after truncating the number of excitations of each bosonic mode at N, the simulation remains challenging due to the curse of dimensionality, as simulating with multiple modes makes scaling exponential with respect to mode number.

5.1 Hilbert space modeling

We work with a photon with M bosonic modes interacting with one two-level atom.

From now on, we will assume that the Hilbert space containing photon-atom states, $\mathscr{H} = \mathscr{H}_F^M \otimes \mathscr{H}_A$, is made finite by truncating the M Fock spaces that describe bosonic states to N excitations. This step is necessary to make $\mathscr{H}$ finite-dimensional and therefore numerically treatable in the first place.

Then, a state in the basis of $\mathscr{H}$ can be written as a product of M bosonic states and one atomic state:

$$|\psi_{sample}\rangle = |n_1\rangle \otimes |n_2\rangle \otimes ... \otimes |n_M\rangle \otimes |s\rangle \equiv |n_1, n_2, ...n_M; s\rangle \tag{5.3}$$

Here, $|s\rangle \in \{g, e\}$ can be either of the two states of the two-level atom (ground or excited), and the n_i refer to the case where mode i has n photons occupying it.

The full state is therefore

$$|\psi\rangle = \sum_{n_1=0}^{N} ... \sum_{n_M=0}^{N} \sum_{s\in\{g,e\}} c_{n_1,...,n_M,s}|n_1, n_2, ...n_M; s\rangle \tag{5.4}$$

As we have discussed, the dimensionality of this space is $dim(\mathscr{H}) = 2(N+1)^M$. This means the curse of dimensionality is present in such a system, as exponential scaling in the number of modes is present.

The approach we have taken to deal with this scaling problem is to devise a way of further truncating the Hilbert space $\mathscr{H}$ such that the curse of dimensionality is alleviated. Still, it remains possible to extract information from the system, especially the qubit (that is, the atom), with high fidelity with respect to the untruncated state. It should be possible to calculate expectation values efficiently and with high accuracy as well.

We design a method based on the fact that changes in the bosonic part of the state are mediated by two-level atom absorption and emission. The only way for the photon state to change is for the atom to absorb an excitation and later re-emit it. Therefore, we hypothesize that, in such an interaction, the emitted photon will have the same energy, and, hence, the same wavenumber, as the photon that was absorbed.

Assuming we initialize the photon at a number state:

$$|\psi(0)\rangle = \sum_k c_{k,g}|1_k; g\rangle + c_{k,e}|1_k; e\rangle \tag{5.5}$$

We can use our hypothesis to guess that a state of the form $|1_k; s\rangle$ can evolve in one of the following ways:

$$|1_k; g\rangle \to |0; e\rangle \text{ mediated by the term } \sigma^+ a_k \text{ in } H_{QR}$$

$$|1_k; e\rangle \to |2_k; e\rangle \text{ mediated by the term } \sigma^- a_k^\dagger \text{ in } H_{QR}$$

$$|1_k; g\rangle \to |2_k; e\rangle \text{ mediated by the term } \sigma^+ a_k^\dagger \text{ in } H_{QR}$$

$$|1_k; g\rangle \to |0; g\rangle \text{ mediated by the term } \sigma^- a_k \text{ in } H_{QR}$$

Notice that at no point do we obtain states with excitations across multiple modes, that is, of the form

$$|1_k; 1_{k'}; s\rangle$$

Our approach involves building a new Hilbert space, which we will write as $\mathscr{H}'$, that excludes any such states with photon excitations spread across modes. A sample basis element can therefore be written as

$$|\psi_{samp}\rangle = |n_k; s\rangle \tag{5.6}$$

And the general expression for any state is

$$|\psi\rangle = \sum_k \sum_{n=0}^{N} \sum_{s\in\{g,e\}} c_{n,k,s}|n_k; s\rangle \tag{5.7}$$

This space only has N bosonic states per mode and atom excitation state. Therefore, its dimensionality is $dim(\mathscr{H}') = 2MN$, constituting a vast improvement over $\mathscr{H}$, since we lose the exponential dependence on M. The improvement makes simulations with large numbers of modes feasible, which opens the way for

complex pulse design. It also makes many-photon simulation significantly easier, as the cost depends only linearly on N (unlike in the untruncated state, where the dependence was $\sim N^M$, which, while polynomial, can still be unfeasible when M is large enough).

$N \setminus M$	1	5	10	20	30
1	8 B	40 B	80 B	160 B	240 B
2	16 B	80 B	160 B	320 B	480 B
5	40 B	200 B	400 B	800 B	1.2 kB
10	80 B	400 B	800 B	1.6 kB	2.4 kB

Table 5.1 Amount of computer memory necessary to store a state with M bosonic modes with N excitations per mode in $\mathscr{H}'$, the Hilbert space where states with excitations across different modes have been truncated away. For reference, compare with table 4.1 for the memory requirements in the untruncated space $\mathscr{H}$.

5.1.1 Ladder operators

It becomes necessary to build new ladder operators on the truncated basis. We have, as with the usual bosonic operators:

$$a_k|n_k; s\rangle = \sqrt{n}|(n-1)_k; s\rangle \tag{5.8}$$

$$a_k^\dagger|n_k; s\rangle = \sqrt{n+1}|(n+1)_k; s\rangle \tag{5.9}$$

Furthermore, this allows for the definition of a mode number operator, which counts the number of excitations in every mode:

$$N_k = a_k^\dagger a_k; \qquad N_k|n_k; s\rangle = n|n_k; s\rangle \tag{5.10}$$

So far, these truncated space bosonic operators for $\mathscr{H}'$ act like those for the untruncated space $\mathscr{H}$. However, due to the exclusion of states with excitations across different modes:

$$a_{k'}^\dagger|n_k; s\rangle = 0 \quad \text{for } k \neq k', n > 0 \tag{5.11}$$

$$a_{k'}|n_k; s\rangle = 0 \quad \text{for } k \neq k', \quad \forall n \tag{5.12}$$

These can be rewritten more succinctly as

$$a_k^\dagger a_{k'}^\dagger = 0 \quad \text{for } k \neq k' \tag{5.13}$$

and

$$a_k a_{k'}^\dagger = 0 \quad \text{for } k \neq k' \tag{5.14}$$

To implement bosonic states in the computer, it is helpful to rewrite them directly as functions of the basis states. Using equations 5.8 and 5.9, each of the annihilation and creation operators can be written as:

$$a_k = \sum_{n=1}^{N} \sum_{s} \sqrt{n} |(n-1)_k; s\rangle\langle n_k; s| = \sum_{s} |0_k; s\rangle\langle 1_k; s| + \sqrt{2}|2_k; s\rangle\langle 1_k; s| + \ ... \ \sqrt{N}|(N-1)_k; s\rangle\langle N_k; s| \tag{5.15}$$

$$a_k^\dagger = \sum_{n=0}^{N-1} \sum_{s} \sqrt{n+1} |(n+1)_k; s\rangle\langle n_k; s| = \sum_{s} |1_k; s\rangle\langle 0_k; s| + \sqrt{2}|2_k; s\rangle\langle 1_k; s| + \ ... \ \sqrt{N}|N_k; s\rangle\langle (N-1)_k; s| \tag{5.16}$$

Collective bosonic operators for the whole state can be defined from these as well:

$$a = \sum_k c_k a_k \tag{5.17}$$

$$a^\dagger = \sum_k c_k^\star a_k^\dagger \tag{5.18}$$

Here, the constants c_k are in principle free with the only constraint being normalization:

$$\sum_k |c_k|^2 = 1 \tag{5.19}$$

However, in practice, we will pick c_k to recreate a Gaussian pulse:

$$c_{\vec{k}} = \frac{2\sigma\sqrt{\pi}}{L} \exp\Big(-\frac{\sigma^2}{2}(\vec{k} - \vec{k}_0)^2 - i\vec{k}\cdot\vec{r} \Big) \tag{5.20}$$

It is interesting to discuss the dimensionality of the matrices representing the operators a, a_k, $a^\dagger$, $a_k^\dagger$. In principle, all of these operators are linear maps:

$$a_{(k)}^{(\dagger)} : \mathscr{H}' \to \mathscr{H}' \tag{5.21}$$

And therefore, have dimension $dim(\mathscr{H}')^2 = 4N^2M^2$. However, suppose we focus on the number of non-zero elements. In that case, we see that the $a_k^{(\dagger)}$ only have $N-1$ elements that are distinct from zero, and, similarly, the $a^{(\dagger)}$ have $M(N-1)$ non-zero elements. This fact means the operators lend themselves to being modeled with sparse matrices, since that allows us to make significant savings on memory of one order of magnitude in N and one or two orders of magnitude in M.

5.1.2 Coherent states

Recall that a coherent state is defined as the eigenstate of the annihilation operator. For a multimode state, that means the collective operator acting on all modes (**??**):

$$a|\alpha\rangle = \alpha|\alpha\rangle; \qquad a = \sum_{\vec{k}} c_{\vec{k}} a_{\vec{k}} \tag{5.22}$$

A naive approach to finding the eigenstate $|\alpha\rangle$ would be to diagonalize a. However, there are two issues with that approach:

- The definition of a coherent state refers to the case where a is understood as an operator on the infinite-dimensional Hilbert space of all bosonic states. However, truncation discretizes the operator, making its spectrum of eigenvalues go from continuous and uncountably infinite to discrete and finite. Then, it becomes impossible to find eigenstates $|\alpha\rangle$ for every complex number α, which makes finding phase state descriptions of time evolutions challenging.
- Diagonalizing a matrix is computationally expensive, as a singular value decomposition is required, with $O(n^3)$ cost.

The displacement operator again offers an alternative formulation for coherent states. For the collective mode operator:

$$D(\alpha) = e^{\alpha a^\dagger - \alpha^\star a} \tag{5.23}$$

$$|\alpha\rangle = D(\alpha)|0\rangle \tag{5.24}$$

We inspect how this collective mode operator acts on the vacuum state:

$$D(\alpha)|0\rangle = e^{\alpha a^\dagger - \alpha^\star a}|0\rangle = \sum_n \frac{1}{n!}(\alpha a^\dagger - \alpha^\star a)^n|0\rangle = \sum_n \frac{1}{n!}\sum_m \binom{n}{m}(\alpha a^\dagger)^m(\alpha^\star a)^{n-m}|0\rangle =$$

$$= \sum_n \frac{1}{n!}\sum_m \binom{n}{m}\alpha^m(\alpha^\star)^{n-m}(a^\dagger)^m a^{n-m}|0\rangle = \sum_n \frac{1}{n!}\sum_m \binom{n}{m}\alpha^m(\alpha^\star)^{n-m}(\sum_k c_k^\star a_k^\dagger)^m(\sum_k c_k a_k)^{n-m}|0\rangle =$$

$$= \sum_n \frac{1}{n!}\sum_m \binom{n}{m}\alpha^m(\alpha^\star)^{n-m}$$

$$\Big(\sum_{m_1,...m_M\geq 0}\frac{m!}{m_1!...m_M!}(c_1^\star)^{n_1}...(c_M^\star)^{n_M}(a_1^\dagger)^{n_1}...(a_M^\dagger)^{n_M}\Big)\Big(\sum_{m_1',...m_M'\geq 0}\frac{(n-m)!}{m_1'!...m_M'!}c_1^{m_1'}...c_M^{m_M'}a_1^{m_1'}...a_M^{m_M'}\Big)|0\rangle =$$

$$= \sum_{n_1}...\sum_{n_M}\gamma_{n_1,...,n_M}|n_1,...,n_M\rangle$$

This means that, by initializing a state as a coherent state, the probability at time zero of finding excitations spread across modes is non-zero, so we should expect less accuracy for the truncation method compared to using it for number photonic states.

We investigate different ways of modeling coherence to preserve accuracy while alleviating the curse of dimensionality.

Independent mode initialization

We aim to make the Hilbert space truncation from $\mathscr{H}$ to $\mathscr{H}'$ compatible with working with coherent states. A straightforward way is to neglect the components of the displacement operator that create the excitations across modes. That is, rewrite the displacement operator as

$$D(\alpha) \approx D_1(\alpha)\otimes D_2(\alpha)\otimes ...\otimes D_M(\alpha) \tag{5.25}$$

Where the $D_k(\alpha)$ are the displacement operators for the bosonic operators on mode k

$$D_k(\alpha) = e^{\alpha a_k^\dagger - \alpha^\star a_k} \tag{5.26}$$

The state can then be equivalently written as

$$\frac{1}{C}D_1(\alpha)\otimes D_2(\alpha)\otimes ...\otimes D_M(\alpha)|0\rangle = |\alpha_1,\alpha_2,...,\alpha_M\rangle \equiv |\alpha\rangle \tag{5.27}$$

Where C is chosen so that the state is normalized. Since $D(\alpha)^\dagger = D(-\alpha)$

$$C = \sqrt{\langle 0|D_1(-\alpha)\otimes D_2(-\alpha)\otimes ...\otimes D_M(-\alpha)D_1(\alpha)\otimes D_2(\alpha)\otimes ...\otimes D_M(\alpha)|0\rangle} \tag{5.28}$$

This method has the advantage of making initialization straightforward and fast. However, a significant part of the information is lost in the part of the Hilbert space with excitations across modes, which is truncated away.

Changing the basis of the Hamiltonian

This approach is based on noticing that for the temporal evolution of a coherent state:

$$|\phi(t)\rangle = U(t)|\psi(0)\rangle = e^{iHt}|\phi(0)\rangle = e^{iHt}|\alpha;s\rangle \tag{5.29}$$

By making use of the unitary nature of the displacement operator ($D(\alpha)^\dagger D(\alpha)$), and writing $D(\alpha)\otimes$ for brevity:

$$e^{iHt}|\alpha;s\rangle = D(\alpha)D(\alpha)^\dagger e^{iHt}D(\alpha)|0;s\rangle \tag{5.30}$$

$$D(\alpha)^\dagger e^{iHt}D(\alpha) = D(\alpha)^\dagger \sum_n \frac{1}{n!}(iHt)^n D(\alpha) = D(\alpha)^\dagger \sum_n \frac{1}{n!} i^n D(\alpha)^\dagger H^n D(\alpha) t^n \tag{5.31}$$

Defining $\tilde{H} = D(\alpha)^\dagger H D(\alpha)$:

$$\tilde{H}^n = D(\alpha)^\dagger H D(\alpha) D(\alpha)^\dagger H ... H D(\alpha) D(\alpha)^\dagger H D(\alpha) = D(\alpha)^\dagger H^n D(\alpha) \tag{5.32}$$

And, therefore:

$$e^{iHt}|\alpha;s\rangle = D(\alpha)e^{i\tilde{H}t}|0;s\rangle \tag{5.33}$$

This allows us to simulate the evolution of a coherent state as the evolution of a vacuum state in a different basis, according to the Hamiltonian $\tilde{H} = D(\alpha)^\dagger H D(\alpha)$.

However, the problem of computing $\tilde{H}$ remains. A numerical calculation of $\tilde{H}$ by multiplying H by the matrices of the discretized displacement operator and its conjugate would be computationally expensive, and it would also run into the same issues regarding loss of information as independent mode initialization, if we truncate the Hilbert space to alleviate the curse of dimensionality.

We analytically derive the Hamiltonian $\tilde{H}$ from the Quantum Rabi Hamiltonian to sidestep this. First, we observe that, using the Baker-Campbell-Hausdorff formula [Rossmann, 2006], we can derive the following identities:

$$D(\alpha)^\dagger a_k^\dagger D(\alpha) = a_k^\dagger + \alpha^\star c_k \tag{5.34}$$

$$D(\alpha)^\dagger a_k D(\alpha) = a_k - \alpha c_k^\star \tag{5.35}$$

We can then see $D(\alpha)^\dagger H D(\alpha)$ as a transformation of $H = H_0 + H_I$:

$$H_0 = \sum_k \hbar\omega a_k^\dagger a_k + \sum_j \hbar\Omega_j \sigma_j^+ \sigma_j^- \rightarrow \sum_k \hbar\omega(a_k^\dagger + \alpha^\star c_k)(a_k - \alpha c_k^\star) + \sum_j \hbar\Omega_j \sigma_j^+ \sigma_j^- = \tag{5.36}$$

$$= H_0 + \sum_k \hbar\omega_k(\alpha^\star c_k a_k - \alpha c_k^\star a_k^\dagger) = \tilde{H}_0$$

$$H_I = \hbar \sum_{j,k} g_{jk}(\sigma_j^+ a_k e^{ikr_j} + \sigma_j^- a_k^\dagger e^{-ikr_j} + \sigma_j^- a_k e^{ikr_j} - \sigma_j^+ a_k^\dagger e^{ikr_j}) \rightarrow \tag{5.37}$$

$$\rightarrow \hbar \sum_{j,k} g_{jk}(\sigma_j^+(a_k - \alpha c_k^\star)e^{ikr_j} + \sigma_j^-(a_k^\dagger + \alpha^\star c_k)e^{-ikr_j} + \sigma_j^-(a_k - \alpha c_k^\star)e^{ikr_j} - \sigma_j^+(a_k^\dagger + \alpha^\star c_k)e^{ikr_j}) =$$

$$= H_I + \sum_{j,k}(\sigma_j^+ + \sigma_j^-)(\alpha c_k^\star e^{ikr_j} + \alpha^\star c_k e^{-ikr_j}) = \tilde{H}_I$$

We can then simulate temporal evolution with the operator $e^{i\tilde{H}t}$, $\tilde{H} = \tilde{H}_0 + \tilde{H}_I$ initializing the boson in the atom state. This Hamiltonian is non-conserving as well; unlike the original quantum Rabi Hamiltonian, though, it has terms that induce excitations in each of the subsystems (photon and atom) separately.

This introduces a stronger foundation for performing simulations of coherent states on truncated Hilbert spaces, as no information is discarded at initialization time.

Simulating in a coherent state basis

The previously introduced methods are based on finding a way to conveniently simulate the evolution of a coherent state in a number basis in an efficient manner, avoiding the curse of dimensionality. However, there remains the option of completely sidestepping that process and choosing to simulate the coherent states in a state tailored for them, built out of a different approach than using the Fock space.

Given that we want to simulate the evolution of a state initialized as a coherent state, using other coherent states as members of the basis of a Hilbert space seems like a natural choice. That is, the goal is to find some way of expressing a generic state as a linear combination of coherent states:

$$|\psi\rangle = \sum_i c_i |\alpha_i\rangle \tag{5.38}$$

An issue one encounters, however, is that coherent states with different α are not orthogonal to each other. Indeed,

$$\langle\beta|\alpha\rangle = \langle 0|D(\beta)D(\alpha)|0\rangle = e^{\frac{1}{2}(|\alpha|^2+|\beta|^2)-\beta^\star\alpha} \sim e^{-\frac{1}{2}|\beta-\alpha|^2}$$

Orthogonality (which can be easily extended to orthonormality) is a key property for bases of Hilbert spaces on which we wish to perform simulations. Indeed, let us consider two possible bases for a Hilbert space:

$$\mathscr{H}_A \quad : B_A = \{|u_i\rangle\}, \quad i = 1, .., n$$

$$\mathscr{H}_B \quad : B_B = \{|v_i\rangle\}, \quad i = 1, .., n$$

If we attempt to calculate the expectation value of an arbitrary observable M in each basis:

$$\langle M\rangle_A = tr[\rho M]_A = \sum_i \langle v_i|\rho M|u_i\rangle$$

$$\langle M\rangle_B = tr[\rho M]_B = \sum_i \langle w_i|\rho M|v_i\rangle$$

If B_A and B_B are orthogonal, then we can find a third basis

$$B_C = \{|w_i\rangle\}, \quad i = 1, .., n$$

such that U_A, U_B exist and are unitary matrices fulfilling

$$|u_i\rangle = U_A|w_i\rangle$$

$$|v_i\rangle = U_B|w_i\rangle$$

And, therefore,

$$tr[\rho M]_A = tr[U_A^\dagger \rho M U_A]_C$$

$$tr[\rho M]_B = tr[U_B^\dagger \rho M U_B]_C$$

Using the cyclical property of the trace, $tr[U^\dagger \rho M U] = tr[\rho M]$. Therefore, for B_A and B_B orthogonal,

$$\langle M\rangle_A = tr[\rho M]_A = tr[\rho M]_B = \langle M\rangle_B \tag{5.39}$$

However, for non-orthogonal bases, this fact is not true, and expectation values therefore depend on the specific choice of basis used to evaluate the expectation value, making non-orthogonal bases ill-suited for simulation, as it is not possible to extract valuable information from simulations performed using them.

This presents a significant challenge to simulation in coherent state bases, since, for two states in the basis, $\langle\alpha_i|\alpha_j\rangle = e^{-\frac{1}{2}|\alpha_i-\alpha_j|^2} \neq 0$.

Nevertheless, the exponential decay in overlap between coherent states as a function of their distance in phase space can be leveraged to build "almost orthogonal" bases. In particular:

$$|\alpha_i - \alpha_j| = 1 \implies e^{-\frac{1}{2}|\alpha_i-\alpha_j|^2} = 0.60653$$

$$|\alpha_i - \alpha_j| = 2 \implies e^{-\frac{1}{2}|\alpha_i-\alpha_j|^2} = 0.13534$$

$$|\alpha_i - \alpha_j| = 3 \implies e^{-\frac{1}{2}|\alpha_i-\alpha_j|^2} = 0.01111$$

$$|\alpha_i - \alpha_j| = 4 \implies e^{-\frac{1}{2}|\alpha_i-\alpha_j|^2} = 0.00003$$

Therefore, to achieve an "orthogonality fidelity" of at least four zeros, coherent states should be sampled for the basis in such a way that the distance between them in phase space is at least 4. This constitutes a significant roadblock in the way to coherent state space simulation: usually, when building bases to describe spaces, basis elements are picked as close to each other as computational resources allow to increase accuracy and representativity. This is challenging to achieve for coherent state bases, as methods to maintain representativity while maintaining the orthogonality of the basis must be devised.

There are no standard methods to construct coherent state bases. Research has been done on the simulation of Fock states in coherent state bases [Marshall and Anand, 2023], where it was found that, for an arbitrary state in the N-excitational Fock basis $\mathscr{H}_n$:

$$|\psi\rangle = \sum_{n=0}^{N} a_n |n\rangle \tag{5.40}$$

A construction in terms of $N+1$ coherent states can be built that approximates $|\psi\rangle$:

$$|\tilde{\psi}\rangle = \frac{1}{\sqrt{C}} \sum_{k=0}^{N} c_k |\epsilon e^{2\pi i k/(N+1)}\rangle \tag{5.41}$$

$$c_k = \frac{e^{\epsilon^2/2}}{N+1} \sum_{n=0}^{N} \sqrt{n!} \frac{a_n}{\epsilon^n} e^{-2\pi i n k/(N+1)} \tag{5.42}$$

Here, C is a normalization parameter, and $\epsilon \in \mathbb{R}$ is a free parameter.
A function of ϵ can bound the fidelity of this reconstruction:

$$|\langle\psi|\tilde{\psi}\rangle|^2 = \frac{1}{C} = \frac{1}{1 + O(\epsilon^{2(N+1)}/(N+1)!)} \tag{5.43}$$

Implying that, by choosing an appropriate ϵ, the error can be made arbitrarily small.
Furthermore, the Fock states $|N\rangle$ can be approximated by

$$|\tilde{N}\rangle = \frac{1}{C} \frac{\sqrt{N!}}{N+1} \frac{e^{\epsilon^2/2}}{\epsilon^N} \sum_{k=0}^{N} e^{-2\pi i k N/(N+1)} |\epsilon e^{2\pi i k/(N+1)}\rangle \tag{5.44}$$

$$C = N! \sum_{k=0}^{\infty} \frac{e^{2k(N+1)}}{(k(N+1)+N)!}$$

Nevertheless, such an approach does not fit well with our goal of simulating coherent states, or states that we at least expect to be close to coherent states, since they evolved from coherent states. Mainly, this is due to the interaction of the error bound and the basis orthogonality requirement: reconstruction fidelity drops when ϵ is made small, but orthogonality imposes a requirement to make ϵ large.

We propose an alternative scheme based on a grid of coherent states. The grid should fulfill two basic requirements:

- The basis states should be at a distance in phase space, $|\alpha_i - \alpha_j|$, that is sufficient so that an assertion that they are almost orthogonal is justified.
- Any state $|\alpha\rangle$ within the grid should be able to be written as a linear combination of some of the grid states. Therefore, there should be at least some $|\alpha_i\rangle$ for which $|\langle\alpha_i|\alpha\rangle|^2 \gg 0$.

The most natural ways to construct such a grid are to build a square grid or a hexagonal grid: both present advantages and disadvantages. On the one hand, a square grid is more straightforward to construct. However, hexagonal grids present advantages that make them a better choice for coherent state simulation, namely the fact that they are the most compact 2D packing of spheres, thus minimizing the distance from a grid point to the centers of the cell. For a distance between grid points of l in both types of grids, the grid point to cell center distance x is:

$$x_{\text{square}} = l/\sqrt{2} \approx 0.707l$$

$$x_{\text{hexagonal}} = l/\sqrt{3} \approx 0.577l$$

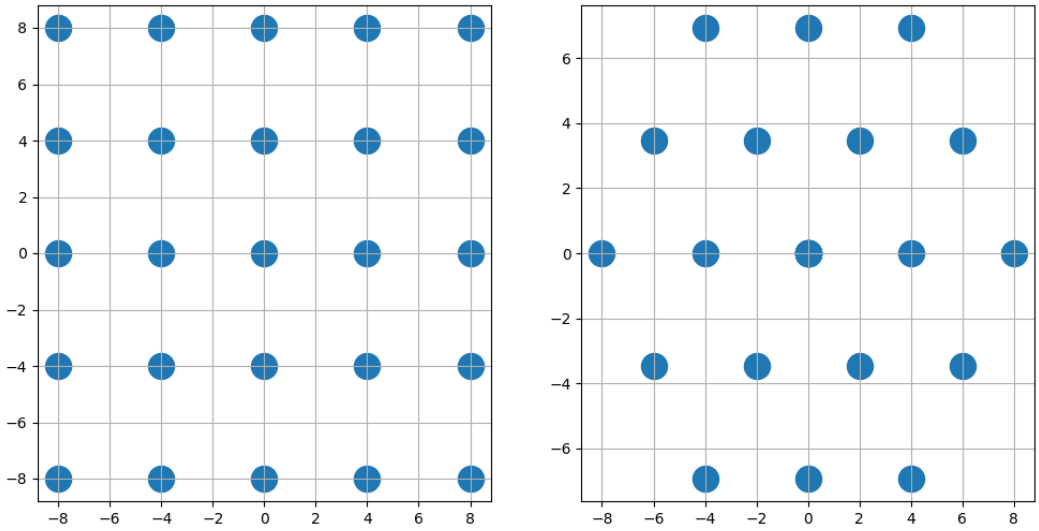

Figure 5.1 Square grid (left) and hexagonal grid (right) of coherent states in phase space. Spacing between states is kept constant at $|\alpha_i - \alpha_j| = 4$.

This is a crucial aspect to leverage in building bases for coherent state simulation since that distance is the furthest an arbitrary coherent state can be from the basis states.

The problem of dealing with many-mode states remains. We do that by rewriting the displacement operator:

$$D(\alpha) = \exp\Big[\alpha a^\dagger - \alpha^\star a\Big] = \exp\Big[\alpha \sum_k c_k^\star a_k^\dagger - \alpha^\star \sum_k c_k a_k\Big] = \exp\Big[\sum_k \Big(\alpha c_k^\star a_k^\dagger - \alpha^\star c_k a_k\Big)\Big] = \tag{5.45}$$

$$= \prod_k \exp\Big[\alpha c_k^\star a_k^\dagger - \alpha^\star c_k a_k\Big]$$

Care must be taken to notice that this is only true for untruncated Hilbert spaces, as otherwise $a_i a_j \neq a_j a_i$ and the sum in the exponential does not factorize into a product, as commutator terms would need to be accounted for.

We inspect how this operator acts on the vacuum state. Noticing that $\exp\Big[\alpha c_k^\star a_k^\dagger - \alpha^\star c_k a_k\Big]$ is equivalent to $\mathbb{I} \otimes ... \otimes D_k(\alpha c_k^\star) \otimes ... \otimes \mathbb{I}$:

$$D(\alpha)|0\rangle = \prod_k D_k(\alpha c_k^\star)|0\rangle = \bigotimes_k |\alpha c_k^\star\rangle = |(ac_1^\star)_1, (ac_2^\star)_2, ..., (ac_M^\star)_M\rangle \tag{5.46}$$

This expression can be used to initialize a computational state as

$$|\psi(0)\rangle = \frac{1}{\sqrt{C}} \sum_{i_1,...,i_M=0}^{N} \langle \alpha_{i_1}, ..., \alpha_{i_M} | (ac_1^\star)_1, (ac_2^\star)_2, ..., (ac_M^\star)_M\rangle |\alpha_{i_1}, ..., \alpha_{i_M}\rangle \tag{5.47}$$

The components of the Hamiltonian matrix can be calculated by exploiting orthogonality. Recalling the definition of the bosonic operators:

$$a|\alpha\rangle = \alpha|\alpha\rangle$$

It is easy to see that:

$$\langle\alpha|a^\dagger = \alpha^\star\langle\alpha| \tag{5.48}$$

And, therefore,

$$\langle\alpha_j|a^\dagger a|\alpha_i\rangle = \alpha_j^\star\alpha_i\delta_{ij} \tag{5.49}$$

$$\langle\alpha_j|a|\alpha_i\rangle = \alpha_i\delta_{ij} \tag{5.50}$$

$$\langle\alpha_j|a^\dagger|\alpha_i\rangle = \alpha_j^\star\delta_{ij} \tag{5.51}$$

Using these, the combinations of operators in the Hamiltonian can be rewritten in the coherent state basis:

$$a_k^\dagger a_k = \sum_i |\alpha_i|^2(|\alpha_{k,i};g\rangle\langle\alpha_{k,i};g| + |\alpha_{k,i};e\rangle\langle\alpha_{k,i};e|) \tag{5.52}$$

$$a_k\sigma^+ = \sum_i \alpha_i|\alpha_{k,i};e\rangle\langle\alpha_{k,i};g| \tag{5.53}$$

$$a_k^\dagger\sigma^- = \sum_i \alpha_i^\star|\alpha_{k,i};g\rangle\langle\alpha_{k,i};e| \tag{5.54}$$

$$a_k\sigma^- = \sum_i \alpha_i|\alpha_{k,i};g\rangle\langle\alpha_{k,i};e| \tag{5.55}$$

$$a_k^\dagger\sigma^+ = \sum_i \alpha_i^\star|\alpha_{k,i};e\rangle\langle\alpha_{k,i};e| \tag{5.56}$$

If we denote the space of M bosonic modes with N coherent basis states per mode as $\mathscr{H}_{coh}$, it follows that

$$\text{dim}(\mathscr{H}_{coh}) = N^M \tag{5.57}$$

This scaling makes the curse of dimensionality apparent again, as memory requirements scale exponentially with the number of modes.

To alleviate this, we suggest a scheme that instead initializes the computational state as follows:

$$|\tilde{\psi}\rangle = \frac{1}{\sqrt{C}}\sum_k\sum_i\langle\alpha_i|\alpha c_k\rangle|\alpha_{i,k}\rangle \tag{5.58}$$

The dimensionality of this truncated space $\mathscr{H}'_{coh}$ is

$$\text{dim}(\mathscr{H}'_{coh}) = NM + 1 \tag{5.59}$$

Where the factor of plus one comes from the fact that the vacuum state must be treated separately, as $|0_k\rangle$ is in fact the same state for every k.

Truncating the coherent state space lets us make the memory requirements of the problem linear in both the number of modes and the number of coherent states included per bosonic mode, therefore sidestepping the curse of dimensionality.

5.2 Calculation of expectation values and other figures of merit

We often wish to use the evolved states we have computed to obtain expectation values and other figures, such as entanglement, that help characterize the systems we work with.

5.2.1 Atom excitation probability

Since we want to apply our work to qubit control, it is vital to understand how qubits excite when subjected to interaction with photons. Given that our system is probabilistic, we can only have access to the probability of, upon measuring the atom, observing whether it is in the ground or in the excited state.

This is equivalent to measuring the expectation value of the operator:

$$A = \mathbf{1} \otimes |e\rangle\langle e| \tag{5.60}$$

Then, the probability of, at time t, finding the atom in the excited state is:

$$p_e(t) = \langle\psi(t)|A|\psi(t)\rangle \tag{5.61}$$

As expected, for $t = 0$, since we initialize the state as a product of pure states:

$$|\psi(0)\rangle = |\psi_\gamma\rangle \otimes |\psi_a\rangle = |\psi_\gamma\rangle(c_g|g\rangle + c_e|e\rangle) \tag{5.62}$$

$$p_e(0) = \langle\psi_\gamma|\langle\psi_a|(\mathbf{1} \otimes |e\rangle\langle e|)|\psi_\gamma\rangle|\psi_a\rangle = \langle\psi_\gamma|\psi_\gamma\rangle\langle\psi_a|e\rangle\langle e|\psi_a\rangle = |c_e|^2 \tag{5.63}$$

However, this probability does not fully characterize the atom's state. More information is required to model properties such as entanglement.

5.2.2 Reduced density matrix

To better understand the behavior of the photon and atom subsystems, it is helpful to extract their reduced density states first. We know that, from the fact that purity is a constant of motion, the whole system state will never become mixed. However, the same cannot be said of the reduced states. Therefore, instead of working with a vector as with the composite system state, we have to resort to the density matrix formalism for the reduced states.

The reduced density states for the photon, ρ_γ, and for the atom, ρ_a, are defined as follows:

$$\rho_\gamma = \mathsf{Tr}_a(\rho) = \sum_s \langle s|_a \, \rho |s\rangle_a \tag{5.64}$$

$$\rho_a = \mathsf{Tr}_\gamma(\rho) = \sum_s \langle s|_\gamma \, \rho |s\rangle_\gamma \tag{5.65}$$

The sum $\sum_s$ refers to summing over all states in the basis of the Hilbert space of the atom and the photon, respectively.

Formally, reduced states are computed from the density matrix of the whole system, which can be obtained directly from the state vector at any given time. Expressed generally, with the sum over s referring to either a number or a coherent state basis comprised of M modes, N basis elements per mode, and a ground or excited state in the atom:

$$|\Psi(t)\rangle = \sum_s c_s|s\rangle \implies \rho(t) = \sum_s \sum_{s'} c^\star_{s'} c_s |s\rangle\langle s'| \tag{5.66}$$

For the photon reduced state, this is simply equivalent to:

$$\rho_\gamma = \mathsf{Tr}_a(\rho) = \langle g|\rho|g\rangle + \langle e|\rho|e\rangle \tag{5.67}$$

Meanwhile, for the atom reduced states, the sum depends on the exact basis we used to describe the bosonic modes. For number states in the truncated Hilbert space, this is equivalent to

$$\rho_a = \mathsf{Tr}_\gamma(\rho) = \sum_{n=1}^{N} \sum_k \langle n_k|\rho|n_k\rangle \tag{5.68}$$

Or, in the analogous coherent state basis:

$$\rho_a = \mathsf{Tr}_\gamma(\rho) = \sum_{i=1}^{N} \sum_k \langle \alpha_{i,k}|\rho|\alpha_{i,k}\rangle \tag{5.69}$$

In practice, though, we want to avoid computing the density matrix for the whole system, as it requires $4N^2M^2$ operations, making it impractical.

Instead, we notice that, for the reduced density matrix of the atom, we only need $4NM$ operations if we calculate it from the whole-system state vector as

$$\rho_{gg} = \sum_k \sum_{s\in\text{basis}} c^\star_{ksg} c_{ksg} \tag{5.70}$$

$$\rho_{ge} = \sum_k \sum_{s\in\text{basis}} c^\star_{kse} c_{ksg} \tag{5.71}$$

$$\rho_{eg} = \sum_k \sum_{s\in\text{basis}} c^\star_{ksg} c_{kse} \tag{5.72}$$

$$\rho_{ee} = \sum_k \sum_{s\in\text{basis}} c^\star_{kse} c_{kse} \tag{5.73}$$

Nevertheless, obtaining the reduced density matrix for the photon is much more costly. Even though there are only two calculations per element, there are N^2M^2 elements to compute, so the total cost is $2N^2M^2$.

5.2.3 Occupation numbers

As mentioned before, calculating the excitation probability for the atom is straightforward, even without access to the reduced density matrix of the atom. Indeed, the probability of measuring the excited state can be identified with one of the elements of the reduced density matrix:

$$p_e(t) = (\rho_a)_{ee} \tag{5.74}$$

It is nevertheless interesting to comment on measurements in the atomic subsystem. Recall that, in section 5.1.2, we reviewed a method to compute time evolution with a coherent state Hamiltonian basis, based on rewriting the evolved state as:

$$|\Psi(t)\rangle = D(\alpha)e^{i\tilde{H}t}|0;s\rangle \tag{5.75}$$

This means that, instead of evolving a coherent state, a vacuum state evolves using a Hamiltonian that depends on the initial state. After that, a displacement operator, equivalent to a matrix-vector product in discretized space, must be computed.

However, let us consider a case in which we are only interested in measurements on the atom. Such measurement operators are of the form:

$$A = \sum_{i,j\in\{g,e\}} \mathbf{1}_\gamma \otimes c_{ij}|s_i\rangle\langle s_j| \tag{5.76}$$

The expectation value of such an operator for the evolved state is thus:

$$\langle A\rangle = \langle 0;s|e^{i\tilde{H}t}D(\alpha)^\dagger\Big(\sum_{i,j\in\{g,e\}} \mathbf{1}_\gamma \otimes c_{ij}|s_i\rangle\langle s_j|\Big)D(\alpha)e^{i\tilde{H}t}|0;s\rangle = \tag{5.77}$$

$$\sum_{i,j\in\{g,e\}} \langle 0;s|e^{i\tilde{H}t}\mathbf{1}_\gamma \otimes c_{ij}|s_i\rangle\langle s_j|e^{i\tilde{H}t}|0;s\rangle$$

Since $D(\alpha)$ is shorthand for $D(\alpha)\otimes\mathbf{1}_a$, it turns out the matrix-vector multiplication can be avoided.

An analogous method can be used to calculate the occupation of each photonic mode. However, it must be considered that the average occupation number is not a probability since there are multiple excitation levels.

There are two separate ways to understand occupation. It can refer either to the number of photons in each mode:

$$N_k = a_k^\dagger a_k \tag{5.78}$$

Or to the total number of photons present in the system:

$$N = a^\dagger a = \Big(\sum_k c_k^\star a_k^\dagger\Big)\Big(\sum_{k'} c_{k'}a_{k'}\Big) = \sum_{k,k'} c_k^\star c_{k'} a_k^\dagger a_{k'} \tag{5.79}$$

Notice that it is not true that $\sum_k N_k = N$.

The expected number of photons per mode can be computed from the state vector as

$$\langle N_k\rangle = \sum_n \langle n_k;g|\rho|n_k;g\rangle + \langle n_k;e|\rho|n_k;e\rangle = \sum_n |c_{nkg}|^2 + |c_{nke}|^2 \tag{5.80}$$

An equivalent formula is also valid for a coherent state basis, with the sum being over basis states instead of over excitations.

By calculating $\langle N_k\rangle(t)$, animations can be created that illustrate how the photon propagates in the box.

5.2.4 Entanglement

Entanglement illustrates to what degree the state of the photon and the state of the atom can be understood as independent. It is related to the size of the interaction, or the effect the photon has on the atom and vice versa.

We quantify entanglement using entanglement entropy. Entanglement entropy leverages the fact that, given a state ρ and its reduced states $\rho_A = Tr_B[\rho]$, $\rho_B = Tr_A[\rho]$ on two subsystems A, B, the following holds:

$$S(\rho_A) = -Tr[\rho_A \log \rho_A] = -Tr[\rho_B \log \rho_B] = S(\rho_B) \tag{5.81}$$

Where $S(\rho)$ is the Von Neumann entropy of a state ρ.
Then, we define entanglement entropy as

$$E(\Psi) = -Tr[\rho_a \log \rho_a] \tag{5.82}$$

For this computation, we use the reduced density matrix for the atom for computational reasons. Firstly, as discussed, it is much less expensive to compute compared to the reduced density matrix for the photon. Secondly, calculating entropy from a matrix requires computing its eigenvalues. Since that has complexity $O(n^3)$, with n the dimensionality of the matrix, doing it for the 2×2 atom density matrix will be much faster than for the $NM \times NM$ matrix of the photon state.

5.2.5 Fidelity

Often, we wish to have a quantitative measure of how similar two states are. The standard way to achieve that is by computing the fidelity between the two states.

Fidelity is defined as the square of the norm of the inner product of the two states:

$$F(\psi, \phi) = |\langle \psi | \phi \rangle|^2 = |\langle \phi | \psi \rangle|^2 \tag{5.83}$$

Alternatively, the previous formula can be extended to density matrix states by

$$F(\rho, \sigma) = \left(Tr \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right)^2 \tag{5.84}$$

Both formulas are equal if $\rho = |\Psi_\rho\rangle\langle\Psi_\rho|$ and $\sigma = |\Psi_\sigma\rangle\langle\Psi_\sigma|$ (both states are pure).

We want to use fidelity quantification to determine if states simulated in truncated Hilbert spaces are close to states in untruncated spaces. However, in that case, the inner product is in principle undefined, and so is the matrix product of matrices whose dimensions do not match.

We provide a fidelity quantification method for two states in different Hilbert spaces. Assuming state $|\Psi_A\rangle$ is in space $\mathscr{H}_A$ and $|\Psi_B\rangle$ is in space $\mathscr{H}_B$, and $\mathscr{H}_A \subset \mathscr{H}_B$, then it is possible to find bases for $\mathscr{H}_A$ and $\mathscr{H}_B$, A, B, that fulfill

$$A = \{u_1, u_2, ..., u_N\} \tag{5.85}$$

$$B = \{u_1, u_2, ..., u_N, v_1, v_2, ..., v_M\}$$

Then, we can associate a state in $\mathscr{H}_A$ to a state in $\mathscr{H}_B$ with the following mapping:

$$|\Psi_A\rangle = \sum_i c_i |u_i\rangle \rightarrow |\Psi'_A\rangle = \sum_i^N c_i |u_i\rangle + \sum_j^M 0 \cdot |v_i\rangle \tag{5.86}$$

Where $|\Psi'_A\rangle$ is now a state in $\mathscr{H}_B$.
The inner product of $|\Psi'_A\rangle$ and a state in $\mathscr{H}_B$ can be computed:

$$|\Psi'_A\rangle = \sum_{i=1}^N a_i |u_i\rangle$$

$$|\Psi_B\rangle = \sum_{i=1}^{N} b_i |u_i\rangle + \sum_{i=N+1}^{N+M} b_i |u_i\rangle$$

$$\implies F(\Psi'_A, \Psi_B) = |\langle \Psi'_A | \Psi_B \rangle|^2 = |\sum_i a_i^\star b_i|^2$$

Two identical states will have fidelity one, while two orthogonal states will have zero fidelity. Generally, for two states that have been initialized identically and evolved according to the same Hamiltonian, we expect fidelity to be very close to one since they should be almost identical.

Often, fidelity is expressed in units of "nines". This refers to the number of leading nines in the value of fidelity expressed as a percentage. That is, a fidelity of 99.99% equals "four nines", and 99.999% is "five nines".

5.3 Simulating in two dimensions

We have also developed methods to extend our simulations to two dimensions, allowing us to recreate more realistic systems. Doing this is also essential to simulate more complex optical devices, such as polarizers, by aiming to place the atoms in the device so that they discriminate between directions of propagation.

We start by generalizing the boundary conditions we have been working with to two-dimensional systems. The box now has dimensions (L_x, L_y), with N_x grid points discretizing it in the x direction and N_y grid points for the y direction. Boundary conditions are periodic but must be applied separately for each propagation axis. Formally, the conditions are

$$\begin{aligned} |\psi(0, y)\rangle &= |\psi(L_x, y)\rangle \qquad \forall y \\ |\psi(x, 0)\rangle &= |\psi(x, L_y)\rangle \qquad \forall x \end{aligned} \tag{5.87}$$

These boundary conditions make the box we work with topologically equivalent to a torus, underscoring the fact that the dimensions of the box have to be large compared to the size of the region where the experiment happens to avoid self-interference from the boundary conditions. The goal is to replicate the results we would obtain in free two-dimensional space as closely as possible.

It is unnecessary to effect any changes to the Hamiltonian we use to model the interaction. We only need to relabel the ladder operators:

$$a_k \to a_{\vec{k}} \quad \vec{k} = (k_x, k_y)$$

The modes in each propagation direction run independently of each other, and have, in general, different spacing and limits. Therefore, especially for boxes with $L_x \neq L_y$, care must be taken to set the number of grid points accordingly.

Occupation numbers can be calculated for each vector $\vec{k} = (k_x, k_y)$ as in the one-dimensional case.

Although formally, simulating in two dimensions is equivalent to doing so in one dimension, it is, in practice, a lot more expensive in terms of memory and computation requirements. That is because the dimensionality of the Hilbert space in question is now, assuming we truncate the bosonic space at N_{exc} excitations:

$$dim(\mathscr{H}) = 2N_x^{N_{exc}} N_y^{N_{exc}} \equiv 2N^{2N_{exc}}$$

The scaling implies that, in practice, memory requirements are almost the square of what would be required to simulate a one-dimensional setup with equal precision.

5.4 Sending multiple pulses

In qubit control processes, trains of laser pulses are emitted towards the qubit so that the qubit can be set at a desired state by the combined action of all pulses arriving one after the other. To recreate these systems, we wish to find a simulation technique where multiple pulses are sent one after the other.

A naive approach would be to explicitly include multiple photons in the Hamiltonian. That is, to simulate P pulses, we would write

$$H = \sum_{p=1}^{P} H_{p,a} \tag{5.88}$$

Each of the $H_{p,a}$ would be a standard Hamiltonian, with a relabeling of ladder operators $a_k \to a_{p,k}$ to indicate that they act on only one of the photons. Every photon would be assigned a Hilbert space, which would then give rise to the simulation Hilbert space when multiplied by the space of the atom.

However, this approach is computationally costly, as it is equivalent to multiplying the truncation cutoff for the number of excitations times the number of pulses we wish to simulate. Furthermore, it does not consider the presence of periodic boundary conditions. Unless the box is made very big, which would also require an impractically large number of modes, each pulse would interact with the atom multiple times.

Instead, we assume that only one pulse can be in the simulation region. After the interaction, we compute the reduced density matrix for the atom's state and re-initialize the system with that state and the next photon state in the pulse.

However, the issue with this approach is that initializing the atom state with a density matrix would require us to perform time evolution with density matrices, which would be very costly. To sidestep that, we compute the eigenstate with the largest eigenvalue in the density matrix and use that to initialize the atom. Again, we expect this approximation to be valid when entanglement is weak.

With this method, we can simulate trains of pulses economically, with an error that depends on the degree of entanglement between photon and atom when the pulses in the train are cut off.

6 Numerical methods

When formulating simulation frameworks, it is necessary to translate the problem setting, which is often expressed in terms of operators acting on functions, to a language that can be implemented in a computer. This is due to the continuous nature of most physical problems. Meanwhile, computers work using bits, which are discrete units that can take only two values. Therefore, methods to discretize problems are required that also guarantee that the error from said discretization remains small. To this end, good numerical discretization methods are of the utmost importance.

Another concern when implementing numerical methods is resource limitations. Often, discretization schemes allow us to achieve results that can be made as good as required by increasing the amount of dedicated resources, but that brings requirements to levels that are unfeasible to achieve using current technology. Methods must use available resources efficiently.

Throughout our work, we attempt to use whatever data structures will be most efficient to represent the systems we operate on. For instance, it suffices to represent the joint photon-atom state vector on which we perform simulations as a regular vector on the computer. However, it is necessary to devise structures to represent operators, including density operators, ladder operators, and Hamiltonians, which have dimensions of matrices, efficiently. Otherwise, the memory requirements become too onerous, and it becomes impossible to simulate systems larger than a few modes.

In this chapter, we discuss the approaches we took to deal with the numerical simulation problems we faced.

6.1 Matrix storage

In our work, finding correct schemes to store matrices and vectors is of the utmost importance. We use vectors to store states, and matrices to store operators such as Hamiltonians and ladder operators. It is key to find good data structures to represent them.

The main criterion for deciding what structure is suitable to hold a given matrix or vector is its sparsity. Sparsity is defined as the ratio of null elements to the total number of elements. Often, the converse of sparsity, density, is used, which is simply the number of non-zero elements compared to the total number of elements.

In general, array storage is classified into dense and sparse, which are suited for different levels of sparsity in data structures. Storing a dense matrix will be more efficient for low-sparsity matrices, while for high-sparsity, a sparse matrix will be more efficient.

The key difference lies in the way they store data. A dense matrix requires explicitly specifying the value of every component. Therefore, to store an $n \times m$ matrix, $n \times m$ values will be required. The indices of every element are, nevertheless, not explicitly stored. Instead, the whole matrix is stored in memory as an array, and the indices are instead made explicit by the position of every element in the array.

Dense matrices can be row-major, when matrix elements are placed in memory one row after the other, or column-major, when they are instead placed in memory one column after the other. The choice of ordering is often determined by the default layout of the programming language used, with Python and C++ using row-major ordering. In contrast, Fortran and MATLAB use column-major ordering.

Sparse matrices instead leverage matrix structure by only specifying non-zero elements and their coordinates. If the degree of sparsity is high enough, this becomes more efficient than storing all values.

We present some of the most popular sparse matrix storage formats [Gao et al., 2024]:

Coordinate List (COO)

The Coordinate List (COO) format, also known as the triplet format, stores a sparse matrix using three arrays: one for non-zero values, one for row indices, and one for column indices. For a matrix with k non-zero elements, COO requires $3k$ storage units (assuming integer indices and floating-point values). For a matrix A, COO stores tuples $(i, j, A_{i,j})$ for each non-zero element. For example, a 3x3 matrix with non-zero elements $A_{1,1} = 5$, $A_{2,3} = 7$, and $A_{3,2} = 9$ would be stored with row indices [1, 2, 3], column indices [1, 3, 2], and values [5, 7, 9]. This format is simple and intuitive, making it easy to implement and ideal for matrix construction, as elements can be appended without restructuring. However, it is inefficient for arithmetic operations like matrix-vector multiplication due to uncoordinated index access and poor cache performance. COO is often used as an intermediate format during matrix assembly, but is typically converted to other formats for computation.

Compressed Sparse Row (CSR)

The Compressed Sparse Row (CSR) format is optimized for row-wise operations, using three arrays: one for non-zero values, one for column indices, and one for row pointers, which indicate the starting index of each row's non-zero elements. For a matrix with k non-zero elements and n rows, CSR stores values (length k), column indices (length k), and row pointers (length $n + 1$). For the example 3x3 matrix, CSR might store values [5, 7, 9], column indices [1, 3, 2], and row pointers [0, 1, 2, 3]. CSR is highly efficient for row-wise operations, such as matrix-vector multiplication ($O(k)$ complexity), common in iterative solvers for FDTD simulations, and offers compact storage ($2k + n + 1$ units) with good cache performance. However, modifying the matrix is costly, as it requires restructuring the arrays, and it is less efficient for column-wise operations.

Compressed Sparse Column (CSC)

The Compressed Sparse Column (CSC) format is similar to CSR but optimized for column-wise operations, storing non-zero values, row indices, and column pointers. For the example matrix, CSC might store values [5, 9, 7], row indices [1, 3, 2], and column pointers [0, 1, 2, 3]. CSC is efficient for column-wise operations and has similar memory efficiency to CSR ($2k + m + 1$ for m columns). However, it is less effective for row-wise operations, and modifications are costly.

List of Lists (LIL)

The List of Lists (LIL) format stores a sparse matrix as a collection of lists, where each row is represented by pairs of column indices and non-zero values. For the example 3x3 matrix, LIL stores row 1 as [(1, 5)], row 2 as [(3, 7)], and row 3 as [(2, 9)], often implemented as linked lists or dynamic arrays per row. LIL is highly efficient for incremental matrix construction and modification, as elements can be added without restructuring the entire matrix, making it ideal for adaptive simulations. However, it is inefficient for arithmetic operations due to non-sequential memory access and has higher memory overhead due to list metadata.

In our work, we exploit matrix sparsity whenever possible. The Hamiltonians we use, both Jaynes-Cummings and Quantum Rabi, only have $O(n)$ non-zero elements for an $n \times n$ matrix, so they are quite suited to being modeled by sparse matrices. However, the same cannot be said for density operators. While density matrices at initialization may be sparse depending on the initial conditions, that will not necessarily be true after time evolution. Since we expect state vectors to be dense, density matrices calculated from them will also be dense.

For our simulations, we have opted to use a LIL format during matrix assembly, which is then converted to CSR for operations to improve efficiency.

6.2 Integration schemes

A central role in our work is played by computing the time evolutions of certain states. This is, in effect, equivalent to solving the Schrödinger equation. Analytical solutions are, however, only available for the simplest problems, making numerical integration schemes a necessity.

Numerical integration schemes are essential for approximating the solutions of differential equations when analytical solutions are intractable. They transform continuous problems into discrete ones, allowing for step-by-step computation of system evolutions over time. However, choosing an appropriate integration scheme requires careful consideration of factors such as stability, accuracy, and computational efficiency.

In particular, attention must be paid to the nature of the differential equations, especially whether they are stiff or non-stiff, as this affects the selection between explicit and implicit methods. Additionally, the step size plays a crucial role in ensuring the stability and accuracy of the scheme. Too large a step size may lead to instability or inaccurate results, while too small a step size may result in excessive computational cost.

Stability is a key concept in numerical analysis, especially when dealing with differential equations. A numerical method is said to be stable if small changes in the initial conditions or intermediate computations do not lead to significant deviations in the result as the computation progresses. Stability is particularly significant in the context of stiff equations, where solutions can vary rapidly.

Stability often imposes a condition on the step size Δt, which must be made sufficiently small, especially for explicit methods. Implicit methods, such as backward Euler or advanced backward difference methods, generally provide better stability properties, allowing for larger step sizes without sacrificing accuracy. This is made up for the fact that solving the scheme at each time step to obtain the next value is often more expensive for implicit methods, since numerically solving an equation is required, unlike for explicit methods, where the computation of the next step is straightforward.

There is an abundance of methods that solve or integrate differential equations. We discuss the advantages and disadvantages of some of them.

6.2.1 Forward Euler Method

The forward Euler method is a simple and explicit technique used for solving ordinary differential equations. It updates the solution using the derivative at the current time step:

$$y_{n+1} = y_n + \Delta t f(t_n, y_n). \tag{6.1}$$

While the method is straightforward and easy to implement, it is an explicit method and can suffer from stability issues, especially when applied to stiff equations. Small time steps are often necessary to maintain stability, which can increase computational cost.

6.2.2 Backward Euler Method

The backward Euler method is an implicit alternative to the forward Euler method, offering greater stability, especially for stiff differential equations:

$$y_{n+1} = y_n + \Delta t f(t_{n+1}, y_{n+1}). \tag{6.2}$$

The implicit nature requires solving an equation at each step, often using iterative techniques, such as Newton's method. Despite the increased computational effort per step, the backward Euler method allows for larger time steps while maintaining stability.

6.3 Runge-Kutta Methods

Runge-Kutta methods are a family of iterative methods characterized by their use of intermediate steps. The most well-known and most frequently used is the fourth-order version (RK4):

$$\begin{aligned} k_1 &= f(t_n, y_n) \\ k_2 &= f(t_n + \frac{h}{2}, y_n + \frac{h}{2}k_1) \\ k_3 &= f(t_n + \frac{h}{2}, y_n + \frac{h}{2}k_2) \\ k_4 &= f(t_n + h, y_n + hk_3) \\ y_{n+1} &= y_n + \frac{h}{6}(k_1 + 2k_2 + 2k_3 + k_4) \end{aligned} \tag{6.3}$$

RK4 improves upon Euler schemes by providing stability without the need for very small time steps.

A general Runge-Kutta method aiming to solve $\dot{y} = f(t, y)$ is of the form

$$y_{n+1} = y_n + h\sum_{i=1}^{s} b_i k_i \tag{6.4}$$

$$\begin{aligned} k_1 &= f(t_n, y_n) \\ k_2 &= f(t_n + c_2 h, y_n + (a_{21}k_1)h) \\ k_3 &= f(t_n + c_3 h, y_n + (a_{31}k_1 + a_{32}k_2)h) \\ &\dots \\ k_s &= f(t_n + c_s h, y_n + h\sum_{i=1}^{s-1} a_{si}k_i) \end{aligned}$$

A particular method is then determined by the matrix A with coefficients a_{ij} and the vectors b and c, with coefficients b_i and c_i respectively.

It can be shown that consistency (the numerical solution approaching the differential equation solution) requires

$$\sum_{i=1}^{s} b_i = 1 \tag{6.5}$$

Depending on the number of stages, additional requirements on A, b, c can be found by requiring the solution to converge with a certain speed. The forward Euler scheme is a Runge-Kutta scheme with $b_1 = 1, c_1 = 0$.

6.4 Advanced Backward Difference Methods

Advanced backward difference methods extend the backward Euler approach to achieve higher-order accuracy and stability. The Backward Differentiation Formula (BDF) is a popular choice:

$$\sum_{j=0}^{k} a_j y_{n+j} = h\beta f(t_{n+k}, y_{n+k}), \tag{6.6}$$

The coefficients a_j and β are chosen so that the numerical solution converges to the actual solution of the differential equation with order k. Nevertheless, it is a requirement that $k < 7$ due to stability [Süli and Mayers, 2003].

An advantage they offer over Runge-Kutta methods is increased stability, especially for stiff differential equations.

We employ a mix of Runge-Kutta and Backward Difference integrators for our simulations. As a default, we opt for Runge-Kutta, with Backward Difference schemes reserved for stiff problems, when we observe that Runge-Kutta is showing stability issues.

7 Results

This section evaluates the effectiveness of numerical methods for simulating qubit dynamics under quantum light control, focusing on the fidelity of atom and photon states and the visualization of wave propagation phenomena. Using number and coherent state bases, we assess the accuracy of simulations in both truncated and untruncated Hilbert spaces to quantify the trade-offs between computational efficiency and fidelity for varying interaction strengths. The atom state fidelity is analyzed through metrics like excitation probability and Bloch sphere representations, revealing high accuracy for weak couplings in truncated spaces. We address the computational challenges of density matrix calculations for photon states by deriving fidelity bounds, highlighting limitations in coherent state simulations. Additionally, we present visualizations of one- and two-dimensional wave propagation, capturing interaction effects such as diffraction and photon re-emission. Finally, we discuss challenges in simulating coherent state bases, including numerical stability issues and parameter sensitivity, underscoring the need for optimized simulation frameworks to model complex light-matter interactions accurately.

7.1 Atom state fidelity

7.1.1 Simulating in number basis

We let the photon propagate and interact with the atom from $t = 0$ to $t = 20$, in both a truncated and an untruncated basis. The atom is initialized in the plus state $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, and the photon is initialized to a number state $|1\rangle$.

Figure 7.1 demonstrates that the approximation by truncation of the Hilbert space yields good results, as long as the coupling strength in the interaction is not too large. We achieve fidelities of four nines as long as $g \lesssim 0.1$ in Planck natural units ($\hbar = c = 1$), while managing a speedup of around two orders of magnitude, as demonstrated by figure 7.2. While truncation does provide better and better speedups as coupling is made stronger beyond this point,

The value of the Planck frequency is $\omega_P = \frac{1}{2\pi}\sqrt{\frac{c^5}{\hbar G}} = E_P/h \approx 2.952 \times 10^{42}$ Hz, putting in context the strength of the coupling required for the truncation approximation to fail. In contrast, quantum computers have spin-spin coupling strengths in the hundreds of MHz [Pita-Vidal et al., 2024], much lower than the region where the truncation approximation starts to break down.

7.1.2 Simulating coherent states

Obtaining consistent results has proved much more elusive for simulations of coherent states. A big issue in simulating coherent states is tuning the simulation framework to the parameters. Simulations are highly sensitive to the value of α, which is related to the mean photon number by $\langle n \rangle = |\alpha|^2$.

For simulations in number state basis, the initialization

$$|\psi(0)\rangle = D(\alpha)|0\rangle$$

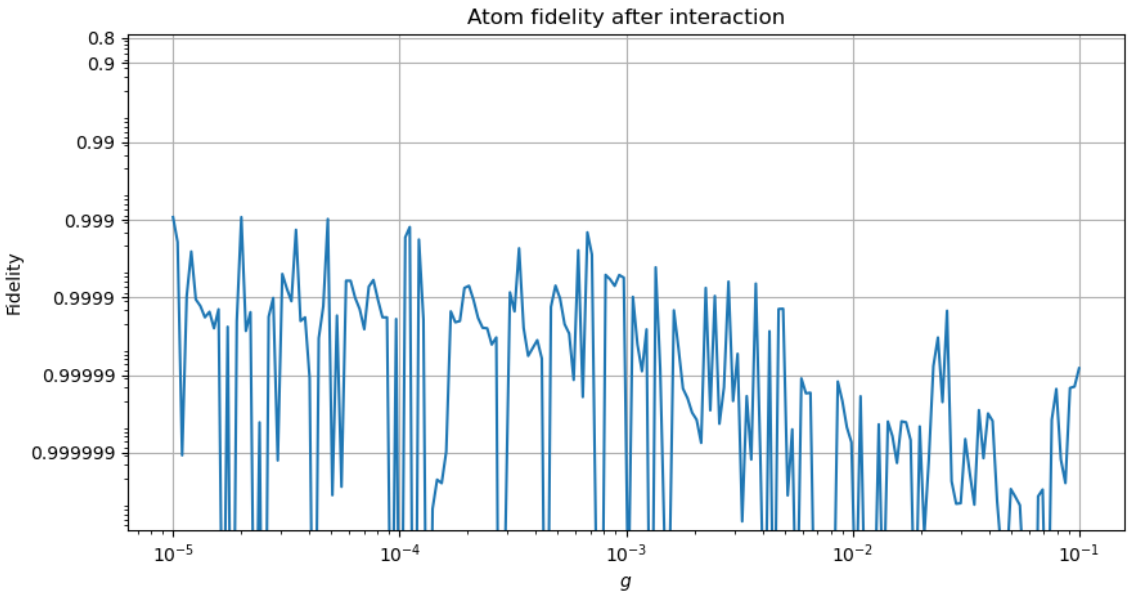

Figure 7.1 Fidelity of the reduced atom state evolved in the truncated Hilbert space versus the untruncated Hilbert space, as a function of interaction strength. The length of the box is $L = 20$, and the photon has been initialized at $x = 0$ and the atom at $x = 10$. The photon is initialized as a Gaussian packet with $k_0 = 1$, propagating to the right, in resonance with the atom. Both spaces are truncated at $N = 3$ excitations.

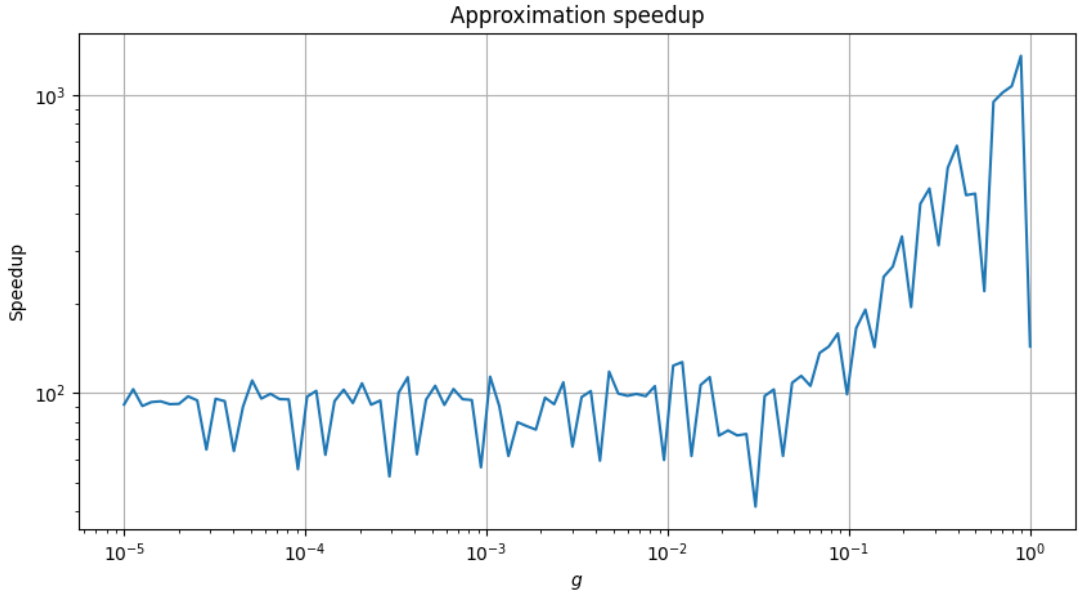

Figure 7.2 Speedup of the simulation on the truncated Hilbert space compared to simulating on the untruncated Hilbert space. Parameters are the same as in figure 7.1.

means that states with excitations across modes, of the form $|n_{k_1}, n_{k_2}\rangle$, are non-zero at initialization, making truncation not a good option.

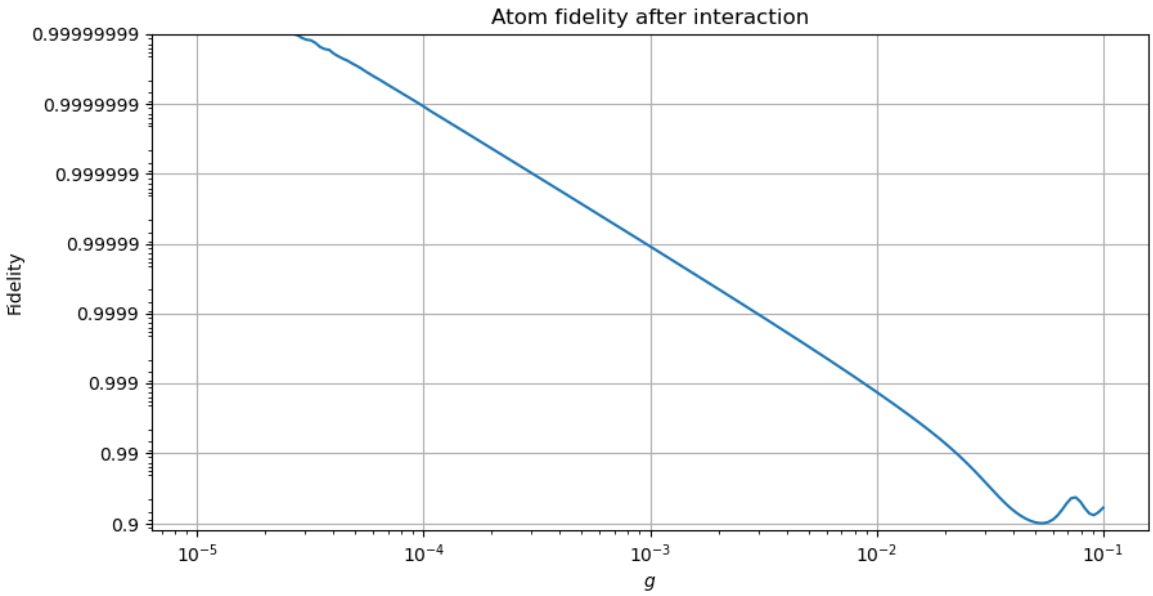

Figure 7.3 Fidelity after interaction as a function of coupling strength of a state simulated in a truncated space of number states vs untruncated. The atom is initialized as a plus state, and the photon is initialized to a coherent state with $\alpha = 1$.

Figure 7.3 shows the fidelity after time evolution of a state evolved in a truncated basis versus an untruncated basis. As expected by the mentioned line of reasoning, our results are slightly worse than initialization at a number state.

7.2 Atom states

The reduced state of the atom after time evolution is expressed by a density operator, which can be understood as a 2×2 matrix. However, matrices are complex to visualize and compare, so we choose two alternative metrics to describe the atom's state.

The most straightforward of the two is excitation probability. It has the advantage of being a positive real number that can be plotted and compared quantitatively.

The other alternative is to use a Bloch sphere representation to visualize the atom's state, which has the advantage of containing all the information in the density matrix in such a way that it is easy to visualize. We have also been able to generate animations, showing how the atomic state changes throughout the interaction.

An arbitrary 2×2 density operator ρ can be depicted on the Bloch sphere by rewriting it as

$$\rho = \frac{1}{2}(I + \vec{n} \cdot \vec{\sigma}) = \frac{1}{2}\begin{pmatrix} 1 + n_z & n_x - in_y \\ n_x + in_y & 1 - n_z \end{pmatrix} \tag{7.1}$$

Here, $\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ can be understood as a vector with the Pauli matrices as components. $\vec{n}$ is the Bloch vector, with $|\vec{n}| \leq 1$, that indicates the state's position in the Bloch sphere.

For a pure state, $tr(\rho^2) = 1 \implies |\vec{n}| = 1$. Therefore, the entanglement of the state can be tracked by looking at how far from the exterior of the Bloch sphere the state moves. In particular, the maximally mixed

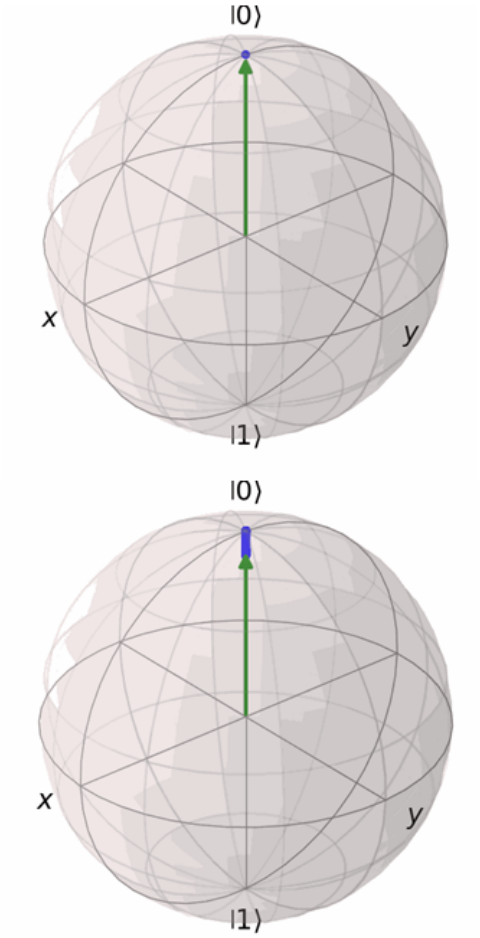

Figure 7.4 Illustration of the evolution of the Bloch sphere representation of an atomic state initialized in the ground state when a photon interacts with it. The topmost sphere corresponds to the state at $t = 0$ and the lowermost sphere to $t = 30$.

state $\rho = I/2$ corresponds to a Bloch vector $\vec{n} = \vec{0}$, meaning that the state would be depicted as located at the center of the sphere.

In Figures 7.4 and 7.5, we show some frames of the animations we have generated that illustrate how the time evolution of the atomic state can be depicted in the Bloch sphere representation. The green arrow indicates the state of the atom when the shot of the frame was taken, while the blue dots represent the previous trajectory of the state.

7.3 Photon state fidelity

We also wish to compute a measure of the quality of our methods for the photon part of the system. We proceed as we did for the atom, obtaining the reduced photonic state for evolution in truncated and untruncated spaces, and then computing the fidelity between them to quantify the quality of the approximation.

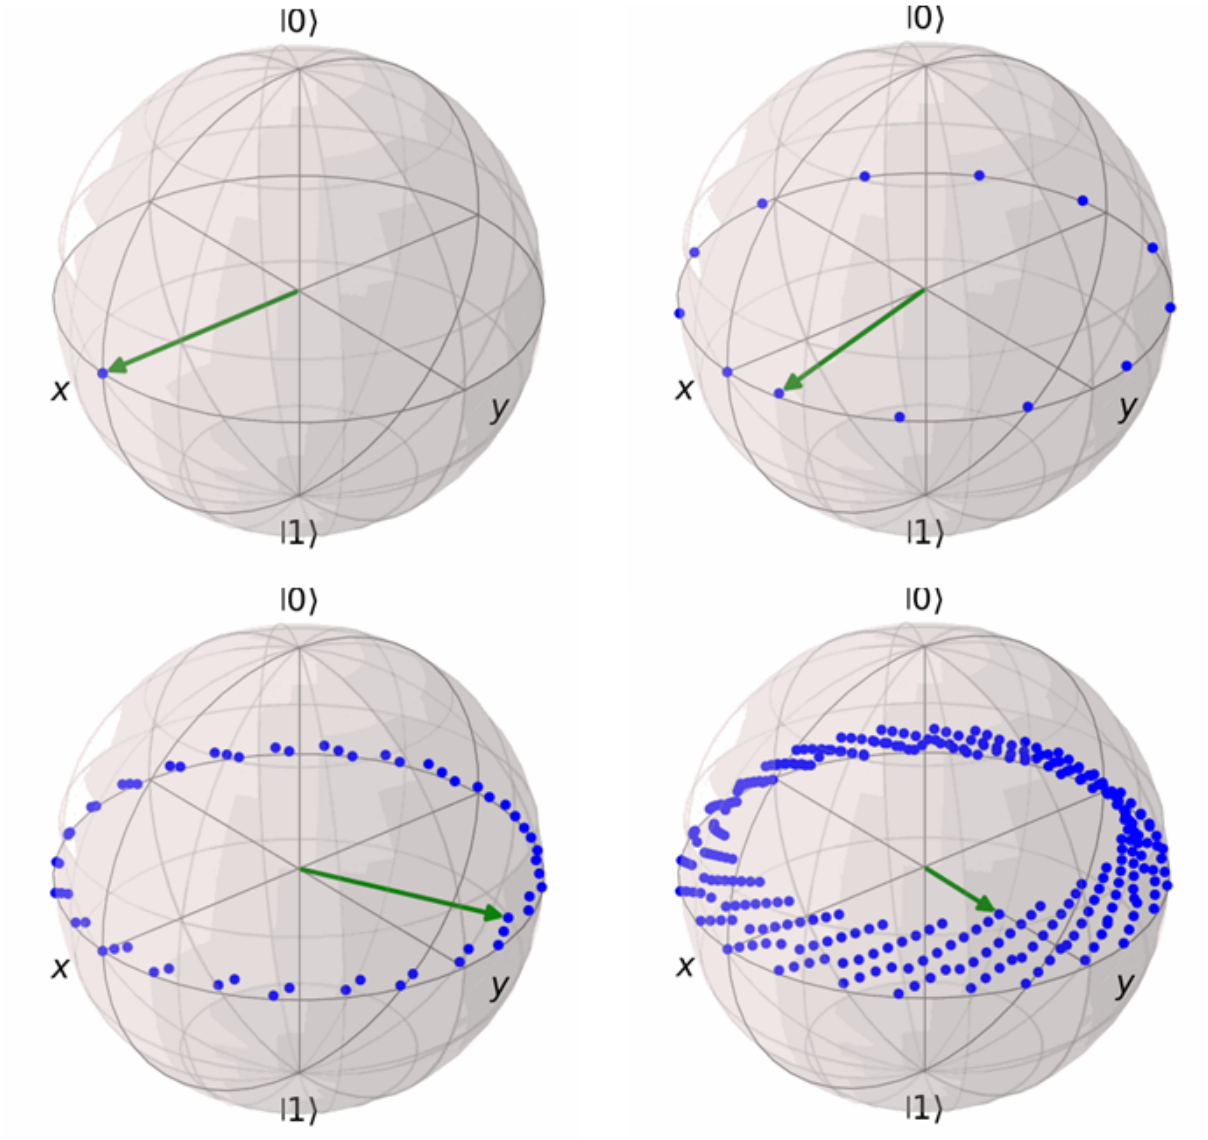

Figure 7.5 Illustration of the evolution of the Bloch sphere representation of an atomic state initialized in the ground state when a photon interacts with it. From left to right and top to bottom, the spheres correspond to the atom states at times $t = 0$, $t = 1.5$, $t = 6$, and $t = 30$.

However, a roadblock to quantifying fidelity for the photonic part of the state is the very high computational cost of computing the density matrices for the photonic part, especially for untruncated Hilbert spaces. Since, for these cases, the curse of dimensionality is very present, it will be impractical to compute a photonic state density matrix, since it will require the square of half of the space in memory that the state vector requires.

Instead, we directly compare the whole state vector, according to the method we outlined in section 5.2.5, allowing us to indirectly measure the fidelity of the photonic part of the system without resorting to computing the full photonic density matrix.

First, let us denote the state obtained by simulating in a truncated Hilbert space as $|\phi\rangle$. The reduced density operators in the atomic and photonic parts of the system are, respectively:

$$\rho_a = Tr_\gamma(|\phi\rangle\langle\phi|)$$
$$\rho_\gamma = Tr_a(|\phi\rangle\langle\phi|)$$

Similarly, we denote the state obtained by simulating in an untruncated space as $|\psi\rangle$, and its reduced density operators:

$$\sigma_a = Tr_\gamma(|\psi\rangle\langle\psi|)$$
$$\sigma_\gamma = Tr_a(|\psi\rangle\langle\psi|)$$

To avoid explicitly computing the fidelity for the photonic part of the system, we instead use the fact that fidelity is monotonic with respect to partial trace [Wilde, 2013]:

$$F(\rho_\gamma, \sigma_\gamma) \geq F(\phi, \psi) \tag{7.2}$$

$$F(\rho_a, \sigma_a) \geq F(\phi, \psi) \tag{7.3}$$

These two inequalities can be multiplied to obtain:

$$F(\rho_\gamma, \sigma_\gamma) \cdot F(\rho_a, \sigma_a) \geq F(\phi, \psi)^2 \tag{7.4}$$

Which can be used to establish a lower bound on photonic state fidelity

$$F(\rho_\gamma, \sigma_\gamma) \geq \frac{F(\phi, \psi)^2}{F(\rho_a, \sigma_a)} \tag{7.5}$$

We have computed this bound for the previous simulation of a system initialized with a number state, the results of which we presented in Figure 7.1. We report the quality of this bound in Figure 7.6.

By comparing the results of Figure 7.6 with those of Figure 7.1, we notice that, while the lower bound on fidelity for the photon is worse than the actual fidelity for the atom, it is still of acceptable quality. While a few nines of fidelity are lost, it is still higher than 99%.

We have also simulated a coherent state in number state space and computed this bound. We present the results in Figure 7.7.

We notice that the bound on fidelity for the photon part of the system is a lot worse than the fidelity we can obtain for the atomic part. That is, we can say that, while the truncation scheme we have devised preserves the dynamics of the atom, making them very similar to the dynamics on the original, untruncated space, the same cannot be said of the photonic state.

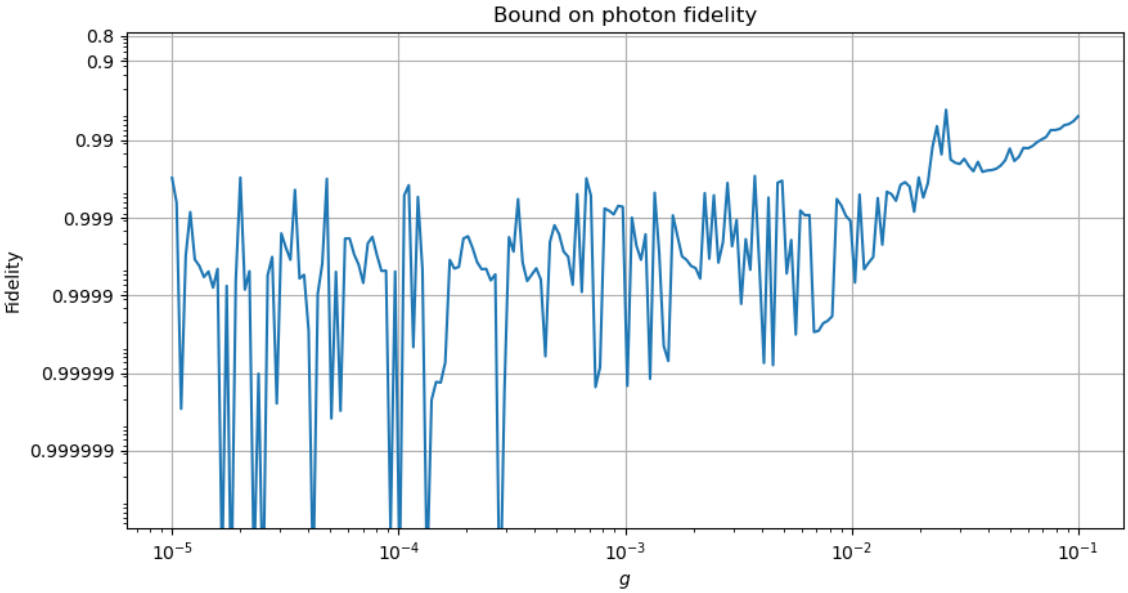

Figure 7.6 Lower bound on the fidelity of the photon part of the system, simulating in a truncated space versus an untruncated space. The photon is initialized to the number state $|1\rangle$.

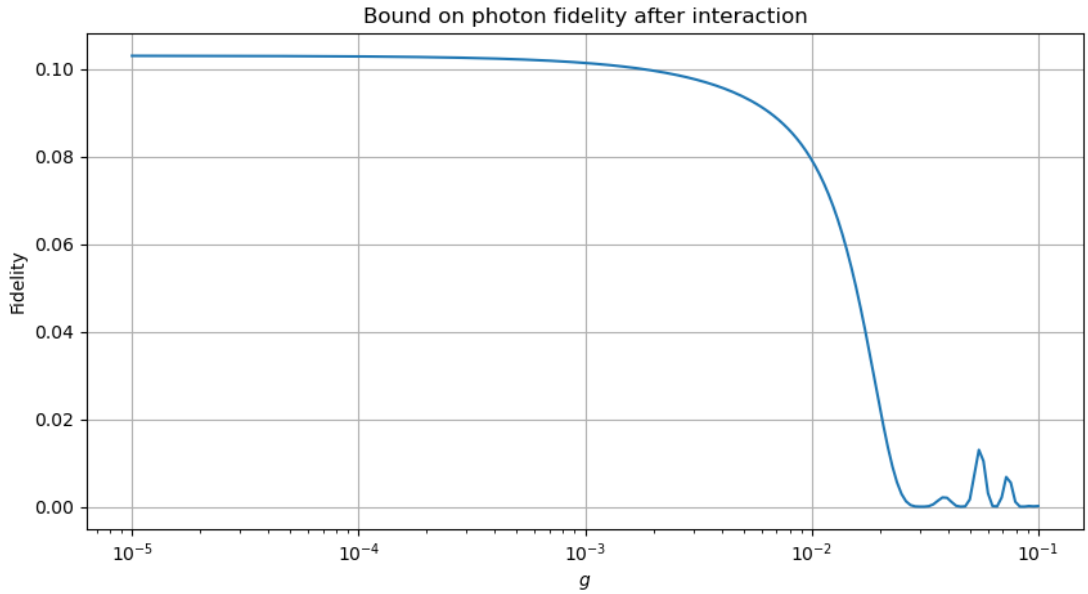

Figure 7.7 Lower bound on photon fidelity in the truncated space versus the untruncated space as a function of coupling strength g.

Furthermore, the quality of the bound on photon fidelity is also much lower than the one we could achieve when we initialized a number state in the boson, as in 7.6. While the previous one allowed us to have confidence when analyzing photon results, the low fidelities achieved for coherent states make it challenging to extract meaningful results from the photonic part of the system.

We observe significant qualitative behavior differences in fidelity when initializing numbers or coherent states because of how the truncation method is constructed.

For number state initialization, the initial state

$$|\psi\rangle = |1_k\rangle$$

is the same in both the untruncated and the truncated spaces. However, the coherent state is:

$$|\psi\rangle = D(\alpha)|0\rangle = \sum_n \frac{1}{n!}(\alpha a^\dagger - \alpha^\star a)^n|0\rangle = \sum_n \frac{1}{n!}\sum_k (\alpha c_k^\star a_k^\dagger - \alpha^\star c_k a_k)|0\rangle$$

This expansion proves that, in a coherent state, excitations across different modes are present, which would then be truncated away by our scheme. Therefore, expecting more information to be lost in this latter case makes sense in comparison to simulating number states.

7.4 Wave propagation

We describe how the photon looks as it propagates. To this end, we compute the expectation value $\langle N_k \rangle$ for each mode k, which informs us of the modes present in the simulation, following the procedure indicated in equation 5.80.

However, it is also desirable to obtain a description of the evolution of the photon as it propagates in position space. To this end, we transform the state vector from the frequency basis in which we have performed the simulations to a position basis using a Fourier transform.

We have generated visualizations of wave propagation, both in one and two dimensions. This allows us to inspect the physical aspect of the wave, which is helpful to ensure that results are coherent with what is expected.

Through analyzing wave propagation, we can also investigate phenomena such as re-emitted waves and diffraction in the presence of the atom. These are especially interesting in two dimensions, due to the appearance of phenomena like interference patterns due to diffraction and emission from the atom.

In Figures 7.8, 7.9, 7.10, and 7.11, we showcase the two-dimensional wave animations we have been able to produce.

We can visualize the processes that intervene in the interaction through the animations. The photon travels freely until it arrives at the atom. Then the interaction happens, at which point the photon's shape stops being Gaussian. The photon then continues traveling in the original propagation direction, albeit with a different wave front shape. By paying close attention to state density in frequency space, re-emitted photons can be observed, corresponding to a tenuous circle of radius $|k| = 5$, equal to the atom's transition frequency.

7.5 Simulating in coherent state bases

Achieving consistent results while working in coherent state bases has proved much more challenging than doing so for number state bases. We review the two methods discussed in sections 5.1.2 and 5.1.2.

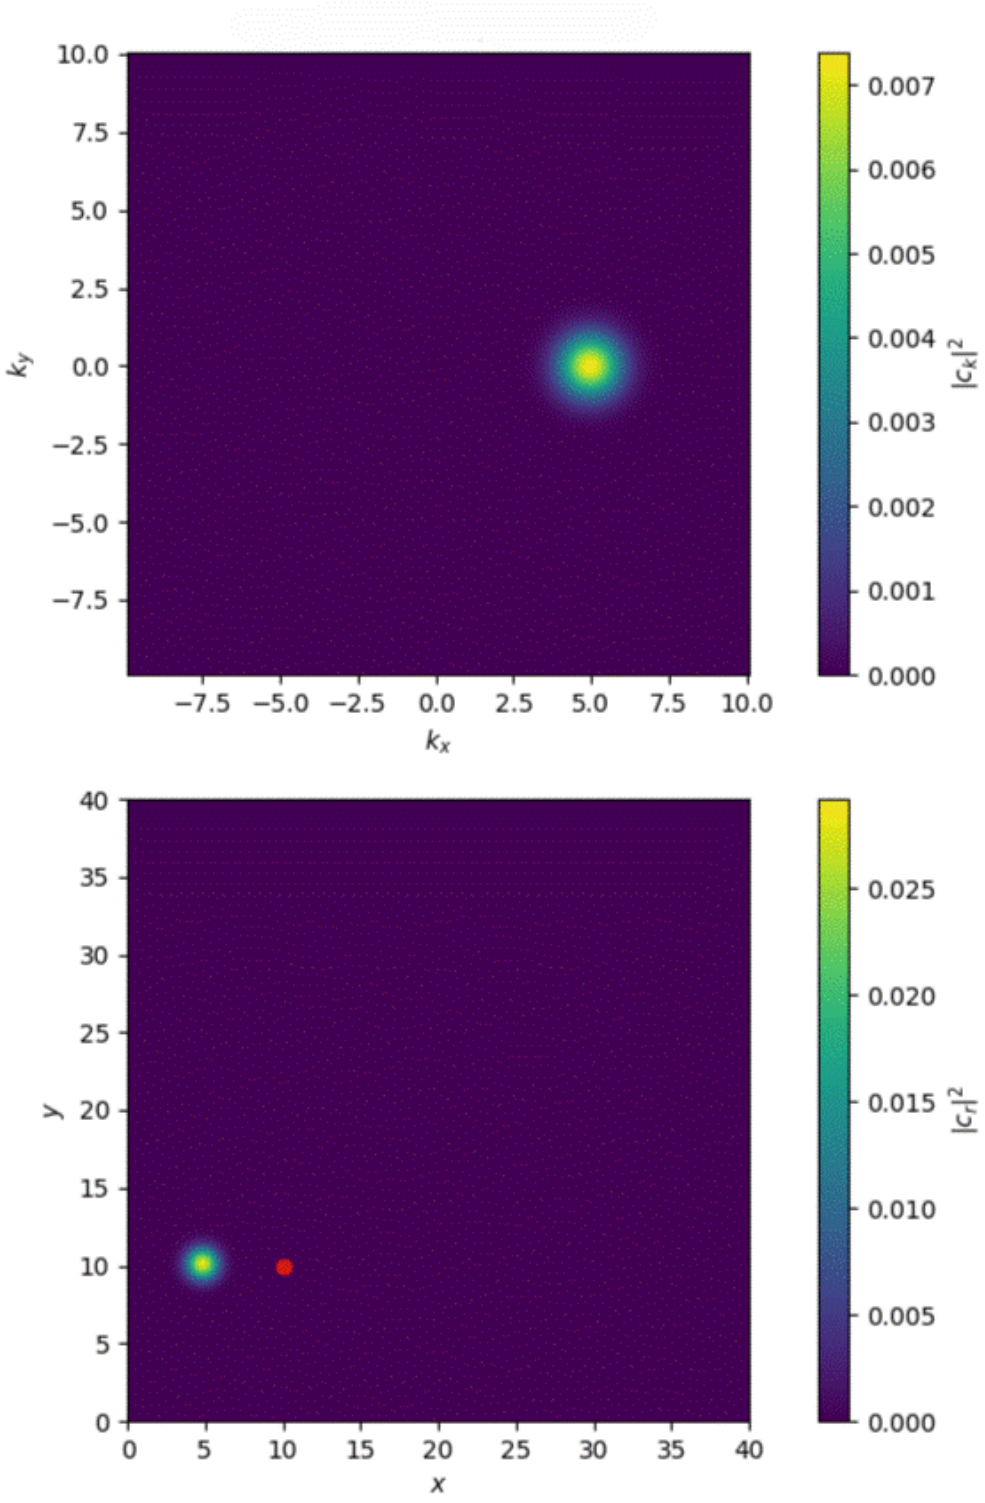

Figure 7.8 Frame of a wave propagation animation in two dimensions, corresponding to $t = 0$. The top plot shows state density in frequency basis, while the bottom plot is in position basis.

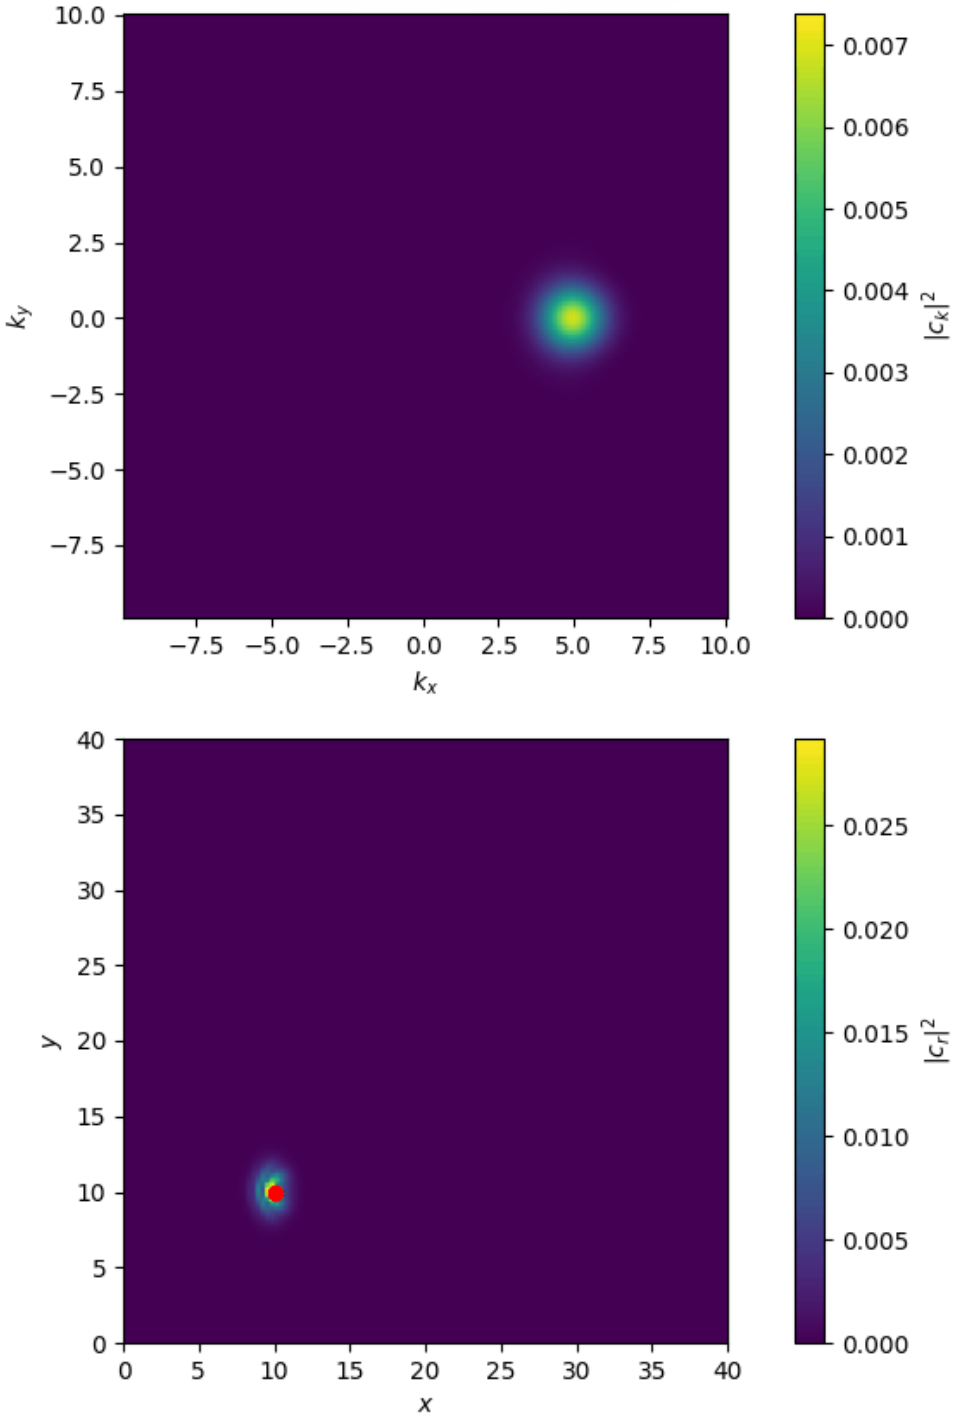

Figure 7.9 Frame of a wave propagation animation in two dimensions, corresponding to $t = 5$. The diffraction patterns generated by interaction with the atom can be noticed. The top plot shows state density in frequency basis, while the bottom plot is in position basis.

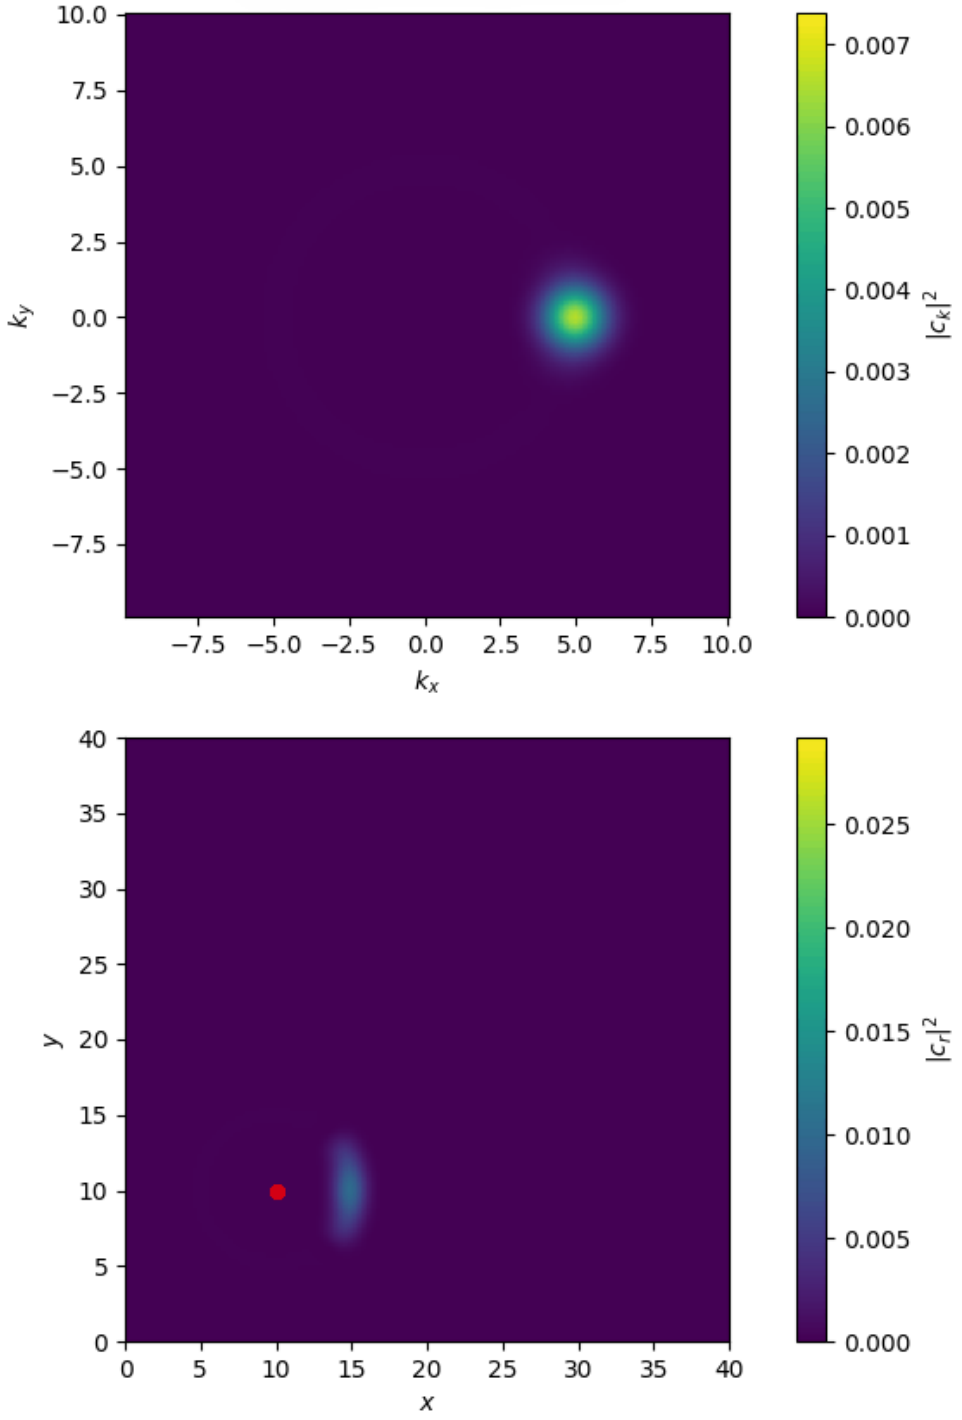

Figure 7.10 Frame of a wave propagation animation in two dimensions, corresponding to $t = 10$. The state of the photon after the interaction can be observed. The top plot shows state density in frequency basis, while the bottom plot is in position basis.

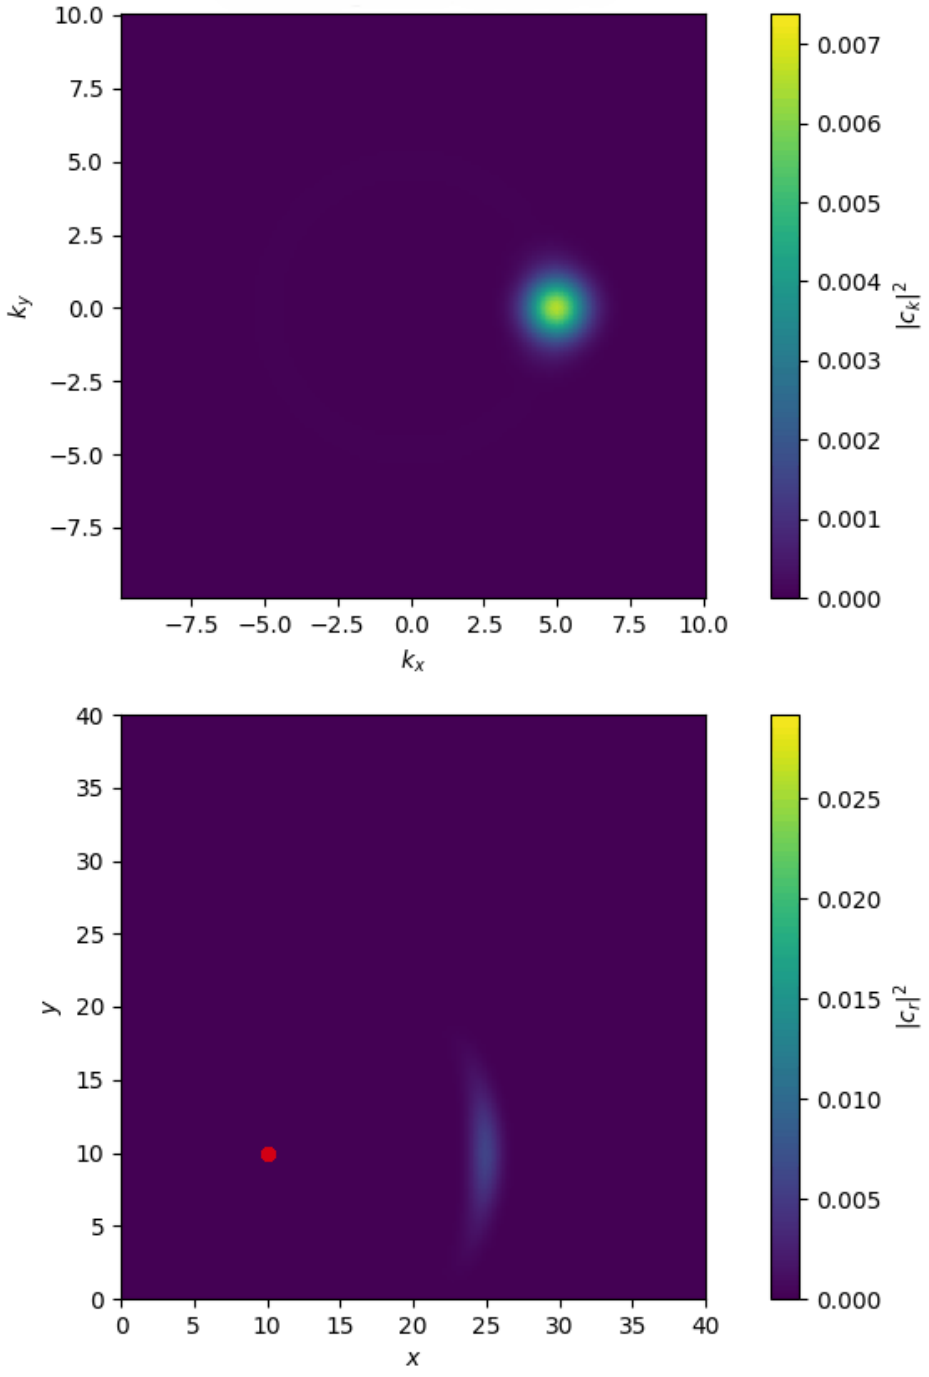

Figure 7.11 Frame of a wave propagation animation in two dimensions, corresponding to $t = 20$. The formation of a wave front can be observed according to Fresnel's principle. The top plot shows state density on a frequency basis, while the bottom plot is on a position basis.

Change of basis with unitary transformation of Hamiltonian

The motivation of this approach is that the time evolution of a coherent state can be written as

$$e^{iHt}|\alpha; s\rangle = D(\alpha)e^{i\tilde{H}t}|0; s\rangle \tag{7.6}$$

$$\tilde{H} = \tilde{H}_0 + \tilde{H}_I \tag{7.7}$$

$$\tilde{H}_0 = H_0 + \sum_k \hbar\omega_k(\alpha^\star c_k a_k - \alpha c_k^\star a_k^\dagger) \tag{7.8}$$

$$\tilde{H}_I = H_I + \sum_{j,k}(\sigma_j^+ + \sigma_j^-)(\alpha c_k^\star e^{ikr_j} + \alpha^\star c_k e^{-ikr_j}) \tag{7.9}$$

Which we can then leverage to recreate the evolution of the coherent state by instead simulating the evolution of a vacuum state according to a different Hamiltonian $\tilde{H}$, which has been obtained by the unitary action of $D(\alpha)$ on H.

Nevertheless, this method presents a significant numerical issue. In the Jaynes-Cummings Hamiltonian, all terms present are of the forms $a^\dagger a,\ a^\dagger\sigma,\ a\sigma^+$, which preserve the number of excitations in the system. In the quantum Rabi Hamiltonian, those are augmented by the $a\sigma, a^\dagger\sigma^+$. The latter, while not conservative, will also not tend to create infinite excitations, since the fact that the atom is only a two-level system means that $(a\sigma)^2 = 0,\ (a^\dagger\sigma^+)^2 = 0$.

The time evolution of the vacuum state can be decomposed, if the time step Δt is small enough, as a succession of n time evolutions:

$$e^{-iHt}|0\rangle \sim e^{-iH\Delta t}...e^{-iH\Delta t}|0\rangle \tag{7.10}$$

Substituting the exponentials by their Taylor expansions up to order n:

$$\sum_{m_1=0}^{n}\frac{1}{m_1!}(-iH\Delta t)^{m_1}...\sum_{m_n=0}^{n}\frac{1}{m_n!}(-iH\Delta t)^{m_n}|0\rangle \tag{7.11}$$

We analyze the problematic term that creates photons, so we can let $H \sim -\hbar w_k c_k^\star \alpha a_k^\dagger$.

$$\sum_{m_1=0}^{n}\frac{1}{m_1!}(i\hbar w_k c_k^\star\alpha\Delta t a_k^\dagger)^{m_1}...\sum_{m_n=0}^{n}\frac{1}{m_n!}(i\hbar w_k c_k^\star\alpha\Delta t a_k^\dagger)^{m_n}|0\rangle = \tag{7.12}$$

$$\sum_{m_1=0,...,m_n=0}^{n}\frac{1}{m_1!...m_n!}(i\hbar w_k c_k^\star\alpha\Delta t a_k^\dagger)^{m_1+...+m_n}|0\rangle$$

We can also use the definition of a Fock state to derive that $(a_k^\dagger)^n|0\rangle = \sqrt{n!}|n_k\rangle$. Now, we focus on the terms generating n photons, so $m_1 + ... + m_n = n$:

$$(i\hbar w_k c_k^\star\alpha\Delta t)^n\sqrt{n!}|n_k\rangle \sum_{\substack{m_1=0,...,m_n=0\\ m1+...+m_n=n}}^{n}\frac{1}{m_1!...m_n!} \tag{7.13}$$

We now realize that $i = e^{i\pi/2}$ and $c_k^\star = |c_k|e^{-i\theta}$, so the term is now:

$$e^{in(\pi/2-\theta)}(\hbar w_k|c_k|\alpha\Delta t)^n\sqrt{n!}|n_k\rangle \sum_{\substack{m_1=0,...,m_n=0\\ m1+...+m_n=n}}^{n} \frac{1}{m_1!...m_n!} \tag{7.14}$$

We are interested only in the absolute value of the coefficients, so we can drop the complex phase.

The sum can be computed through some algebra, leveraging the Taylor expansion of exponential functions. We start with the product of n exponentials with variables $x_1, x_2, ..., x_n$:

$$e^{x_1}e^{x_2}...e^{x_n} = \sum_{m_1=0}^{\infty}\frac{1}{m_1!}\sum_{m_2=0}^{\infty}\frac{1}{m_2!}...\sum_{m_1=0}^{\infty}\frac{1}{m_n!} = \sum_{m_1=0,...,m_n=0}^{\infty}\frac{(x_1)^{m_1}(x_2)^{m_2}...(x_n)^{m_n}}{m_1!m_2!...m_n!} \tag{7.15}$$

We now let $x_1 = x_2 = ... = x_n = x$:

$$\sum_{k=0}^{\infty}\frac{(nx)^k}{k!} = e^{nx} = \sum_{m_1=0,...,m_n=0}^{\infty}\frac{x^{m_1+m_2+...+m_n}}{m_1!m_2!...m_n!} = \sum_{k=0}^{\infty}\sum_{\substack{m_1=0,...,m_n=0\\ m1+...+m_n=n}}^{k}\frac{x^{m_1+m_2+...+m_n}}{m_1!m_2!...m_n!} \tag{7.16}$$

Comparing the first and last terms, the value of the sum can be identified:

$$\sum_{\substack{m_1=0,...,m_n=0\\ m1+...+m_n=n}}^{k}\frac{x^{m_1+m_2+...+m_n}}{m_1!m_2!...m_n!} = \frac{(nx)^k}{k!} \tag{7.17}$$

And, setting $x = 1$ and $k = n$, we arrive at the result we were looking for:

$$\sum_{\substack{m_1=0,...,m_n=0\\ m1+...+m_n=n}}^{n}\frac{1}{m_1!m_2!...m_n!} = \frac{n^n}{n!} \tag{7.18}$$

Going back to the time evolution coefficients we were studying, we are left with:

$$(\hbar w_k|c_k|\alpha\Delta t)^n\sqrt{n!}\frac{n^n}{n!}|n_k\rangle = (\hbar w_k|c_k|\alpha\Delta t)^n\frac{n^n}{\sqrt{n!}}|n_k\rangle \tag{7.19}$$

The method's stability requires that excitations do not keep adding up until infinity, which is equivalent to a condition stating that the norm of the coefficients of the corresponding terms in the time evolution expansion must remain bounded.

$$(\hbar w_k|c_k|\alpha\Delta t)^n\frac{n^n}{\sqrt{n!}} < \infty \quad \text{for } n \to \infty \tag{7.20}$$

This is equivalent to studying the limit of a succession of the form:

$$S_n(\beta) = \beta^n\frac{n^n}{\sqrt{n!}} \qquad \beta = \hbar\Delta t\omega_k\alpha|c_k| \tag{7.21}$$

To guarantee convergence, it is necessary that

$$\lim_{n\to\infty} S_n(\beta) < \infty$$

We can use Stirling's formula to analyze the term with the factorial:

$$n! \sim \sqrt{2\pi n}\Big(\frac{n}{e}\Big)^n = \sqrt{2\pi n} n^n e^{-n} \tag{7.22}$$

$$\sqrt{n!} \sim (2\pi n)^{1/4} n^{n/2} e^{-n/2} \tag{7.23}$$

Therefore, we can rewrite S_n for large n as:

$$S_n(\beta) \sim \beta^n n^n (2\pi n)^{-1/4} n^{-n/2} e^{+n/2} \sim e^{nln\beta} e^{-ln(n)/4} e^{nln(n)/2} e^{n/2} = e^{ln(n)(n/2-1/4)+n(1/2+ln(\beta))} \tag{7.24}$$

The dominant term in the argument of the exponential grows like $nln(n)$. Therefore, the limit is equal to, for any $\beta > 0$:

$$\lim_{n\to\infty} S_n(\beta) = +\infty \tag{7.25}$$

Going back to the original problem, this means that the coefficients of the large photon terms in the state vector diverge:

$$\sum_k (\hbar \Delta t \omega_k \alpha c_k)^n \frac{n^n}{\sqrt{n!}} |n_k\rangle \to \infty \tag{7.26}$$

And, therefore, the scheme will be numerically unstable, no matter the choice of parameters, or how small the time step and the mean photon number in the state are selected.

This makes using the scheme to simulate coherent states unfeasible, since the numerical solution will never converge to the actual physical solution.

Direct simulation in a constructed basis of coherent states

In this alternative approach, we directly construct a basis of coherent states, where we perform the simulation, forgoing Fock states altogether. A basis on a hexagonal grid is constructed for every mode.

The grid is constructed so that the points in coherent phase space are equally spaced from each other, and are located far apart enough to be assumed to be orthogonal. The hexagonal packing of grid points guarantees that points in the middle of the grid will be as close as possible to the grid points, which means they will be well-represented by linear combinations of grid points.

Despite these steps taken to optimize the scheme, we have observed that the simulations yield results highly dependent on scheme parameters, such as the number of modes and the number of points in the grid basis. Furthermore, verifying the method presents issues related to computational cost: at least six basis states per mode are necessary, equivalent in computational requirements to simulating six excitations on a number state basis. This makes simulation with more than a few modes too costly to carry out in non-truncated bases, so it is hard to determine if the error originates from the construction of the coherent state basis or the posterior truncation of the Hilbert space. Simulating the simplest possible hexagonal grid, consisting of the origin in phase space and six points around it, becomes impractical already with eight modes in the system, which is nowhere near enough for a realistic propagating wave.

To realize this scheme, more rigorous methods should be developed to determine optimal parameters for simulation.

8 Conclusion

This thesis aimed to develop numerical methods for modeling qubit dynamics under quantum light control, addressing the curse of dimensionality inherent in light-matter interactions. Accurate simulation of these dynamics is crucial for advancing fault-tolerant quantum computing, particularly in architectures like superconducting circuits and trapped ions, where precise photonic control is essential. By focusing on the Quantum Rabi Hamiltonian and exploring techniques such as Hilbert space truncation and coherent state basis methods, this work sought to achieve high-fidelity simulations with manageable computational resources.

We have developed a method to accelerate the simulation of quantum Rabi Hamiltonians by truncating the Hilbert space to exclude basis states with photon excitations spread across different modes. We have applied this scheme successfully to simulate the time evolution of the atom, achieving fidelities of up to four nines for weak coupling ($g \lesssim 0.1$). Significant speedups are achieved, of about two orders of magnitude.

Nevertheless, preserving accuracy for photonic states has proved more challenging. Due to the computational complexity of directly computing fidelities for many-mode photons, we develop an easy-to-compute bound to indirectly determine accuracy. While the truncation method also shows promising results for photons initialized to number states, obtaining high fidelities has proved more difficult for coherent states. It is particularly challenging to get consistent results for the photonic subsystem compared to the atom, leading to tradeoffs that require consideration when choosing an appropriate simulation framework. We have attempted to sidestep this with methods simulating in bases constructed in different ways, from changes of basis to direct definition of a grid basis, but these all showed their own issues and pitfalls.

We have developed methods of visualizing time evolution for both atom and photon, ranging from Bloch sphere animation to state density representations of how the photon wave travels through position space.

Our work has implications in quantum computing, particularly in quantum control. Gaining a better understanding of atom-matter interaction allows for more accurate pulse design by enabling design with lower physical error rates. Our methods can find application in a wide range of physical implementations, ranging from superconducting circuits to trapped ions.

8.1 Limitations and challenges

Despite the successes, the methods we have developed face limitations. The scheme breaks down for large interaction coupling strengths ($g \gtrsim 0.1$). Furthermore, it has been challenging to create methods that can consistently simulate coherent states for a wide range of system parameters.

Furthermore, dealing with photonic states has required us to use methods that sacrifice precision to make problems computationally tractable. We have exclusively determined fidelity for the photon part through bounds, which are inherently less precise than exact fidelity computations. While in some cases the bound was good enough to guarantee a certain level of quality of the results, as in the case of number states, in other cases the low quality of the bound makes simulation results unusable, as the results cannot be trusted to resemble those for the actual, unapproximated space.

8.2 Future work

There are multiple avenues where future work could make a substantial impact. A key focus should be on developing more robust methods to simulate coherent states. A possible approach would be to optimize the parameters of the hexagonal grid-based approach to simulate systems that are initialized with coherent beams accurately. To achieve this, approaches based on hyperparameter optimization would probably be necessary. Another option would be to explore alternative changes of basis or conditioning schemes that yield systems with better numerical properties, allowing for simulation at higher mean photon number regimes.

Further performance optimizations could also be explored in the code to achieve even higher speedups. Currently, the simulations run exclusively sequentially, so headroom exists to move to distributed execution. While parallelizing time evolution is more challenging, due to its inherently sequential nature, we believe there is low-hanging fruit to be picked by making reduced state computations as well as Hamiltonian initialization parallel.

Another interesting research direction would be a more detailed study of multi-pulse simulations. While we have provided the tools to simulate pulse trains, insights remain to be discovered on the behavior of coupled photon-atom systems under driving. This would have immediate practical application in qubit control, since, by means of repeated simulation, efficient pulse design methods could be constructed.

Yet another family of schemes that could be explored is schemes with only partial truncation. Instead of truncating all states with excitations across different modes, states fulfilling a number of conditions could be kept. An option would be to explore keeping states with photon frequencies close to resonance in the simulation: it is reasonable to expect them to become populated from the re-emission of atom excitations. Schemes tailored to the parameters of individual problems could also be developed, with corresponding truncation strategies: from the wave visualizations we have generated, we know that most modes see close to no excitation, so not including them in the simulation should also yield similar results.

It should also be possible to investigate methods to mitigate the breakdown in fidelity at large coupling strengths. Partial truncation may also be a solution, including, for example, states with two excitations across modes, but not more, to keep computational requirements reasonable.

Another method that could be worth exploring is, instead of using truncation to reduce the size of the Hilbert space, using a quantum autoencoder architecture, along the guidelines of [Romero et al., 2017]. The state vector in the large-dimensional Hilbert space could be compressed optimally to a small-dimensional latent space, and then time evolution could be computed on the latent space. As a disadvantage, further research would need to be carried out on methods to efficiently time-evolve latent space states.

8.3 Concluding remarks

This thesis has advanced the numerical modeling of qubit dynamics controlled by quantum light, achieving high-fidelity simulations for weak enough coupling regimes and providing valuable insights into light-matter interactions through wave propagation visualizations. This work lays a foundation for scalable, accurate simulation frameworks critical for fault-tolerant quantum computing by addressing the curse of dimensionality with innovative truncation and basis methods. Continued research building on these findings will be essential to overcome current limitations and realize the full potential of quantum technologies.

Bibliography

[Bacon, 2003] Bacon, D. (2003). Decoherence, control, and symmetry in quantum computers.

[Barenco et al., 1995] Barenco, A., Bennett, C. H., Cleve, R., DiVincenzo, D. P., Margolus, N., Shor, P., Sleator, T., Smolin, J. A., and Weinfurter, H. (1995). Elementary gates for quantum computation. *Physical Review A*, 52(5):3457–3467.

[Barone and Paterno, 1982] Barone, A. and Paterno, G. (1982). *Physics and Applications of the Josephson Effect*. A Wiley-interscience publication. Wiley.

[Bruzewicz et al., 2019] Bruzewicz, C. D., Chiaverini, J., McConnell, R., and Sage, J. M. (2019). Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2).

[Büttiker, 1987] Büttiker, M. (1987). Zero-current persistent potential drop across small-capacitance josephson junctions. *Phys. Rev. B*, 36:3548–3555.

[Case, 2008] Case, W. B. (2008). Wigner functions and weyl transforms for pedestrians. *American Journal of Physics*, 76(10):937–946.

[Chae et al., 2024] Chae, E., Choi, J., and Kim, J. (2024). An elementary review on basic principles and developments of qubits for quantum computing. *Nano Convergence*, 11(11).

[Cirac and Zoller, 1995] Cirac, J. I. and Zoller, P. (1995). Quantum computations with cold trapped ions. *Phys. Rev. Lett.*, 74:4091–4094.

[Clerk et al., 2010] Clerk, A. A., Devoret, M. H., Girvin, S. M., Marquardt, F., and Schoelkopf, R. J. (2010). Introduction to quantum noise, measurement, and amplification. *Reviews of Modern Physics*, 82(2):1155–1208.

[Devoret et al., 2004] Devoret, M. H., Wallraff, A., and Martinis, J. M. (2004). Superconducting qubits: A short review. *arXiv: Mesoscale and Nanoscale Physics*.

[DiVincenzo, 2000] DiVincenzo, D. P. (2000). The physical implementation of quantum computation. *Fortschritte der Physik*, 48:771–783.

[Fowler et al., 2012] Fowler, A. G., Mariantoni, M., Martinis, J. M., and Cleland, A. N. (2012). Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3).

[Gao et al., 2024] Gao, J., Liu, B., Ji, W., and Huang, H. (2024). A systematic literature survey of sparse matrix-vector multiplication.

[Havukainen et al., 1999] Havukainen, M., Drobný, G., Stenholm, S., and Bužek, V. (1999). Quantum simulations of optical systems. *Journal of Modern Optics*, 46(9):1343–1367.

[Krantz et al., 2019] Krantz, P., Kjaergaard, M., Yan, F., Orlando, T. P., Gustavsson, S., and Oliver, W. D. (2019). A quantum engineer's guide to superconducting qubits. *Applied Physics Reviews*, 6(2).

[Larson and Mavrogordatos, 2024] Larson, J. and Mavrogordatos, T. (2024). *The Jaynes–Cummings Model and its Descendants (Second Edition): Modern research directions*. IOP Publishing.

[Mandel and Wolf, 1995] Mandel, L. and Wolf, E. (1995). *Optical Coherence and Quantum Optics*. Cambridge University Press.

[Marshall and Anand, 2023] Marshall, J. and Anand, N. (2023). Simulation of quantum optics by coherent state decomposition. *Optica Quantum*, 1(2):78.

[Nielsen and Chuang, 2010] Nielsen, M. A. and Chuang, I. L. (2010). *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press.

[Oba et al., 2024] Oba, J., Kajita, S., and Soeda, A. (2024). Fast simulation for multi-photon, atomic-ensemble quantum model of linear optical systems addressing the curse of dimensionality. *Scientific Reports*, 14(1).

[Picken et al., 2023] Picken, C. J., Despard, I., Kelly, A., Pritchard, J. D., Bain, J. R. P., Hempler, N., Maker, G. T., and Malcolm, G. P. A. (2023). Laser systems for high fidelity control and entanglement of neutral atomic qubits.

[Pita-Vidal et al., 2024] Pita-Vidal, M., Wesdorp, J. J., Splitthoff, L. J., Bargerbos, A., Liu, Y., Kouwenhoven, L. P., and Andersen, C. K. (2024). Strong tunable coupling between two distant superconducting spin qubits. *Nature Physics*, 20(7):1158–1163.

[Roffe, 2019] Roffe, J. (2019). Quantum error correction: an introductory guide. *Contemporary Physics*, 60(3):226–245.

[Romero et al., 2017] Romero, J., Olson, J. P., and Aspuru-Guzik, A. (2017). Quantum autoencoders for efficient compression of quantum data. *Quantum Science and Technology*, 2(4):045001.

[Rossmann, 2006] Rossmann, W. (2006). *Lie Groups: An Introduction Through Linear Groups*. Oxford Mathematics. Oxford University Press.

[Sakurai, 1994] Sakurai, J. J. (1994). *Modern quantum mechanics; rev. ed.* Addison-Wesley, Reading, MA.

[Schlosshauer, 2019] Schlosshauer, M. (2019). Quantum decoherence. *Physics Reports*, 831:1–57.

[Schmidt-Kaler et al., 2003] Schmidt-Kaler, F., Häffner, H., Riebe, M., Gulde, S., Lancaster, G. P. T., Deuschle, T., Becher, C., Roos, C. F., Eschner, J., and Blatt, R. (2003). Realization of the cirac–zoller controlled-not quantum gate. *Nature*, 422(6930):408–411. Received 18 December 2002; Accepted 17 February 2003.

[Schumacher, 1995] Schumacher, B. (1995). Quantum coding. *Phys. Rev. A*, 51:2738–2747.

[Shah, 2010] Shah, D. (2010). *Design of Regular Reversible Quantum Circuits Design of Regular Reversible Quantum Circuits*. PhD thesis, Portland State University.

[Shankar, 1980] Shankar, R. (1980). *Principles of quantum mechanics*. Plenum, New York, NY.

[Süli and Mayers, 2003] Süli, E. and Mayers, D. (2003). *An Introduction to Numerical Analysis*. Cambridge University Press.

[Weinbub and Ferry, 2018] Weinbub, J. and Ferry, D. K. (2018). Recent advances in wigner function approaches. *Applied Physics Reviews*, 5(4):041104.

[Wilde, 2013] Wilde, M. M. (2013). *Distance Measures*, page 203–231. Cambridge University Press.

[Wong, 2025] Wong, H. (2025). *Quantum Computing Architecture and Hardware for Engineers: Step by Step*. Springer Nature Switzerland.

[Wootters et al., 1982] Wootters, W. K., Wootters, W. K., and Zurek, W. H. (1982). A single quantum cannot be cloned. *Nature*, 299:802–803.