

Aspects of Quantum Information in Quantum Field Theory:

Particle detector models, entanglement, and complexity

by

Bruno de Souza Leão Torres

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics (Quantum Information)

Waterloo, Ontario, Canada, 2024

© Bruno de Souza Leão Torres 2024

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner: Lucas Hackl
Senior Lecturer, School of Mathematics and Statistics & Physics,
University of Melbourne

Supervisors: Eduardo Martín-Martínez
Professor, Dept. of Applied Mathematics,
University of Waterloo

Beni Yoshida
Professor, Perimeter Institute for Theoretical Physics,
Adjunct Professor, Dept. of Physics & Astronomy,
University of Waterloo

Internal Member: Niayesh Afshordi
Professor, Dept. of Physics & Astronomy,
University of Waterloo

Internal-External Member: Achim Kempf
Professor, Dept. of Applied Mathematics,
University of Waterloo

Other Member(s): Robert Myers
Professor, Perimeter Institute for Theoretical Physics,
Adjunct Professor, Dept. of Physics & Astronomy,
University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

This thesis contains results from the following works:

- (1) **Bruno de S. L. Torres**, “Particle detector models from path integrals of localized quantum fields”, [Phys. Rev. D **109**, 065004 \(2024\)](#) [[177](#)];
- (2) Daniel Grimmer, **Bruno de S. L. Torres**, Eduardo Martín-Martínez, “Measurements in QFT: Weakly coupled local particle detectors and entanglement harvesting”, [Phys. Rev. D **104**, 085014 \(2021\)](#) [[86](#)];
- (3) T. Rick Perche, José Polo-Gómez, **Bruno de S. L. Torres**, Eduardo Martín-Martínez, “Fully relativistic entanglement harvesting”, [Phys. Rev. D **109**, 045018 \(2024\)](#) [[136](#)];
- (4) **Bruno de S. L. Torres**, Kelly Wurtz, José Polo-Gómez, Eduardo Martín-Martínez, “Entanglement structure of quantum fields through local probes”, [J. High Energy Phys. **2023**, 58 \(2023\)](#) [[59](#)];
- (5) **Bruno de S. L. Torres**, Eduardo Martín-Martínez, “Toward a physically motivated notion of Gaussian complexity geometry”, [Phys. Rev. D **110**, 025006 \(2024\)](#) [[178](#)].

The results from paper (1) appear in all chapters of Part [I](#) of this thesis. This was a single-author paper, with all contributions therein being my responsibility.

Results from papers (2) and (3) appear in Chapter [5](#) of Part [II](#). I was the main responsible for the calculations in Sec. [5.2](#) and [5.4.1](#). The bulk of the calculations and plots in the remainder of Chapter [5](#) were first due to Daniel Grimmer (in Section [5.3](#)) and T. Rick Perche (in Section [5.4.2](#)). The text was prepared and edited collaboratively by all co-authors.

Paper (4) appears in Chapter [6](#) of Part [II](#). I was responsible for the proof of the main result in Section [6.3](#), as well as for writing the first draft of the paper (with the exception of the text and plots of Section [6.5](#), which were written at first by José Polo-Gómez). The remaining plots were produced by Kelly Wurtz. The full text was then edited collaboratively by all co-authors.

Finally, paper (5) appears in all chapters of Part [III](#). I performed all of the calculations and was also responsible for writing the first version of the draft, which was then edited collaboratively by both co-authors.

Additional contributions published during my PhD but not included explicitly in this thesis:

- (6) T. Rick Perche, José Polo-Gómez, **Bruno de S. L. Torres**, Eduardo Martín-Martínez, “Particle detectors from localized quantum field theories”, [Phys. Rev. D **109**, 045013 \(2024\)](#);
- (7) I. M. Burbano, T. Rick Perche, **Bruno de S. L. Torres**, “A path integral formulation for particle detectors: the Unruh-DeWitt model as a line defect”, [J. High Energy Phys. **2021**, 76 \(2021\)](#).
- (8) Eduardo Martín-Martínez, T. Rick Perche, **Bruno de S. L. Torres**, “Broken covariance of particle detector models in relativistic quantum information”, [Phys. Rev. D **103**, 025007 \(2021\)](#).
- (9) Lin-Qing Chen, Wan Zhen Chua, Shuwei Liu, Antony J. Speranza, **Bruno de S. L. Torres**, “Virasoro Hair and Entropy for Axisymmetric Killing Horizons”, [Phys. Rev. Lett. **125**, 241302 \(2020\)](#).
- (10) **Bruno de S. L. Torres**, T. Rick Perche, André G. S. Landulfo, George E. A. Matsas. “Neutrino flavor oscillations without flavor states”, [Phys. Rev. D **102**, 093003 \(2020\)](#).

Abstract

This thesis explores three themes in the interface between quantum information and quantum field theory (QFT).

Part I is devoted to *particle detector models*, which are one of the key ingredients in formulations of a measurement theory for quantum fields. In tune with recent efforts to devise a fully local and relativistic measurement theory for quantum fields, we present a simple model of a local probe that is, itself, formulated in terms of a field theory. We then proceed to show how to systematically reduce this field-theoretic description of the probe to an effective theory restricted to a finite set of modes. The resulting dynamics at leading order in perturbation theory are given precisely by the widely adopted models for detectors based on nonrelativistic probe systems. These results pave the way to bridge the gap between the fully field-theoretic and the detector-based approach to measurements in QFT, and give particle detector models an effective field theory flavor.

Part II then focuses on the concept of *entanglement* in quantum field theory. We start in the first half of Part II by studying a protocol known as *entanglement harvesting*, which allows two localized probes to extract entanglement from a quantum field even before they have time to exchange causal signals. Recent works on field-theoretic models for local probes in relativistic quantum information have raised objections against the possibility of entanglement harvesting at weak coupling between the probes and the field when the probes themselves consist of localized degrees of freedom of a field theory. We address the origins of these concerns and show that, for an appropriate choice of modes used as probes for the quantum field, it is indeed possible to harvest entanglement using localized probes described along the lines of the formalism presented in Part I. In the second half of Part II, we also show how to best couple to a quantum field in order to most accurately reproduce its entanglement structure. This helps to establish limits on how efficient entanglement harvesting between complementary subregions can be, and also suggests further directions for how to address the problem of probing the structure of entanglement in field theory in more general scenarios.

Finally, in Part III we delve into the concept of *computational complexity* of Gaussian states, which are a special class of quantum states that is pervasive in many contexts in QFT. Following Nielsen’s geometric approach to circuit complexity, we devise a general class of metrics on the space of Gaussian unitary circuits which allows the circuit complexity of any pure Gaussian state to be characterized in a unified fashion. This gives us a generalized framework that can accommodate additional physical constraints on the notion of complexity adopted; we comment on a few examples where these additional features can be of physical relevance.

Acknowledgments

I could not start by thanking anyone but my wife, Amanda Barros. Amanda, words cannot describe how grateful I am for the privilege of sharing my life with you. Nothing that I did in these past 5 years would have been possible without all your love, support, and patience; and I feel that, in a very real sense, getting to this point in my PhD has been as much your merit as it has been mine.

I should also thank my supervisors, Eduardo Martín-Martínez and Beni Yoshida. Edu, your enthusiasm and passion for physics will never cease to amaze me, and your deep sense of responsibility and care for your students is something that I will always admire. Beni, even though we did not manage to write a paper together during my PhD, I have always looked up to you as an example of how to do theoretical physics—to be curious, think deeply, and not be afraid of tackling the hard questions. You both never failed to provide me with all the support and assistance that I needed to thrive, while also giving me full independence to explore my own interests. It has truly been a privilege to have both of you as mentors during my PhD, and I feel that academia would be a better place if more supervisors were like the two of you.

I have also been fortunate to meet several people that made my time as a PhD student all the more enjoyable. Thanks to José Polo-Gómez, Boris Ragula, Tales Rick Perche, Adam Teixidó-Bonfill, Evan Peters, Erickson Tjoa, Kelly Wurtz, and Matheus Zambianco, who helped make Waterloo feel more like home; to Achim Kempf, who, despite not having me as a student, was kind enough to give me access to the Physics of Information Lab—probably the best place on the University of Waterloo campus to work on theoretical physics; to all members of Beni’s journal club on Quantum Information and Quantum Gravity at Perimeter Institute (in particular Charlie Cummings, Dongjin Lee, Shuwei Liu, Elisa Tabor, and Jessica Yeh), from whom I learned so much about quantum information and holography throughout the years; and finally, to Alexander Jahn and Marios Christodoulou, who hosted me during two very enriching visits to Freie Universität Berlin and IQOQI Vienna in the last year of my PhD.

Finally, I thank the financial support from the Institute for Quantum Computing through the Mike & Ophelia Lazaridis Fellowship, which was essential for me to have enough time to dedicate what motivated me the most.

Dedication

To my grandmother, *vovó* Zezé—an endless source of wisdom and inspiration, and an example of a life well lived;
and to my son Theo, for the bright future that lies ahead.

Table of Contents

Examining Committee Membership	ii
Author's Declaration	iii
Statement of Contributions	iv
Abstract	vi
Acknowledgments	vii
Dedication	viii
List of Figures	xiii
Preamble	xvii
I Particle detector models	1
1 Introduction	2
2 Background	7
2.1 Particle detectors: general lore and UDW model	7
2.2 Causality and covariance in particle detector models	12

3	From quantum fields to particle detectors	17
3.1	Localized quantum fields	17
3.2	From probe field to particle detectors via the path integral	21
3.2.1	Setting up the path integral	21
3.2.2	Reduction of probe field to a finite set of modes	28
3.3	Generalizations to multiple trajectories	33
3.3.1	General strategy	33
3.3.2	Examples	36
3.4	Concluding remarks	39
II	Entanglement in QFT	41
4	Motivation: why entanglement in QFT is hard	42
4.1	Entanglement in quantum mechanics	42
4.2	What makes QFTs different	46
5	Local particle detectors and entanglement harvesting	53
5.1	Review of entanglement harvesting	53
5.2	Can local particle detectors really harvest entanglement?	58
5.3	Estimating realistic levels of probe mixedness	64
5.3.1	Minimally mixed local field modes	64
5.3.2	Realistic levels of mixedness	72
5.4	Entanglement harvesting with modes of localized quantum fields	81
5.4.1	Typical mixedness of modes in localized quantum fields	81
5.4.2	Results for entanglement harvesting	85

6	Optimal entanglement extraction from a QFT	90
6.1	Extracting modes of the field with particle detectors	91
6.2	Entanglement in the ground state of a harmonic chain	93
6.3	Upper bound to entanglement extraction with a finite set of modes	98
6.4	Entanglement area laws from finite set of modes	103
6.5	The continuum limit	105
6.6	Summary	114
III	Complexity geometry of Gaussian states	119
7	Introduction	120
8	Review of complexity geometry and Gaussian states	124
8.1	Nielsen’s geometric approach to circuit complexity	124
8.2	Parametrization of pure Gaussian states	128
9	Generalized Gaussian complexity	135
9.1	General features of Gaussian state complexity geometry	135
9.2	A non-reversible measure of complexity	141
9.2.1	Formalism for pure states	141
9.2.2	Generalizations to mixed states	143
9.3	Examples	144
9.3.1	Right-invariant metric on the group manifold	144
9.3.2	Complexity of coherent states	146
9.3.3	Weyl transformation of right-invariant metric	147
9.3.4	Non-reversible complexity for a single bosonic mode	150

IV	Conclusions	153
10	Summary and outlook	154
	References	161
	APPENDICES	177
A	Appendices for Chapter 3	178
A.1	Computation of the Jacobian determinant $\mathcal{N}$ in Eq. (3.34)	178
A.2	Normalization (3.10) and canonical commutation relations	180
B	Appendix for Chapter 6	183
B.1	Proof of Equation (6.30) and interlacing theorem for symplectic eigenvalues	183

List of Figures

5.1	Schematic view of a possible configuration of the coupling zones K_A , K_B and processing regions N_A and N_B in a spacetime diagram.	59
5.2	The power spectrum, $ \tilde{z}_{m,\kappa}(k) ^2$, given by (5.54) is shown for several values of m and κ . The peaks are paired mirror-wise. From inside to outside we have $(m = 3, \kappa = 5\pi)$, $(m = 4, \kappa = 15\pi)$, $(m = 5, \kappa = 25\pi)$, $(m = 6, \kappa = 35\pi)$, and $(m = 7, \kappa = 45\pi)$ with $\ell = 1$ throughout. As κ increases the peaks are displaced farther from the origin. As m is increased the polynomial tails of the distributions are suppressed. Away from the peaks $ \tilde{z}_{m,\kappa}(k) ^2$ decays polynomially as $ k ^{4-2m}$. Increasing m also widens the peaks. The peak width is proportional to m	70
5.3	The profiles of the local modes considered in this chapter. Specifically, profiles 1, 2, and 3 show Equation (5.62) with $m = 1, 2, 3$. Likewise, profiles 4, 5, and 6 show Equation (5.63) with $m = 4, 5, 6$. Unlike the previous profiles, the final three profiles are 3 dimensional. Specifically, profiles 7, 8, and 9 show Equation (5.64) with $m = 2, 3, 4$	74
5.4	The purity of select local modes as we vary the size of their support (vertical axis) and the ambient temperature of the field (horizontal axis). The top row of subfigures has an IR regulator introduced by means of a field mass M to which the support size ℓ and ambient temperature T are compared. The bottom row of subfigures has an IR regulator introduced by means of cavity walls a distance L apart from each other. The columns show the purity of the local mode given by profiles 1, 2, and 3 in Figure 5.3.	75

5.5	The purity of select local modes as we vary the size of their support (vertical axis) and the ambient temperature of the field (horizontal axis). The top row of subfigures has an IR regulator introduced by means of a field mass M to which the support size ℓ and ambient temperature T are compared. The bottom row of subfigures has an IR regulator introduced by means of cavity walls a distance L apart from each other. The columns show the purity of the local mode given by profiles 4, 5, and 6 in Figure 5.3.	77
5.6	The purity of select local modes as we vary the size of their support (vertical axis) and the ambient temperature of the field (horizontal axis). Each subfigure has an IR regulator introduced by means of a field mass M to which the support size ℓ and ambient temperature T are compared. The subfigures show the purity of the local mode given by profiles 7, 8, and 9 in Figure 5.3.	78
5.7	The purity of three different local modes is shown without an IR regulator. In subfigure a) the field is 1+1 dimensional and in b) the field is 3+1 dimensional. In figure a) the three lines show the purity of the local mode given by profiles 4, 5, and 6 in Figure 5.3. b) The three lines show the purity of the local mode given by profiles 7, 8, and 9 in Figure 5.3.	80
5.8	Symplectic eigenvalue of the mode defined with the spatial profile (5.81) in 1, 2, and 3 dimensions, as a function of the ratio σ/ℓ . For concreteness, the mass of the field was set so that $m\ell = 10$ in all cases.	86
5.9	The negativity of the state of two localized quantum fields confined by a quadratic potential when restricted to their lowest energy after interacting with a massless scalar field. The negativity is plotted as a function of the energy of the modes $\Omega = \sqrt{m^2 + 3/\ell^2}$. The time duration of the interaction T is used as a scale. The separation between the interaction regions of each detector for these plots is $L = 5T$	88
5.10	The negativity of the state of two localized quantum fields in boxes of sides d when restricted to their lowest energy mode after spacelike interaction with a massless scalar field. The negativity is plotted as a function of the energy of the localized mode, $\Omega = \sqrt{m^2 + 3\pi/d^2}$. The time duration of the interaction T is used as a scale. The separation between the interaction regions of each detector for these plots is $L = 4.5T$	88

6.1	Spatial profile of the most mixed normal mode of a given region (in red) and its complement (in blue) as given by the occupation function F_i . Subfigures a) and b) show this for a lattice in 1 and 2 spatial dimensions respectively. Subfigures c) and d) are both in 3D, with c) being a full three-dimensional visualization of the mode and its complement, and d) showing two-dimensional slices at different heights of the plot at the top.	100
6.2	Entanglement entropy between a subregion and its complement as a function of the ratio between the subregion's side length and the total side length of the lattice. Subfigures a), b), and c) show this in 1, 2, and 3 spatial dimensions, respectively. If we take the number of boundary sites to be N_B , then each colored line corresponds to the entanglement captured by the given fraction of N_B modes. The total entanglement entropy is saturated, and we observe an area-law behavior in each case, when we count the entanglement coming from N_B modes.	104
6.3	Entanglement entropy between one third of a 1D chain and its complement as a function of the total number of sites in the lattice, i.e. the inverse UV cutoff, as here we keep the length of the chain fixed. As the density of lattice sites increases, we see that the dependence of the entanglement entropy on the effective short-distance cutoff quickly converges to the behavior that is expected from other methods in field theory [133].	110
6.4	Dimensionless approximated position and momentum profiles, $\lambda_p g_N(x)$ and $\lambda_q f_N(x)$, of the most mixed mode of a region comprising the central third of a box, for a 1+1-dimensional massive quantum scalar field theory of mass $m = 1/L$ in the vacuum state. The profiles were calculated using Equations (6.73) and (6.74), using the spatial profiles obtained for the corresponding 1-dimensional lattice field theories with UV parameters $N = 500, 1000$ and 2000 , and setting the parameters ω and α via Equation (6.63) with $d = 1$. The second row shows an enlarged view of the profiles close to the left boundary.	111

6.5	Dimensionless approximated position and momentum profiles, $\lambda_p g_N(x)$ and $\lambda_q f_N(x)$, of the partner mode of the most mixed mode of a region comprising the central third of a box, for a 1+1-dimensional massive quantum scalar field theory of mass $m = 1/L$ in the vacuum state. The profiles were calculated using Equations (6.73) and (6.74), using the spatial profiles obtained for the corresponding 1-dimensional lattice field theories with UV parameters $N = 500, 1000$ and 2000 , and setting the parameters ω and α via Equation (6.63) with $d = 1$. Below the profiles, a zoom-in of the details of the profiles close to the left boundary.	113
6.6	Dimensionless approximated position and momentum profiles, $\lambda_p g_N(x)$ and $\lambda_q f_N(x)$, of the most mixed mode of a region comprising the central square of side length $L/3$ of a square of side length L , for a 2+1-dimensional massive quantum scalar field theory of mass $m = 1/L$ in the vacuum state. The profiles were calculated using Equations (6.73) and (6.74), using the spatial profiles obtained for the corresponding 2-dimensional lattice field theories with UV parameters $N = 15, 30$ and 60 , and setting the parameters ω and α via Equation (6.63) with $d = 2$	115
6.7	Dimensionless approximated position and momentum profiles of the partner mode of the mode represented in Fig. 6.6.	116
9.1	Pictorial representation of the geometrical ingredients involved in defining state complexity from a metric in the group manifold. After fixing the reference state (and thus the stabilizer subgroup $\text{Sta}(N)$ in G), one shoots off geodesics in the orthogonal directions to the generators of the stabilizer at the identity, thus constructing the submanifold $\exp(\mathfrak{sta}_\perp(N))$. The state complexity of the target state $ \psi_T\rangle$ given by the equivalence class of some Gaussian transformation $M \in G$ is then visualized as the length of the geodesic connecting $\mathbb{1}$ to the point where the equivalence class $[M]$ intersects $\exp(\mathfrak{sta}_\perp(N))$	140

Preamble

This thesis is divided into three parts, each of them addressing a different corner of the interface between quantum information and quantum field theory. The overarching principle connecting all three corners is the idea that understanding operationally motivated notions of quantum information in the context of QFT can be a powerful way of shedding light on how to think of quantum field theory on a more foundational level.

In **Part I**, we show how it is possible to draw a connection between localized quantum field theories and more commonly used models of local probes in relativistic quantum information (RQI). By integrating over and then tracing out the inaccessible modes of the localized field being used as a probe, we show that, at leading order in perturbation theory, the dynamics of any finite number of modes of the probe field is exactly that of a finite number of harmonic-oscillator Unruh-DeWitt (UDW) detectors. The equivalence is valid for a rather general class of input states of the probe-target field system, as well as for any arbitrary number of modes included as detectors. Our approach will also provide a closed-form expression which gives us a systematic way of obtaining the corrections to the UDW model at higher orders in perturbation theory due to the existence of the additional modes that have been traced out. This vindicates and extends a recently proposed bridge between detector-based and field-theory-based measurement frameworks for quantum field theory, and the technical tools used (most importantly, the Schwinger-Keldysh path integral) point to promising connections between the emergence of effective particle detector models in RQI and other areas of physics where path integral methods are more commonplace—in particular, the Wilsonian approach to the renormalization group and effective field theories.

In **Part II**, we study a very well-known protocol in RQI known as *entanglement harvesting*. This is a protocol by which two localized probes that couple to a quantum field in local regions of spacetime can get entangled with each other, even if they start in a separable state and their respective coupling regions are spacelike separated. Partially building on the ideas of Part I, we study entanglement harvesting when the probes that harvest entanglement from the field emerge from fully relativistic quantum field theories

themselves. We will show that entanglement harvesting is indeed possible with probes made of modes of relativistic fields, despite earlier claims in the literature that suggested otherwise. In the second half of Part II, we also comment on how to optimize the extraction of entanglement from a quantum field, by studying how entanglement between complementary regions of a QFT (or more specifically, an approximation of the QFT as a harmonic chain) can be distributed into modes associated with each respective region. Those results point to the optimal choice of coupling that detectors must aim to implement, in order to best reflect in their final state the entanglement structure originally present in the field.

Finally, in Part III we consider the problem of *computational complexity*, which has also gained a lot of attention in recent years due to its many implications in several areas of theoretical physics, ranging from quantum information to many-body physics, quantum thermodynamics, and even quantum gravity. Motivated by the desire to have a general framework for complexity in QFT, we study the notion of circuit complexity for Gaussian states (both bosonic and fermionic), building on the formalism devised by Michael Nielsen where the circuit complexity is understood in terms of Riemannian geometry. We lay out general conditions that a Riemannian metric function on the space of Gaussian states should satisfy in order for it to yield a physically reasonable measure of complexity, and show how previous analytical results on the characterization of circuit complexity can actually be extended to a much more general class of complexity geometries under appropriate assumptions. This general formalism can naturally accommodate modifications to complexity geometries that arise from cost functions that depend nontrivially on the instantaneous state and the direction in circuit space at each point. We explore these modifications and, as a particular case, show how to account for time-reversal symmetry breaking in measures of complexity, which is often natural from an experimental (and thermodynamical) perspective, but is absent in commonly studied complexity measures. Our goal is to establish a first step towards building a quantitative, geometric notion of complexity that faithfully mimics what is experienced as “easy” or “hard” to implement in a lab from a physically motivated point of view.

Notation and conventions

Spacetime is given by a pair $(\mathcal{M}, g_{ab})$, where $\mathcal{M}$ is a $(d + 1)$ -dimensional differentiable manifold and g_{ab} is a Lorentzian metric on $\mathcal{M}$. For simplicity, it will always be assumed that the background spacetime in question is globally hyperbolic. The signature convention for the metric is such that $g_{ab}T^aT^b < 0$ if T^a is a timelike vector. We will represent abstract points in $\mathcal{M}$ with sans-serif font x , and reserve the normal math font x for the collection

of spacetime coordinates associated with the point $\mathbf{x}$ in a given coordinate system. We will denote by ∇_a the Levi-Civita connection, defined as the unique torsion-free covariant derivative that is compatible with the metric g_{ab} . dV is the volume form associated to g_{ab} , which in any coordinate system takes the form $dV = \sqrt{-g} d^{d+1}x$, where $g \equiv \det(g_{\mu\nu})$. Quantum operators acting on a Hilbert space will always be written with hats, to clearly distinguish them from their classical counterparts: for instance, $\hat{\phi}_D(t, \mathbf{x})$ in Eq. (3.18) is the operator (or, if you want to be more precise, the operator-valued distribution) corresponding to the quantized version of $\phi_D(t, \mathbf{x})$ —which, when written as in Eq. (3.11), for example, is just a classical field configuration. Finally, we adopt natural units with $\hbar = c = 1$, except when clearly stated otherwise.

Part I

Particle detector models

Chapter 1

Introduction

Our modern understanding of the elementary building blocks of the physical world is ultimately based on the concept of *fields*. General relativity—currently our most successful theory of gravity—is fundamentally a theory about the dynamics of the gravitational field and how it manifests as the geometry of spacetime. Similarly, the Standard Model of particle physics—our most complete description of the fundamental constituents of matter and all elementary forces of nature (excluding gravity)—is entirely formulated in the language of quantum field theory (QFT), where the concept of a quantum field is essential to combine the principles of quantum mechanics and special relativity into a single, consistent framework. This appears to be one of the most profound lessons of modern theoretical physics.

It also appears to be a deep truth that the only way we can acquire information about fields is through some form of interaction between the field of interest and an auxiliary system that can couple to it in localized regions of space and time. We learn about the structure of the gravitational field by studying how it affects the trajectories of (approximately) pointlike test particles; we learn about the electromagnetic field by tracking the dynamics of localized charges and currents; and although all the elementary constituents of matter ultimately emerge from relativistic quantum fields themselves, our knowledge about their most fundamental properties actually comes from the localized excitations detected by the countless sensors that make up particle accelerators such as the LHC—which, directly or indirectly, carries information about the intricate way those fields interact. It should therefore come as no surprise that a consistent framework for how to extract and process information from quantum fields will in general be intimately tied to how local probes can interact with quantum fields.

This general principle is embodied most concretely in the context of Relativistic Quantum Information (RQI). This is an area of research—mostly focused on exploring information-theoretic features of quantum field theory and gravity, as well as studying the role that relativity may play in information-processing tasks—where local probes (usually referred to as *particle detectors*) play a central role. Particle detectors in RQI consist of idealized versions of localized quantum systems that can controllably couple to quantum fields in local regions of spacetime. From a theoretical perspective, these have proven to be very useful for investigations on several aspects of the interplay between quantum information, quantum field theory, and gravity: applications range from a measurement framework for QFT which can be made to respect the underlying principles of relativistic locality and causality [143], operationally motivated protocols for quantifying and extracting entanglement from QFTs [150, 151, 144, 59], and detector-based approaches to important phenomena at the foundations of QFT in both flat and curved spacetimes, such as the Unruh effect and Hawking radiation [182, 44, 61]. From a more practical point of view, particle detectors also capture important features of physical setups of great experimental significance, such as the dynamics of atomic probes coupled to the electromagnetic field in atomic physics and quantum optics [116]. They are, therefore, very well-adapted as theoretical tools for the study of relativistic information-processing tasks that can (at least in principle) actually be performed in a lab.

For practical purposes, it is common to assume that the internal dynamics of the system being used as a detector can be approximated as non-relativistic. This is a natural assumption to make when one thinks of archetypical examples of particle detectors given, for instance, by atomic probes coupled to external electromagnetic fields: after all, in many regimes of physical interest, the internal structure of the atom can be well-approximated by the non-relativistic Schrodinger equation applied to the electrons under the influence of the atomic nucleus.

Approximating the probe as an internally non-relativistic system is also advantageous from a purely theoretical point of view. It is well-known that building a framework for local measurements entirely within relativistic QFT is highly nontrivial, since most immediate versions of measurement postulates imported from non-relativistic quantum mechanics often lead to irreconcilable conflicts with relativistic locality and causality [165, 64, 20]. In contrast, if we can approximate the probe’s internal dynamics as non-relativistic, we are justified in applying the standard measurement framework from non-relativistic quantum mechanics to the physical system being used as a detector. This dramatically alleviates the burden of directly measuring a relativistic field, by replacing this problem with the (arguably easier) problem of coupling the field to a set of detectors, and then then indirectly inferring properties of the target field from results of measurements on the detectors.

However, there is clearly something slightly unsatisfying about ending the story here. After all, our most fundamental theory for the structure of matter teaches us that the internal constituents of any physical system—including the probe itself—fundamentally emerge from quantum fields. This has led to efforts aimed at showing how to model measurements in quantum field theory using fully relativistic probes.

Progress in this direction is spearheaded by the Fewster-Verch (FV) framework [72, 73], which is an approach to the measurement problem in QFT where the probe system is also treated as a fully relativistic field theory, formulated in the language of algebraic QFT in general (possibly curved) spacetimes. The FV framework has been remarkably successful at providing a mathematically rigorous measurement scheme for QFT that is fully compatible with the relativistic nature of the theory. In particular, the formalism provides an elegant solution to long-standing problems regarding how to make sense of local measurements on quantum fields while taking relativistic locality and causality into account [21, 71].¹ However, since the probe is also a quantum field in this case, the FV framework is agnostic to how an external agent (say, an experimentalist taking notes on their notepad) can in practice extract information from their local measurements. Strictly speaking, the framework only describes the chain of information flow from one relativistic quantum field (the target field we are actually interested in measuring) to another relativistic quantum field (the probe system being used as a detector). It is, therefore, less directly applicable to more realistic measurement settings, especially when compared to the detector-based framework of [143].

Our goal in this part of the thesis will be to show how to connect a simple, fully relativistic model of a local probe—namely, a scalar field that is effectively confined to a localized region of space by an external potential—to a system of local probes that is closer in spirit to the usual particle detector models adopted in RQI. This is an attempt towards a more comprehensive understanding of the relation between particle detector models and field-theoretic descriptions of local probes, with the aim of narrowing the gap between the FV framework and the detector-based approaches to RQI.²

The formulation of fully covariant and field-theoretic versions of particle detector models commonly used in RQI has been the subject of active research in recent years [154, 80, 137]. Our work here will then provide yet another example where the comparison between the fully relativistic and non-relativistic versions of probes can be studied explicitly, and illustrate

¹We should emphasize that the detector-based approach of [143] is also capable of addressing these problems; this was, in fact, part of the motivation for its inception. However, the approximation of the probe as a non-relativistic system comes at a price. For more details, see e.g. [57]—or wait until Subsection 2.2.

²On a related note, we also draw attention to recent accounts [84, 134] which thoroughly discuss how particle detector models and the FV framework both fit within the bigger picture of measurements in QFT, seen from the point of view of foundations and philosophy of physics.

the use of new methods which we believe can be useful in systematically showing how the latter can emerge from the former.

The main technical tool used to derive our results will be the *Schwinger-Keldysh path integral* [158, 108, 74]. This can be roughly thought of as an upgraded version of the Feynman path integral that is better suited for describing the non-equilibrium dynamics of a system interacting with an inaccessible environment. In this case, the role of the “environment” is played, in a loose sense, by a set of degrees of freedom of the probe system that are deemed inaccessible, and are therefore traced out in the final state of the system of interest. The result is an effective description of the detector-field system that only contains finitely many degrees of freedom of the probe, in a way that resembles what would be obtained if one started with a finite number of detectors coupled to the field. Due to the simplicity of the model for the probe, the effect of integrating over the inaccessible degrees of freedom can be worked out in full detail; this provides a simple setting in which to compare predictions from particle detector models and more field-theoretic descriptions for probe systems.

The Schwinger-Keldysh path integral has already been suggested as an important tool for the development of field-theory-based formulations of RQI [9, 8]. The way it is being used here, however, highlights its promising role as a method of explicitly connecting fully relativistic probes to effectively non-relativistic particle detectors, which is something that has not been proposed before (to the best of the author’s knowledge at the time of writing). This also paves the way for further investigations on the interplay between measurement schemes in RQI and techniques from effective field theory [142, 37]—which, on a conceptual level, is clearly the correct framework in which to address the apparent tension between fundamentally relativistic dynamics of the probes and non-relativistic descriptions of measurements in experimentally relevant scales.

The remainder of Part I is organized as follows:

- In Chapter 2 we review the basics of particle detector models that are most commonly used in RQI, with a particular focus on the Unruh-DeWitt (UDW) model. We then briefly discuss some of the shortcomings of these models when taking into account what one would expect from a fully relativistic theory.
- In Chapter 3 we describe a simple concrete model of a fully relativistic probe, consisting of a scalar field that is effectively localized by some external potential. Despite its simplicity, the model is already capable of bypassing the main limitations of the usual UDW model that were explained in Chapter 2.

- We then move on to show how to describe the dynamics of the joint field-detector system via the Schwinger-Keldysh path integral. By explicitly integrating over and tracing out a set of inaccessible degrees of freedom of the probe, we see that the effective dynamics of the remaining modes are, at leading order in perturbation theory, precisely the same as what would be obtained in the UDW model; at higher orders, the path integral also provides a closed-form expression which can be used to systematically calculate the deviations between the two models. This provides a rather simple illustration of a model that can quantitatively interpolate between a fully field-theoretic framework for local probes, and a simpler model that is more similar to what one might consider in effectively non-relativistic systems.
- We close by commenting on how to extend the framework to encompass cases where the same underlying quantum field can lead to effective detectors in multiple trajectories, and mention a few scenarios (one of which will re-appear later in the thesis) where this extension may be of physical interest.

Chapter 2

Background

In this Chapter, we review some of the general aspects of particle detector models, with the Unruh-DeWitt (UDW) model being introduced as a paradigmatic example. We also discuss some of the drawbacks and limitations of these models—especially pertaining to considerations about covariance and causality—which have been used to advocate for the necessity of a fully relativistic, field-theoretic version of local probes for quantum fields.

2.1 Particle detectors: general lore and UDW model

The general philosophy when using particle detectors as probes for quantum fields is to extract information from the quantum field of interest by coupling a detector to it, and studying how the evolution of the state of the detector indirectly depends on features of the quantum field being probed. In very general terms, the dynamics of the joint detector-field system can be described through an action of the form

$$S = S_{\text{F}} + S_{\text{D}} + S_{\text{I}}, \quad (2.1)$$

where S_{F} and S_{D} provide the free dynamics of field and detector respectively, and S_{I} encodes the coupling between the two systems.

The action of the field is typically expressed as an integral of a Lagrangian over spacetime,

$$S_{\text{F}} = \int dV \mathcal{L}_{\text{F}}(\psi^A, \nabla_a \psi^A, \dots), \quad (2.2)$$

where the Lagrangian $\mathcal{L}_F$ is a scalar function of the dynamical fields ψ^A and its covariant derivatives, and the index A encodes any possible collection of tensor/spinor indices for the dynamical fields in question. $\mathcal{L}_F$ also generally depends on some external background fields (such as the spacetime metric g_{ab} and any additional external potentials that are treated as non-dynamical), and may also include some explicit dependence on higher derivatives of the metric through terms such as non-minimal coupling to curvature.

In its simplest version, the detector is pictured as a localized system following a fixed classical trajectory $\mathbf{z}(\tau)$ on $\mathcal{M}$, where τ has the interpretation of the detector's proper time. The system is also endowed with a set of quantum internal degrees of freedom q^i , whose dynamics can be effectively prescribed by an action that takes the form of an integral over proper time of some Lagrangian,

$$S_D = \int d\tau \mathcal{L}_D(q^i, \frac{dq^i}{d\tau}, \dots). \quad (2.3)$$

Similar to $\mathcal{L}_F$, the detector's Lagrangian $\mathcal{L}_D$ is a function of the system's internal degrees of freedom and its derivatives, and it will also generically depend on features of the extrinsic geometry of how $\mathbf{z}(\tau)$ is embedded in $\mathcal{M}$ (for instance, the proper acceleration of the detector's trajectory), as well as other geometrical aspects of the background spacetime (for instance, the Riemann curvature tensor of the metric g_{ab} evaluated along $\mathbf{z}(\tau)$). There can also be extensions of the formalism where the detector is put to evolve in a *superposition* of trajectories [76, 12]; this can be achieved dynamically, for instance, by letting the detector's center of mass also be a quantum degree of freedom [168, 78]. For the most part, the simpler setting where the detector's position is treated as classical will be enough for our purposes in this thesis. However, we shall come back to briefly comment on a similar extension to detectors associated with more than one classical trajectory in Section 3.3.

Finally, the coupling between detector and field can be expressed through an interaction action of the form

$$S_I = \lambda \int dV \Lambda(\mathbf{x}) \mu_A(\tau(\mathbf{x})) \mathcal{O}^A(\mathbf{x}), \quad (2.4)$$

where $\mu_A(\tau)$ represents an observable of the detector, $\mathcal{O}^A(\mathbf{x})$ is an observable of the field, $\Lambda(\mathbf{x})$ is a spacetime smearing function that dictates the strength of the interaction between probe and field in space and time, and λ is an overall coupling constant.¹ The spacetime smearing function $\Lambda(\mathbf{x})$ should be pictured as most strongly supported around the detector's trajectory

¹When writing Equation (2.4), we are implicitly assuming that $\Lambda(\mathbf{x})$, $\mu_A(\tau)$, and $\mathcal{O}^A(\mathbf{x})$ are all real, and that upon quantization both $\hat{\mu}_A(\tau)$ and $\hat{\mathcal{O}}^A(\mathbf{x})$ are Hermitian observables on the respective Hilbert spaces of detector and field. An obvious—and most importantly, sometimes physically motivated [145, 179, 138]—generalization of this setup would allow for $\Lambda(\mathbf{x})$ to be complex, and $\hat{\mu}_A(\tau)$ and $\hat{\mathcal{O}}^A(\mathbf{x})$ to be non-Hermitian;

$z(\tau)$. In Equation (2.4), $\tau(x)$ corresponds to the Fermi normal coordinate time relative to the trajectory $z(\tau)$, which extends the proper time parameter τ (originally only defined for points along the detector’s trajectory) to a timelike coordinate that can be assigned to any point x in a sufficiently small neighborhood of $z(\tau)$ [140]. Once again, the index A in the field observable $\mathcal{O}^A(x)$ can comprise any possible collection of tensorial/spinorial indices, with the detector’s observable $\mu_A(\tau)$ then being an element of the dual space to $\mathcal{O}^A$ in order for the contraction $\mu_A \mathcal{O}^A$ to form a Lorentz scalar.

If one thinks about the physical origins of the smearing function of a detector, two main mechanisms may come to mind. The first is the existence of some internal structure of the detector, which gives it some small but finite spatial extension. The second is the limitation on one’s ability to exactly pinpoint the location where the coupling between probe and field takes place, due to some uncertainty (either classical or quantum) in the position of the system being used as a probe. In practice, when applying particle detector models with the flavor described above to concrete physical scenarios, it is more usual to think of the smearing as coming from the former mechanism: that is, the smearing often encodes some inherent spatial extension of the probe. In the case of atomic physics and quantum optics, for instance, the size of the probe would roughly correspond to the typical radius of the atom or ion in an optical trap. Unlike elementary particles (such as electrons), no one would claim that an ion in an optical trap is truly pointlike; therefore, the smearing of the probe in this case would physically correspond to the first aforementioned mechanism, and not the second. A similar reasoning would apply to most physical implementations of particle detectors that one would reasonably imagine being implemented in laboratory settings.

There is a sense, however, in which the difference between these two mechanisms is actually not always sharply defined. After all, the finite size of the probe can itself be interpreted as also coming from what we described above as “quantum uncertainty”. This is because the “size” of an atom—which, in the case of a typical hydrogen atom, is of the order of 10^{-8} centimeters—ultimately stems from the characteristic length scale determining the region of appreciable support to the electron’s wavefunction within the atom, which is quantum in nature. Therefore, even though thinking in terms of an effective size as a true spatial extension of the probe can be a simple and useful way of modeling how the smearing comes about, it may very well be the case that this size is also a consequence of fundamentally quantum uncertainty, albeit in a much smaller scale than what we are interested/capable of directly probing. Of course, this should not come as a surprise, given that our modern understanding is that all matter emerges from quantum fields whose

in this case, the interaction action between detector and field would contain a sum of terms which include the right-hand side of (2.4) and its complex conjugate.

interactions are strictly local in spacetime.²

Let us now go back to the generalities about how to model the dynamics of a (generally smeared) probe with a quantum field through the interaction from Equation (2.4). Given a coordinate system $(t, \mathbf{x})$ which foliates $\mathcal{M}$ by a family of Cauchy surfaces $\mathcal{E}_t$ labeled by constant values of the timelike coordinate t , the quantum dynamics of the joint system of detector and field interacting through the action (2.4) can also be described in terms of an interaction Hamiltonian

$$\begin{aligned}\hat{H}_I(t) &= -\lambda \int_{\mathcal{E}_t} d^d \mathbf{x} \sqrt{-g} \Lambda(\mathbf{x}) \hat{\mu}_A(\tau(\mathbf{x})) \hat{\mathcal{O}}^A(\mathbf{x}) \\ &= \int_{\mathcal{E}_t} d^d \mathbf{x} \sqrt{-g} \hat{h}_I(\mathbf{x}),\end{aligned}\tag{2.5}$$

where we have conveniently defined the Hamiltonian scalar density

$$\hat{h}_I(\mathbf{x}) = -\lambda \Lambda(\mathbf{x}) \hat{\mu}_A(\tau(\mathbf{x})) \hat{\mathcal{O}}^A(\mathbf{x}).\tag{2.6}$$

This is the most common starting point in concrete applications of particle detectors in RQI. In this context, it is also customary to adopt the interaction picture, in which case the operators $\hat{\mu}_A(\tau(\mathbf{x}))$ and $\hat{\mathcal{O}}^A(\mathbf{x})$ appearing in (2.5) should be understood as the resulting time evolution implied by the free actions S_D and S_F , respectively. In the interaction picture, the joint state of the detector-field system then evolves according to the unitary time evolution operator

$$\hat{\mathcal{U}}_I = \mathcal{T}_t \exp \left(-i \int dt \hat{H}_I(t) \right),\tag{2.7}$$

where $\mathcal{T}_t \exp$ denotes the time-ordering exponential with respect to the time parameter t .³ It is natural to take the initial state $\hat{\rho}_0$ of the system to be a product state between detector and field,

$$\hat{\rho}_0 = \hat{\rho}_F \otimes \hat{\rho}_{D,0},\tag{2.8}$$

with $\hat{\rho}_F$ and $\hat{\rho}_{D,0}$ being the initial states for the field and the detector, respectively. With that, the final state of the detector can be expressed as

$$\hat{\rho}_D = \text{Tr}_F \left[\hat{\mathcal{U}}_I (\hat{\rho}_F \otimes \hat{\rho}_{D,0}) \hat{\mathcal{U}}_I^\dagger \right],\tag{2.9}$$

²That is, unless we believe that the elementary constituents of matter are themselves described by extended objects such as strings [141]—but let us not go there!

³The reason why we are emphasizing the dependence on the time coordinate used to perform the time-ordering in Equation (2.7) will be made clear in Section 2.2.

where Tr_F denotes the partial trace over the Hilbert space associated with the field. By then judiciously engineering the interaction action (2.4) and analyzing the final state of the detector (2.9), it is possible to indirectly obtain information about features of the field. In a nutshell, this is how particle detectors can be concretely used as probes for quantum field theories.

The most commonly explored version of this otherwise fairly general setting for the interaction between localized probes and quantum fields is the so-called Unruh-DeWitt (UDW) model [182, 61]. The model consists of a localized quantum system interacting with a free, real scalar field, with the coupling to the detector taken to be linear in the field amplitude. The free action of the field is therefore

$$S_F = -\frac{1}{2} \int dV (g^{ab} \nabla_a \psi \nabla_b \psi + M^2 \psi^2 + \xi R \psi^2), \quad (2.10)$$

and the interaction between field and detector is most commonly given by

$$S_I = \lambda \int dV \Lambda(x) \mu(\tau(x)) \psi(x). \quad (2.11)$$

The action (2.10) describes a Klein-Gordon field $\psi(x)$ with mass M and possibly non-minimal coupling to curvature, where R is the Ricci scalar of the background metric and ξ is a dimensionless constant. As for the detector, a number of variants for its internal degrees of freedom have been explored, with different types of probes being referred to as UDW (or UDW-like) detectors. In [182] Unruh considered two models, the first consisting of a free particle in a box and the second actually built with a quantum field confined in a cavity with Dirichlet boundary conditions.⁴ Not long after that, Unruh and Wald [184] considered a simpler version of the setup with a particle in a box, where the relevant dynamics of the probe are restricted to just two of its internal energy levels. In more modern parlance, the term “UDW detector” without further specifications is most commonly used to describe a local probe given by a two-level system (i.e., a qubit)—see. e.g., [117, 118, 55, 122, 160, 175, 54, 36]. For the purposes of this thesis, however, the most relevant variant will be the harmonic oscillator UDW detector [183, 101, 114, 32, 31], which describes a probe with an internal degree of freedom $q(\tau)$ whose action is given by

$$S_D = \frac{1}{2} \int d\tau \left[\left(\frac{dq}{d\tau} \right)^2 - \omega^2 q^2 \right], \quad (2.12)$$

⁴Note that the second model considered by Unruh in [182] is very close in spirit to the fully relativistic model that we will describe in the next chapter.

and we take the observable $\mu(\tau)$ that couples to the field in (2.11) to simply be $q(\tau)$, so that we have

$$S_I = \lambda \int dV \Lambda(\mathbf{x}) q(\tau(\mathbf{x})) \phi(\mathbf{x}). \quad (2.13)$$

In (2.12), the parameter ω corresponds to the characteristic frequency of the harmonic oscillator, which is also the energy gap between any two consecutive energy eigenstates of the detector's free Hamiltonian in its proper frame. Note that we have absorbed a spurious mass in the Lagrangian of the harmonic oscillator in (2.12)—which would most commonly read $\mathcal{L}_D = \frac{1}{2}(m(dQ/d\tau)^2 + m\omega^2 Q^2)$ —by redefining $q = \sqrt{m}Q$; this gives q a normalization which matches that of a scalar field in $d + 1$ dimensions, particularized to the case $d = 0$ (as is well-known, quantum mechanics can equivalently be seen as a $(0 + 1)$ -dimensional field theory).

The main appeal of the harmonic oscillator UDW model (as opposed, say, to the more standard qubit detector) resides in the fact that the full action of the system in this case is quadratic in all dynamical variables, and therefore the detector-field dynamics is *Gaussian*. The assumption of Gaussianity vastly simplifies the description of the time evolution of the detector-field system, thanks to the plethora of powerful tools and results (collectively referred to as *Gaussian methods* or *Gaussian quantum mechanics*) which apply to systems undergoing Gaussian dynamics [25, 192, 2, 90]. Indeed, this is one of the few known cases where the time evolution of particle detectors interacting with quantum fields can be studied nonperturbatively [32, 31, 189].⁵ As we will see shortly, the harmonic oscillator UDW detector is also the archetypical example of a detector that naturally emerges from the simplest model of a probe system that we can formulate based on a fully local quantum field theory. This is therefore the model of reference that we will make a connection to when exploring the relation between localized quantum field theories and particle detector models later on.

2.2 Causality and covariance in particle detector models

In Subsection 2.1, we made an effort to introduce the basic elements of general particle detector models in terms of action functionals and Lagrangians. This approach has a number of advantages when compared to the slightly more usual presentations in RQI, which tend to introduce detectors at the level of Hamiltonians. First, in this formulation, the framework is more immediately adapted to general relativistic settings: the setup

⁵In fact, properties of Gaussian states will be a recurring theme, in various levels of sophistication, throughout this whole thesis.

from Eqs. (2.1)-(2.4) applies unchanged to detectors in arbitrary trajectories in general spacetimes, and its generalization to the case of multiple detectors in arbitrary relative states of motion is straightforwardly obtained by simply adding the free dynamics and interaction terms of each detector at the level of the total action.⁶ In that same vein, results such as the transformation properties of the interaction Hamiltonian (2.5) under general changes of coordinates [125, 123] can be immediately derived once the interaction Hamiltonian is treated as a byproduct of the interaction action (2.4)—which is itself written in a manifestly covariant way as a spacetime integral of a scalar. Lastly, talking in terms of actions and Lagrangians also allows for a more direct connection with relativistic field theory, where describing the dynamics of interacting fields at the level of Lagrangians (as opposed to Hamiltonians) is often more natural, and generally preferred if one wishes to keep the symmetries of the system manifest.

There is, however, an important subtlety regarding the particle detector models of the kind described in Subsection 2.1 which makes them, in a strict sense, not consistent as fully relativistic theories. The problem stems from the fact that, for generic spacetime smearings $\Lambda(\mathbf{x})$ supported in the neighborhood of the detector’s trajectory $\mathbf{z}(\tau)$, the coupling of a particle detector to a quantum field according to the interaction (2.4) actually violates microcausality—that is, the interaction Hamiltonian density $\hat{h}_I(\mathbf{x})$ prescribed by Equation (2.6) is such that $[\hat{h}_I(\mathbf{x}), \hat{h}_I(\mathbf{y})] \neq 0$ even when the spacetime points $\mathbf{x}$ and $\mathbf{y}$ are spacelike-separated [124]. The condition of microcausality—i.e., the property that any two local operators assigned to spacelike-separated points of spacetime must commute—guarantees that spacelike-separated events cannot influence each other, and is one of the basic requirements that a theory that respects a relativistic causal structure should satisfy.⁷ The fact that microcausality is *not* satisfied by the coupling derived from the interaction Hamiltonian (2.6) therefore means that, strictly speaking, the UDW model generally violates relativistic causality.

The Hamiltonian density (2.6) not commuting with itself at spacelike separation also impacts the properties of the time evolution operator (2.7) under general changes of coordinates. Given two Cauchy surfaces Σ_1 and Σ_2 (representing some “initial” and “final” times), there are infinitely many choices of timelike coordinate t which can be used to foliate the spacetime region between Σ_1 and Σ_2 by a family of Cauchy surfaces labeled by constant

⁶Of course, this can also be done at the Hamiltonian level; however, in that case, one has to be mindful of the fact that, whenever the observers in question do not share a common frame, the Hamiltonians associated with different observers will generate time evolution with respect to distinct notions of “time” [125, 123].

⁷Note that this does not preclude spacelike points from being *correlated*; in fact, it is well-known that the vacuum state of any relativistic field theory is highly entangled across spatial bipartitions [169, 197], and we will explore this fact in depth in Part II of this thesis. However, entanglement alone is not enough for communication or transmission of causal signals. As the maxim goes, correlation does not imply causation.

values of t . Different choices of timelike coordinate will in general assign different temporal orders between spacelike separated events, which could in principle lead to different time evolution operators from Σ_1 to Σ_2 since Equation (2.7) explicitly depends on time-ordering. Fortunately, this ambiguity in the time ordering of spacelike-separated events actually makes no difference if the theory respects microcausality: if the interaction Hamiltonian $\hat{H}_I(t)$ is the integral of a Hamiltonian density $\hat{h}_I(\mathbf{x})$ that commutes with itself at spacelike separation, the ordering assigned to spacelike-separated events is irrelevant, and any timelike coordinate t used to evaluate (2.7) will result in the same time evolution operator. This is what happens in typical relativistic field theories, where interactions are generally local and microcausal by construction. For particle detector models with $\hat{h}_I(\mathbf{x})$ given by (2.6), however, different time-orderings will generally lead to different time evolution operators. Therefore, a purely arbitrary choice of time coordinate used to evaluate (2.7) can make a difference in the results obtained for the final state of the detector at Σ_2 starting from the same input state at Σ_1 —and in this sense, particle detector models are not generally covariant.

These issues of the particle detector models described in 2.1 are rooted in the fact that the interaction (2.4) couples one degree of freedom of the detector (i.e., the observable $\mu_A(\tau)$) to many spacelike-separated observables of the field (i.e., the field observables $\mathcal{O}^A(\mathbf{x})$ in the support of $\Lambda(\mathbf{x})$ at fixed τ). This creates a mild level of non-locality in the theory, since the detector at a given value of its proper time τ is able to probe points that are spacelike to its position at that time. Relatedly, the coupling (2.4) also seems to implicitly select a preferred definition of “simultaneity”—namely, the one provided by the Fermi normal coordinate time relative to the detector’s worldline $\mathbf{z}(\tau)$ —, which in turn seems to be privileging one choice of time coordinate from the outset.⁸

It must be emphasized that these issues are not fatal problems for the majority of the concrete applications of particle detector models. After all, these models are not purported to be fundamental descriptions of relativistic systems all the way down to microscopic scales; rather, they are *effective* descriptions of systems which, in some regimes, can be accurately approximated as a single degree of freedom smeared in a small (but finite) region of space. Particle detectors can be very useful, for instance, as tools for describing the physics of the light-matter interaction, where the role of the detector is played by atoms coupled to the electromagnetic field [116]. In this case, the scale for the size of the detector is determined by the spread of the wavefunctions of the electron in the atom, and the approximation of

⁸In a sense, this can also be seen as a suggestion for how to deal with the noncovariance issue in practice: the presence of a preferred observer (i.e., the detector) provides, at least for practical calculations, a preferred time coordinate relative to which the time evolution operator is most naturally computed (i.e., the time coordinate naturally associated with the detector).

assigning a single quantum degree of freedom to an atom-sized spatial region is justified as long as we cannot resolve time intervals that are of the order of (or shorter than) the time it takes for light to travel between the nucleus and the outermost electron.

More generally, the impacts of the issues with causality and covariance on the actual predictions based on particle detector models are well-understood and can be quantified in detail [124, 57]. In particular, it is possible to show that the non-locality and breaking of covariance are well under control for most cases of interest. For the typical regimes where one would expect particle detectors to be good descriptions of the phenomenology at hand (essentially, when the spatial extension of the detector in its proper frame is much smaller than all the other relevant length scales of the problem, and when we consider the dynamics for times that are longer than the light-crossing time of the detector), these effects are ultimately negligible.

From a fundamental point of view, however, the mild violation of causality and breaking of covariance are certainly drawbacks of the formalism, and it would be more satisfactory to have a setup for local probes interacting with quantum fields which still fully respects relativistic causality and general covariance. From the discussion in the previous paragraphs, we see that there are essentially two ways of achieving this. The first is to demand that the detector be *pointlike*—i.e., the coupling between detector and field only occurs at a single point in space for each value of the detector’s proper time. In other words, given the detector’s trajectory $\mathbf{z}(\tau)$, we restrict ourselves to spacetime smearings $\Lambda(\mathbf{x})$ of the form

$$\Lambda(\mathbf{x}) = \int d\tau \chi(\tau) \frac{\delta^{(d+1)}(x - z(\tau))}{\sqrt{-g}}, \quad (2.14)$$

in which case the interaction (2.4) reduces to

$$S_I = \lambda \int d\tau \chi(\tau) \mu_A(\tau) \mathcal{O}^A(\mathbf{z}(\tau)). \quad (2.15)$$

The function $\chi(\tau)$ modulates the time dependence of the coupling in the detector’s proper frame, and is commonly referred to as the *switching function*. When the coupling is of this form, at no time can the detector probe points that are spacelike to it. Therefore, none of the issues with covariance and causality arise: pointlike detectors are thus fully causal and covariant.

Restricting to pointlike detectors can be sufficient for many purposes, especially when the size of the detector plays no particular role in the phenomenon we are interested in studying. However, the singular nature of the coupling can sometimes lead to UV divergences, and in actual physical scenarios where we may want to employ detector-based approaches, the

physical system playing the role of the detector is certainly not pointlike. This leads us to the second approach for solving the causality and covariance issues, which is to assign different degrees of freedom of the probe system to different spacelike-separated points in spacetime. This essentially amounts to a description of the probe that is fundamentally based on a relativistic field theory. Explaining how to make sense of this, and showing how to systematically connect a field-theory-based description of probe systems to the more usual UDW model, will be the goal of the rest of the next chapter.

Chapter 3

From quantum fields to particle detectors

3.1 Localized quantum fields

In line with the philosophy of describing probes for quantum fields as fully local and relativistic systems themselves, we will start with a model for a probe formulated as a full-fledged quantum field theory. In order to be useful as a *local* particle detector, we would like the field theory describing the probe to have the special property of being “confined to a finite region of space”. The most economical way of achieving this is to take the free action of the probe field ϕ_{D} (the detector) to be given by

$$S_{\text{D}}[\phi_{\text{D}}] = -\frac{1}{2} \int dV (g^{ab} (\nabla_a \phi_{\text{D}}) (\nabla_b \phi_{\text{D}}) + 2U(\mathbf{x}) \phi_{\text{D}}^2). \quad (3.1)$$

The first term in the action above is just the standard kinetic term for any scalar field, and the second term is an additional spacetime-dependent function $U(\mathbf{x})$ which acts as an external potential providing some localization profile to the probe. The intuition is that we want the potential to suppress all field configurations where ϕ_{D} does not go to zero as we move away from a finite region on any spatial slice; it is in that sense that the probe field is “confined” to a finite region of space.

To make this a bit more precise, we assume that the potential is such that, when restricted to any spacelike Cauchy surface Σ , $U(\mathbf{x})$ reaches a minimum (or a set of minima) within a region of finite volume in Σ , and $U(\mathbf{x})$ goes to infinity as the distance of the point

$\mathbf{x} \in \Sigma$ to this region along Σ goes to infinity. Another way of describing this, as done in [137], is to say that there is some timelike curve $\mathbf{z}(\tau)$ on $\mathcal{M}$ such that, for every value of the proper time parameter τ , it holds that $U(\mathbf{x}) \rightarrow +\infty$ when the squared geodesic distance between the event $\mathbf{x}$ and the point $\mathbf{z}(\tau)$, or equivalently the Synge's world function $\sigma(\mathbf{z}(\tau), \mathbf{x})$, goes to $+\infty$.¹ Given a foliation of $\mathcal{M}$ by a family of Cauchy surfaces Σ_τ and denoting by R_{Σ_τ} the smallest connected region containing all the minima of $U(\mathbf{x})$ in the Cauchy slice Σ_τ , the curve $\mathbf{z}(\tau)$ can be roughly pictured as the trajectory of the “center” of the region R_{Σ_τ} as τ varies.

We will now make some further assumptions about the setup in order to make things simpler and more explicit. First, we will take the background spacetime to be static—i.e., we assume there is a timelike Killing vector field χ^a that is orthogonal to a family of Cauchy surfaces. With that, we can define a coordinate system $(t, \mathbf{x})$ on $\mathcal{M}$, where the timelike coordinate t (which we will refer to as *Killing time* from now on) parametrizes the flow of the Killing vector field χ^a , and $\mathbf{x} = (x^1, \dots, x^d)$ are coordinates on the spatial slices of constant t . In terms of these coordinates, the timelike Killing vector takes the simple expression $\chi^a = (\partial_t)^a$, and the metric can be written as

$$ds^2 = -N^2(\mathbf{x})dt^2 + h_{ij}(\mathbf{x})dx^i dx^j. \quad (3.2)$$

The function $N(\mathbf{x})$ is known as the *lapse function*, and $h_{ij}(\mathbf{x})$ is the induced spatial metric on the surfaces of constant t . We will also assume that the confining potential $U(\mathbf{x})$ is invariant under the flow of the timelike Killing vector field χ^a . In coordinate-independent form, this is expressed as $\chi^a \nabla_a U = 0$; in the coordinates $(t, \mathbf{x})$, this just means that U is independent of the Killing time coordinate, $U(\mathbf{x}) = U(\mathbf{x})$. The assumption that U is confining then implies that $U(\mathbf{x}) \rightarrow +\infty$ as the distance between $\mathbf{x}$ and the point where the minimum of U is found goes to $+\infty$.

For the metric written as in (3.2) with the coordinates $(t, \mathbf{x})$, we can write

$$g^{ab} (\nabla_a \phi_D) (\nabla_b \phi_D) = -\frac{1}{N^2} (\partial_t \phi_D)^2 + h^{ij} (\partial_i \phi_D) (\partial_j \phi_D), \quad (3.3)$$

where h^{ij} is the inverse of the spatial metric. Writing the volume element $dV = \sqrt{-g} d^{d+1}x$

¹Note that Synge's world function (which is *half* the squared geodesic distance between two points) is negative when evaluated between timelike-separated events, so the requirement that it grows to $+\infty$ necessarily selects only points that are spacelike to $\mathbf{z}(\tau)$. For more details about the formal definition, properties, and applications of Synge's world function, see [140].

explicitly as $dt d^d \mathbf{x} N \sqrt{h}$, the free action of the probe then becomes

$$S_{\text{D}}[\phi_{\text{D}}] = \frac{1}{2} \int dt d^d \mathbf{x} \frac{\sqrt{h}}{N} (\partial_t \phi_{\text{D}})^2 - \int dt d^d \mathbf{x} \phi_{\text{D}} N \sqrt{h} U(\mathbf{x}) \phi_{\text{D}}^2 - \frac{1}{2} \int dt d^d \mathbf{x} N \sqrt{h} h^{ij} (\partial_i \phi_{\text{D}}) (\partial_j \phi_{\text{D}}). \quad (3.4)$$

Because of the confining potential $U(\mathbf{x})$, in every surface at a constant value of Killing time t , we can restrict ourselves to field configurations $\phi_{\text{D}}(t, \mathbf{x})$ that go to zero at large spacelike separation from a finite region of space. This allows us to integrate the last term of Equation (3.4) by parts at each slice of constant t , throw away the boundary term, and re-write the probe action as

$$S_{\text{D}}[\phi_{\text{D}}] = \frac{1}{2} \int dt d^d \mathbf{x} \frac{\sqrt{h}}{N} ((\partial_t \phi_{\text{D}})^2 - \phi_{\text{D}} E^2(\mathbf{x}) \phi_{\text{D}}), \quad (3.5)$$

where $E^2(\mathbf{x})$ is a differential operator that we define as acting on field configurations φ according to

$$\begin{aligned} E^2(\mathbf{x})\varphi &= 2N^2(\mathbf{x})U(\mathbf{x})\varphi - \frac{N}{\sqrt{h}}\partial_i \left(N\sqrt{h}h^{ij}\partial_j\varphi \right) \\ &= N^2(\mathbf{x}) \left(2U(\mathbf{x}) - \mathcal{D}_i \mathcal{D}^i \right) \varphi - \frac{1}{2}h^{ij}\partial_i (N^2) \partial_j \varphi. \end{aligned} \quad (3.6)$$

In the expression above, $\mathcal{D}_i \mathcal{D}^i$ is the spatial Laplacian on the surfaces of constant Killing time,

$$\mathcal{D}_i \mathcal{D}^i \varphi = \frac{1}{\sqrt{h}} \partial_i \left(\sqrt{h} h^{ij} \partial_j \varphi \right). \quad (3.7)$$

Note that the differential operator $E^2(\mathbf{x})$ is independent of Killing time, since U , h_{ij} , and N are. It also does not contain any time derivatives, and therefore, it is a “purely spatial” operator—i.e., $E^2(\mathbf{x})\varphi(t, \mathbf{x})$ evaluated on a surface Σ_{t_i} of constant time $t = t_i$ only depends on the pullback of $\varphi(t, \mathbf{x})$ to Σ_{t_i} , and is insensitive to how $\varphi(t, \mathbf{x})$ varies as we move away from Σ_{t_i} . This fact makes it convenient to think of $E^2(\mathbf{x})$ as a differential operator only on the surfaces Σ_t characterized by constant values of Killing time. Having that in mind, on each surface Σ_t , we can define an inner product on the space of scalar functions $f, g : \Sigma_t \rightarrow \mathbb{R}$ as

$$(f, g) = \int d^d \mathbf{x} \frac{\sqrt{h}}{N} f(\mathbf{x}) g(\mathbf{x}). \quad (3.8)$$

The operator $E^2(\mathbf{x})$ is clearly symmetric with respect to the inner product (3.8), and we will assume that there is an appropriate domain where $E^2(\mathbf{x})$ is self-adjoint. This can be obtained, for instance, by only considering test functions that decay sufficiently fast at spatial infinity. We will also demand that the space of admissible field configurations $\varphi(t, \mathbf{x})$ be defined such that, for any fixed Killing time $t = t_i$, the pullback of $\varphi(t, \mathbf{x})$ to Σ_{t_i} is in the domain of self-adjointness of $E^2(\mathbf{x})$. Finally, in order for the probe field to be stable, we will assume that E^2 is positive-definite: i.e., $(f, E^2 f) > 0$ for every test function $f(\mathbf{x})$ in the domain of E^2 . Putting all of this together, we conclude that $E^2(\mathbf{x})$ possesses a set of eigenfunctions $\{v_{\mathbf{n}}(\mathbf{x})\}$,

$$E^2(\mathbf{x})v_{\mathbf{n}}(\mathbf{x}) = \omega_{\mathbf{n}}^2 v_{\mathbf{n}}(\mathbf{x}), \quad (3.9)$$

which are orthonormal according to the inner product (3.8),

$$(v_{\mathbf{n}}, v_{\mathbf{n}'}) = \int d^d \mathbf{x} \frac{\sqrt{h}}{N} v_{\mathbf{n}}(\mathbf{x}) v_{\mathbf{n}'}(\mathbf{x}) = \delta_{\mathbf{n}, \mathbf{n}'}, \quad (3.10)$$

and form a basis for the set of admissible field configurations $\varphi(t, \mathbf{x})$ at any fixed time t . The assumption that the potential $U(\mathbf{x})$ is confining ensures that E^2 has a discrete spectrum [148, 161], which in turn means that the label $\mathbf{n}$ ranges over a discrete set of values. We are thus justified in using the orthogonality condition (3.10) in terms of a Kronecker delta, as opposed to a Dirac delta (which would be needed if the label $\mathbf{n}$ were continuous). The assumption that E^2 is positive-definite further justifies denoting the eigenvalues of E^2 by $\omega_{\mathbf{n}}^2$ —which was, in fact, the reason why the operator was defined with the very suggestive notation E^2 to begin with.

All eigenfunctions $v_{\mathbf{n}}(\mathbf{x})$ go to zero as $U(\mathbf{x}) \rightarrow +\infty$, and are most strongly supported around the regions where $U(\mathbf{x})$ attains its minima. Since, by assumption, all admissible field configurations $\phi_{\text{D}}(t, \mathbf{x})$ must go to zero as we move away from the minima of $U(\mathbf{x})$ along any slice Σ_t of constant Killing time t , we can expand ϕ_{D} as

$$\phi_{\text{D}}(t, \mathbf{x}) = \sum_{\mathbf{n}} \phi_{\mathbf{n}}(t) v_{\mathbf{n}}(\mathbf{x}), \quad (3.11)$$

where the full time dependence of the field configuration is incorporated in the amplitudes $\phi_{\mathbf{n}}(t)$ in the expansion above. Plugging this back into Equation (3.5) and using the fact that the set of functions $\{v_{\mathbf{n}}(\mathbf{x})\}$ is orthonormal in the sense of the inner product from

Equation (3.8), the action for the probe field becomes

$$S_{\text{D}}[\phi_{\text{D}}] = \sum_{\mathbf{n}} S_{\mathbf{n}}[\phi_{\mathbf{n}}], \quad (3.12)$$

$$S_{\mathbf{n}}[\phi_{\mathbf{n}}] = \frac{1}{2} \int dt \left(\dot{\phi}_{\mathbf{n}}^2 - \omega_{\mathbf{n}}^2 \phi_{\mathbf{n}}^2 \right).$$

where we have defined $\dot{\phi}_{\mathbf{n}} \equiv d\phi_{\mathbf{n}}/dt$. We recognize (3.12) as the action for a set of decoupled harmonic oscillators (which we will also refer to as *modes* of the field) of frequency $\omega_{\mathbf{n}}$, and with amplitudes given by $\phi_{\mathbf{n}}(t)$, as we first wrote in (2.12).

For concreteness, assume that, on every surface of constant Killing time, the potential $U(\mathbf{x})$ has a unique minimum, with spatial coordinates given by $\mathbf{x}_0$. We can normalize the timelike Killing vector field χ^a by a constant such that $\chi^a \chi_a = -1$ at $\mathbf{x} = \mathbf{x}_0$. With this choice, the Killing time t becomes a proper time parameter for the trajectory of the minimum of the potential, which, in the coordinates $(t, \mathbf{x})$, takes the very simple form $z_{\text{D}}^{\mu}(t) = (t, \mathbf{x}_0)$. Then, we can interpret the action for the localized probe field $\phi_{\text{D}}(\mathbf{x})$ as simply that of an infinite (but discrete) tower of simple harmonic oscillators with proper time t and proper energy gap $\omega_{\mathbf{n}}$, which follow a static trajectory $\mathbf{z}_{\text{D}}(t)$ in spacetime. We can subsequently obtain an effective theory for a finite number of such harmonic oscillators by tracing out all but finitely many of them. This is the logic that will allow us to reduce the theory of a probe field coupled to a Klein-Gordon field to that of a series of particle detectors coupled to a quantum field via the UDW-like coupling that is commonly found in the literature on RQI, as we explain in the next section.

3.2 From probe field to particle detectors via the path integral

In this section we will show how to reduce the probe model given by the localized quantum field of Section 3.1 to the more standard harmonic-oscillator UDW model, using the Schwinger-Keldysh path integral.

3.2.1 Setting up the path integral

Consider that the probe field ϕ_{D} will be used as a detector that couples to a real Klein-Gordon field ψ . For brevity, we will often refer to ψ simply as the *target* field. The joint

system of probe and target field can be described by the action

$$S[\phi_D, \psi] = S_D[\phi_D] + S_F[\psi] + S_I[\phi_D, \psi], \quad (3.13)$$

where S_D and S_F are the free actions of the ϕ_D and ψ , and S_I encodes the interaction between the two. For the target field ψ , we take

$$S_F[\psi] = -\frac{1}{2} \int dV \left(g^{ab} \nabla_a \psi \nabla_b \psi + M^2 \psi^2 + \xi R \psi^2 \right). \quad (3.14)$$

This is the usual action of a Klein-Gordon field ψ with mass M and possibly non-minimal coupling to curvature that we first considered in Section 2.1 in Equation (2.10). The probe action $S_D[\phi_D]$ will be the exact same one that we introduced in Section 3.1, given by Equation (3.1). Lastly, the interaction between the probe and target field is given by

$$S_I[\phi_D, \psi] = \lambda \int dV \zeta(\mathbf{x}) \psi(\mathbf{x}) \phi_D(\mathbf{x}), \quad (3.15)$$

where $\zeta(\mathbf{x})$ is a function that dictates the localization of the coupling between probe and target field in spacetime, and λ is an overall coupling constant.

We will make use of the same simplifying assumptions studied in Section 3.1—namely, we take the background metric to be static, and we assume that the confining potential of the probe is independent of the Killing time t parametrizing the flow of the timelike Killing vector field $\chi^a = (\partial_t)^a$. The function $\zeta(\mathbf{x})$ in the interaction (3.15), however, can still depend on t ; this simply corresponds to the case where the strength of the interaction between the probe and the system of interest can be switched on and off with time.

Let us now present the quantum theory of the Klein-Gordon field coupled to the probe. The central object of the theory will be taken to be the Schwinger-Keldysh path integral [158, 108, 74], which in our case can be written as

$$\begin{aligned} \mathcal{Z}[\mathbf{J}, \mathbf{J}'] = & \int [d\Phi_f][d\Psi_f] \int [d\Phi_i][d\Psi_i] \int [d\Phi'_i][d\Psi'_i] \\ & \times \rho(\Phi_i, \Psi_i; \Phi'_i, \Psi'_i) Z[\Phi_f, \Psi_f; \Phi_i, \Psi_i | \mathbf{J}] Z[\Phi_f, \Psi_f; \Phi'_i, \Psi'_i | \mathbf{J}']^*. \end{aligned} \quad (3.16)$$

There is quite a bit of notation to unpack in Equation (3.16), so let us take some time to clarify what each of the terms means.

For concreteness, consider the dynamics of the theory between two Cauchy surfaces characterized by constant values of Killing time t_i and t_f , with $t_i < t_f$. If we want, we can

take $t_i \rightarrow -\infty$ and $t_f \rightarrow +\infty$, and picture the system as evolving from the asymptotic past to the asymptotic future. The capital Φ s and Ψ s in Equation (3.16) denote classical field configurations $\Phi(\mathbf{x})$, $\Psi(\mathbf{x})$ at constant-time slices for the fields ϕ_D and ψ respectively, with the configurations labeled with the subscript “ i ” being defined at the initial Killing time t_i , and those with the subscript “ f ” being defined at the final Killing time t_f . The quantity $\rho(\Phi_i, \Psi_i; \Phi'_i, \Psi'_i)$ denotes matrix elements of a given initial state for the probe and target field in a basis of field eigenstates: that is, we take a state $\hat{\rho}$ describing the joint system of probe and target field at the initial Cauchy surface Σ_{t_i} , and define

$$\rho(\Phi_i, \Psi_i; \Phi'_i, \Psi'_i) := \langle \Phi_i, \Psi_i | \hat{\rho} | \Phi'_i, \Psi'_i \rangle, \quad (3.17)$$

where we have $|\Phi_i, \Psi_i\rangle \equiv |\Phi_i\rangle \otimes |\Psi_i\rangle$, and the kets $|\Phi_i\rangle$ and $|\Psi_i\rangle$ are states in the Hilbert spaces of the probe and the target field satisfying

$$\hat{\phi}_D(t_i, \mathbf{x}) |\Phi_i\rangle = \Phi_i(\mathbf{x}) |\Phi_i\rangle, \quad (3.18)$$

$$\hat{\psi}(t_i, \mathbf{x}) |\Psi_i\rangle = \Psi_i(\mathbf{x}) |\Psi_i\rangle. \quad (3.19)$$

In short, $|\Phi_i\rangle$ and $|\Psi_i\rangle$ are field eigenstates of $\hat{\phi}_D(t_i, \mathbf{x})$ and $\hat{\psi}(t_i, \mathbf{x})$ associated to the classical field configurations $\Phi_i(\mathbf{x})$ and $\Psi_i(\mathbf{x})$ at time t_i .

The quantity $Z[\Phi_f, \Psi_f; \Phi_i, \Psi_i | \mathbf{J}]$ is the Feynman path integral, which gives transition amplitudes from initial field configurations Φ_i, Ψ_i at time $t = t_i$ to final field configurations Φ_f, Ψ_f at $t = t_f$, in the presence of sources $\mathbf{J} \equiv (J_D, J)$ for the probe and target field, respectively. It can be expressed as

$$Z[\Phi_f, \Psi_f; \Phi_i, \Psi_i | \mathbf{J}] = \int \mathcal{D}\phi_D \mathcal{D}\psi e^{iS[\phi_D, \psi | \mathbf{J}]} \quad (3.20)$$

where we have abbreviated

$$S[\phi_D, \psi | \mathbf{J}] := S[\phi_D, \psi] + \int dV (J_D(\mathbf{x}) \phi_D(\mathbf{x}) + J(\mathbf{x}) \psi(\mathbf{x})) \quad (3.21)$$

and in Equation (3.20) the path integral is performed with the boundary conditions

$$\phi_D(t_i, \mathbf{x}) = \Phi_i(\mathbf{x}), \quad \psi(t_i, \mathbf{x}) = \Psi_i(\mathbf{x}), \quad (3.22)$$

$$\phi_D(t_f, \mathbf{x}) = \Phi_f(\mathbf{x}), \quad \psi(t_f, \mathbf{x}) = \Psi_f(\mathbf{x}). \quad (3.23)$$

Finally, the several integrals over $[d\Phi]$ s and $[d\Psi]$ s in Equation (3.16) denote sums over classical field configurations (according to some measure) at the corresponding constant-time

slices.

It is important to not confuse the path integral measure $\mathcal{D}\phi_{\text{D}}$ with the measure $[\text{d}\Phi]$: the former is a measure on field configurations $\phi_{\text{D}}(t, \mathbf{x})$ supported on the entire spacetime region between times t_i and t_f , whereas the latter simply ranges over fixed-time field configurations $\Phi(\mathbf{x})$. If we were dealing with a $(0 + 1)$ -dimensional field theory of one scalar field $q(t)$ —or equivalently, if we were doing quantum mechanics of a single bosonic particle—then $[\text{d}\Phi]$ would essentially reduce to a simple measure $\text{d}q$ on the real line, whereas $\mathcal{D}\phi_{\text{D}}$ would become the quantum-mechanical path integral measure $\mathcal{D}q$ for a particle with one position degree of freedom q .²

Despite its perhaps uninviting looks, the Schwinger-Keldysh path integral (3.16) is actually a rather simple object at a conceptual level. If we have the Schrodinger picture in mind³, we can understand Equation (3.16) as arising from the following construction (for a very detailed review, see e.g. [15]). First, we take some initial state $\hat{\rho}$ for the probe and target field at an early time t_i . Next, we define the operator $\hat{U}_{\mathbf{J}}(t_f, t_i) \hat{\rho} \hat{U}_{\mathbf{J}'}(t_f, t_i)^\dagger$, where $\hat{U}_{\mathbf{J}}(t_f, t_i)$ and $\hat{U}_{\mathbf{J}'}(t_f, t_i)$ are the time evolution operators from the Cauchy surface at Killing time t_i to the Cauchy surface at Killing time t_f in the presence of sources $\mathbf{J}$ and $\mathbf{J}'$ for the probe and target fields. We then expand the initial state on a common basis of field eigenstates for the probe and target field at time t_i , and take the trace of the resulting operator $\hat{U}_{\mathbf{J}}(t_f, t_i) \hat{\rho} \hat{U}_{\mathbf{J}'}(t_f, t_i)^\dagger$ on a basis of field eigenstates at the later time t_f . The Feynman path integral appears because it corresponds precisely to the transition amplitudes between field eigenstates from early to late times—or equivalently, they are just the matrix elements of the time evolution operator in the basis of field eigenstates at fixed time slices. The addition of the sources J_{D}, J in the Feynman path integral for the probe and target just serves the purpose of allowing us to compute expectation values of observables in the evolved state, by the usual trick of taking functional derivatives with respect to the sources

²We note that, rigorously speaking, the quantum-mechanical path integral measure is not a fully well-defined mathematical object. Since we do not intend to address this particular problem in this thesis, we will limit ourselves to the most common physicist’s treatment, where an implicit choice of path integral measure is justified *a posteriori* by verifying that the results obtained via the path integral are physically reasonable. For a much less heuristic and more detailed discussion on formally defining the quantum-mechanical path integral measure, however, see [35].

³Of course, a completely analogous interpretation could be provided in terms of the Heisenberg picture, which is the one that is most useful/convenient for most purposes in field theory. I just felt the summary presented in this paragraph was more succinctly stated in the Schrodinger picture.

and then setting the sources to zero. For instance, in terms of $\mathcal{Z}[\mathbf{J}, \mathbf{J}']$, we can write

$$\langle \hat{\psi}(\mathbf{x}) \rangle_{\hat{\rho}} = \frac{(-i)}{\sqrt{-g(x)}} \frac{\delta \mathcal{Z}}{\delta J(x)} \Big|_{(\mathbf{J}, \mathbf{J}')=0} \quad (3.24)$$

$$\langle \hat{\psi}(\mathbf{x}) \hat{\psi}(\mathbf{x}') \rangle_{\hat{\rho}} = \frac{1}{\sqrt{-g(x)} \sqrt{-g(x')}} \frac{\delta^2 \mathcal{Z}}{\delta J(x) \delta J'(x')} \Big|_{(\mathbf{J}, \mathbf{J}')=0}, \quad (3.25)$$

$$\langle \mathcal{T}(\hat{\psi}(\mathbf{x}) \hat{\psi}(\mathbf{x}')) \rangle_{\hat{\rho}} = \frac{(-i)^2}{\sqrt{-g(x)} \sqrt{-g(x')}} \frac{\delta^2 \mathcal{Z}}{\delta J(x) \delta J(x')} \Big|_{(\mathbf{J}, \mathbf{J}')=0}, \quad (3.26)$$

and so on. Of course, similar expressions also hold for expectation values involving the probe field, if we take functional derivatives with respect to the probe's source J_{D} . We also used $\mathcal{T}$ to denote the time-ordering operation—which, based on the discussion in Subsection 2.2, can be defined with respect to any timelike coordinate that yields the same time orientation (i.e., the same notion of past and future) as the chosen Killing time t , since both the probe and the target fields are fully relativistic and causal by construction [124, 57]. It is also worth pointing out that, unlike the bare Feynman path integral that we are probably most familiar with, the Schwinger-Keldysh path integral (3.16) can yield arbitrary (i.e., not necessarily time-ordered) expectation values, as emphasized, for instance, by Equation (3.25). This is due to the presence of two independent sources $\mathbf{J}, \mathbf{J}'$ in the definition of (3.16)—note that the only difference between (3.25) and (3.26) is that the functional derivative is taken with respect to the same current J in the latter case, but with respect to different currents J and J' in the former. For this reason, the Schwinger-Keldysh path integral is sometimes described as the *expectation-value* generating functional.

For later reference, it is sometimes also useful to introduce the so-called *influence phase* $\mathcal{W}(\mathbf{J}, \mathbf{J}')$, defined such that

$$\mathcal{Z}[\mathbf{J}, \mathbf{J}'] = e^{i\mathcal{W}[\mathbf{J}, \mathbf{J}']}. \quad (3.27)$$

Just like $\mathcal{Z}[\mathbf{J}, \mathbf{J}']$ is the generating functional for expectation values, $\mathcal{W}[\mathbf{J}, \mathbf{J}']$ is the generating functional for *connected* expectation values.

The Schwinger-Keldysh path integral often appears with different names, depending on where it shows up in the physics literature. In many-body physics and open quantum systems, it is common to refer to it as the Feynman-Vernon influence functional [147, 41] or decoherence functional [65]. It also comes under the guise of the so-called *closed-time-path*

or *in-in formalism* [40, 3], which is a general framework for the non-equilibrium dynamics of quantum systems where the Schwinger-Keldysh path integral plays a central role. These names are probably motivated by the fact that $\mathcal{Z}$ can be visualized as a path integral over two copies of the original spacetime region between the Cauchy surfaces Σ_i and Σ_f , where the integral over the first copy “runs forwards” in time (from t_i to t_f) and the integral over the second copy “runs backwards” (from t_f to t_i), thus bringing us back to the initial time t_i where the input state $\hat{\rho}$ was originally defined. This is in contrast with the standard Feynman path integral (3.20), which is most readily applicable to the computation of transition amplitudes between early (in) states and late (out) states, with time only going in one direction.

Now that we know what each of the terms in (3.16) means, we can proceed with the Schwinger-Keldysh path integral for the probe-field system. We saw in Section 3.1 that the action for the probe can be written as

$$S_{\text{D}}[\phi_{\text{D}}] = \sum_{\mathbf{n}} \left[\frac{1}{2} \int dt \left(\dot{\phi}_{\mathbf{n}}^2 - \omega_{\mathbf{n}}^2 \phi_{\mathbf{n}}^2 \right) \right], \quad (3.28)$$

where we recall that $\phi_{\mathbf{n}}(t)$ are the amplitudes for the field configuration $\phi_{\text{D}}(t, \mathbf{x})$ expressed as (3.11),

$$\phi_{\text{D}}(t, \mathbf{x}) = \sum_{\mathbf{n}} \phi_{\mathbf{n}}(t) v_{\mathbf{n}}(\mathbf{x}), \quad (3.29)$$

and $\{v_{\mathbf{n}}(\mathbf{x})\}$ is a basis of eigenfunctions of the spatial differential operator $E^2(\mathbf{x})$ defined in Equation (3.6). By applying the same expansion (3.29) to the coupling of the probe with the target field ψ and the probe’s source J_{D} in Equation (3.21), we directly obtain

$$\lambda \int dV \zeta(\mathbf{x}) \psi(\mathbf{x}) \phi_{\text{D}}(\mathbf{x}) = \sum_{\mathbf{n}} \lambda \int dt \phi_{\mathbf{n}}(t) \psi_{\mathbf{n}}(t), \quad (3.30)$$

$$\int dV J_{\text{D}}(\mathbf{x}) \phi_{\text{D}}(\mathbf{x}) = \sum_{\mathbf{n}} \int dt \phi_{\mathbf{n}}(t) J_{\mathbf{n}}(t) \quad (3.31)$$

where clearly

$$\psi_{\mathbf{n}}(t) = \int d^d \mathbf{x} N \sqrt{h} v_{\mathbf{n}}(\mathbf{x}) \zeta(t, \mathbf{x}) \psi(t, \mathbf{x}), \quad (3.32)$$

$$J_{\mathbf{n}}(t) = \int d^d \mathbf{x} N \sqrt{h} v_{\mathbf{n}}(\mathbf{x}) J_{\text{D}}(t, \mathbf{x}). \quad (3.33)$$

Now, from the point of view of the path integral, the expansion (3.29) can essentially be seen as a change of dynamical variables of the probe system, from the field configurations $\phi_{\text{D}}(\mathbf{x})$ to the amplitudes $\phi_{\mathbf{n}}(t)$. This helps us in two important ways. First, it allows us to schematically write the integration measure $\mathcal{D}\phi_{\text{D}}$ in the Feynman path integral (3.20) as

$$\mathcal{D}\phi_{\text{D}} = \mathcal{N} \prod_{\mathbf{n}} \mathcal{D}\phi_{\mathbf{n}}, \quad (3.34)$$

where $\mathcal{N}$ plays the role of the Jacobian determinant from the change of variables. It can be shown by a simple comparison argument that we must have $\mathcal{N} = 1$ (see Appendix A.1 for details).⁴ Second, since every field eigenstate $|\Phi\rangle$ at a surface of constant Killing time can be identified with a unique set of coefficients $\varphi_{\mathbf{n}}$ in the mode expansion of $\Phi(\mathbf{x})$ in terms of the eigenfunctions $\{v_{\mathbf{n}}(\mathbf{x})\}$, we always have

$$|\Phi\rangle = \bigotimes_{\mathbf{n}} |\varphi_{\mathbf{n}}\rangle \quad (3.35)$$

where each $|\varphi_{\mathbf{n}}\rangle$ is an eigenstate of $\hat{\phi}_{\mathbf{n}}$ with eigenvalue $\varphi_{\mathbf{n}}$ corresponding to the component of $\Phi(\mathbf{x})$ in the basis $\{v_{\mathbf{n}}(\mathbf{x})\}$: i.e.,

$$\Phi(\mathbf{x}) = \sum_{\mathbf{n}} \varphi_{\mathbf{n}} v_{\mathbf{n}}(\mathbf{x}). \quad (3.36)$$

With this, we can write the Feynman path integral as

$$\begin{aligned} Z[\Phi_f, \Psi_f; \Phi_i, \Psi_i | \mathbf{J}] &= \int \mathcal{D}\psi e^{iS_{\text{F}}} e^{i \int dV J(\mathbf{x}) \psi(\mathbf{x})} \\ &\times \prod_{\mathbf{n}} Z_{\text{D}}[\varphi_{\mathbf{n},f}; \varphi_{\mathbf{n},0} | J_{\mathbf{n}} + \lambda \psi_{\mathbf{n}}], \end{aligned} \quad (3.37)$$

⁴One might worry that equations such as (3.34) with $\mathcal{N} = 1$, as well as (3.35) or (3.41), are not dimensionally consistent—after all, the units of a scalar field in $d + 1$ dimensions are certainly not those of a product of infinitely many scalar fields in $0 + 1$ dimensions. The reason why these equations are fine is that, as it turns out, the more rigorous definitions of both the integration measure $\mathcal{D}\phi_{\text{D}}$ and the inner product between field configurations at a constant time (which indirectly impacts the normalizations of $|\Phi\rangle$ and $[d\Phi]$) actually carry implicit choices of dimensionful scales. If one carefully keeps track of these choices of scale in both $d = 0$ and higher dimensions, it can be shown [35] that things work out in such a way that (3.34) (with $\mathcal{N} = 1$), as well as (3.35) and (3.41), all have the correct units. I am deeply thankful to Iván M. Burbano for clarifications on this.

where we have defined

$$Z_{\text{D}}[\varphi_{\mathbf{n},f}; \varphi_{\mathbf{n},0} | J_{\mathbf{n}} + \lambda \psi_{\mathbf{n}}] = \int \mathcal{D}\phi_{\mathbf{n}} e^{iS_{\mathbf{n}}[\phi_{\mathbf{n}}]} \times e^{i \int dt \phi_{\mathbf{n}}(t)(\lambda \psi_{\mathbf{n}}(t) + J_{\mathbf{n}}(t))}. \quad (3.38)$$

The path integral in the target field ψ is performed with the boundary conditions

$$\psi(t_i, \mathbf{x}) = \Psi_i(\mathbf{x}), \quad \psi(t_f, \mathbf{x}) = \Psi_f(\mathbf{x}) \quad (3.39)$$

and similarly, the boundary conditions for the Feynman path integrals corresponding to each mode $\phi_{\mathbf{n}}$ are

$$\phi_{\mathbf{n}}(t_i) = \varphi_{\mathbf{n},0}, \quad \phi_{\mathbf{n}}(t_f) = \varphi_{\mathbf{n},f} \quad (3.40)$$

where $\varphi_{\mathbf{n},0}$ and $\varphi_{\mathbf{n},f}$ are the components of the probe field configurations $\Phi_i(\mathbf{x})$ and $\Phi_f(\mathbf{x})$ in terms of the basis $\{v_{\mathbf{n}}(\mathbf{x})\}$ at times t_i and t_f , respectively. Finally, the measure $[\text{d}\Phi]$ on fixed-time field configurations in the Schwinger-Keldysh path integral (3.16) can be simply interpreted as

$$[\text{d}\Phi] = \prod_{\mathbf{n}} \text{d}\varphi_{\mathbf{n}}. \quad (3.41)$$

With this, we are now finally ready to see how to reduce the probe field to a finite number of modes at the level of the Schwinger-Keldysh path integral.

3.2.2 Reduction of probe field to a finite set of modes

Consider a subset of the probe's degrees of freedom consisting of modes indexed by $\mathbf{n}$ such that $\mathbf{n} \in A$, where A comprises some finite set of labels. As a slight abuse of language and for the sake of brevity, we will often refer to these modes simply as “the modes in A ”. The full Hilbert space of the probe and target field (which can be described initially as $\mathcal{H} = \mathcal{H}_{\psi} \otimes \mathcal{H}_{\text{D}}$, where $\mathcal{H}_{\psi}$ denotes the Hilbert space for the target field, and $\mathcal{H}_{\text{D}}$ is the Hilbert space of the probe) can then be further decomposed as $\mathcal{H} = \mathcal{H}_{\psi} \otimes \mathcal{H}_A \otimes \mathcal{H}_{\bar{A}}$, where $\mathcal{H}_A$ corresponds to the subsystem comprising the modes in A , and $\mathcal{H}_{\bar{A}}$ contains all of the remaining modes of the probe field which are not included in A .

The reduction of the theory of the probe field to a finite number of modes will consist of two steps. First, we will assume that we can only turn on sources J_{D} for the probe system such that $J_{\mathbf{n}} = 0$ for $\mathbf{n} \notin A$. Under this assumption, the Feynman path integral (3.37) can

be written explicitly as

$$\begin{aligned}
Z[\Phi_f, \Psi_f; \Phi_i, \Psi_i | \mathbf{J}] &= \int \mathcal{D}\psi e^{iS_F} e^{i \int dV J(\mathbf{x}) \psi(\mathbf{x})} \\
&\times \prod_{\mathbf{n} \in A} Z_D[\varphi_{\mathbf{n},f}; \varphi_{\mathbf{n},0} | J_{\mathbf{n}} + \lambda \psi_{\mathbf{n}}], \\
&\times \prod_{\mathbf{n} \notin A} Z_D[\varphi_{\mathbf{n},f}; \varphi_{\mathbf{n},0} | \lambda \psi_{\mathbf{n}}].
\end{aligned} \tag{3.42}$$

Next, we will assume that at early times, the field+probe system is initialized in a state that contains no correlations between the subsystems $\mathcal{H}_\psi \otimes \mathcal{H}_A$ and $\mathcal{H}_{\bar{A}}$. In other words, the initial state at time $t = t_i$ can be expressed as

$$\hat{\rho} = \hat{\rho}_{\psi,A} \otimes \hat{\rho}_{\bar{A}}, \tag{3.43}$$

where $\hat{\rho}_{\bar{A}}$ is a fixed state on the Hilbert space of all the modes of the probe field excluding the ones in A . For concreteness, we will take this state to be the vacuum, i.e., we have

$$\hat{\rho}_{\bar{A}} = \bigotimes_{\mathbf{n} \notin A} |0_{\mathbf{n}}\rangle \langle 0_{\mathbf{n}}| \tag{3.44}$$

with $|0_{\mathbf{n}}\rangle$ being the vacuum state of the mode $\mathbf{n}$. As usual, this is defined for any mode as the unique state that is annihilated by the annihilation operator associated to the mode $\mathbf{n}$ —i.e., it is the state such that

$$\hat{a}_{\mathbf{n}} |0_{\mathbf{n}}\rangle = 0 \tag{3.45}$$

where we define

$$\hat{a}_{\mathbf{n}} = \frac{1}{\sqrt{2}}(\hat{\phi}_{\mathbf{n}} + i\hat{\pi}_{\mathbf{n}}) \tag{3.46}$$

and $\hat{\pi}_{\mathbf{n}} := d\hat{\phi}_{\mathbf{n}}/dt$ is the canonically conjugate momentum to $\hat{\phi}_{\mathbf{n}}$.

If we now replace the state (3.43) and the Feynman path integral (3.42) explicitly in Equation (3.16), we find that the Schwinger-Keldysh path integral can be evaluated in the following way. For brevity of notation, we will abbreviate by $\mathbf{J}_A \equiv (\{J_{\mathbf{n} \in A}\}, J)$ the collection of sources for the probe+field system that excludes the modes not in A , and also

use the convenient notations

$$\phi_A \equiv \{\phi_{\mathbf{n}}, \mathbf{n} \in A\}, \quad (3.47)$$

$$d\varphi_A \equiv \prod_{\mathbf{n} \in A} d\varphi_{\mathbf{n}}, \quad (3.48)$$

$$\mathcal{D}\phi_A \equiv \prod_{\mathbf{n} \in A} \mathcal{D}\phi_{\mathbf{n}}, \quad (3.49)$$

$$\begin{aligned} S[\phi_A, \psi | \mathbf{J}_A] &\equiv S_F[\psi] + \sum_{\mathbf{n} \in A} S_{\mathbf{n}}[\phi_{\mathbf{n}}] \\ &+ \sum_{\mathbf{n} \in A} \lambda \int dt \psi_{\mathbf{n}}(t) \phi_{\mathbf{n}}(t) \\ &+ \sum_{\mathbf{n} \in A} \int dt J_{\mathbf{n}}(t) \phi_{\mathbf{n}}(t) + \int dV J(\mathbf{x}) \psi(\mathbf{x}). \end{aligned} \quad (3.50)$$

With these conventions, we can write the Schwinger-Keldysh path integral for the modes in A and the target field as

$$\begin{aligned} \mathcal{Z}[\mathbf{J}_A, \mathbf{J}'_A] &= \int d\varphi_{A,f} d\varphi_{A,i} d\varphi'_{A,i} \\ &\times \int [d\Psi_f] [d\Psi_i] [d\Psi'_i] \\ &\times \rho_{\psi,A}(\varphi_{A,i}, \Psi_i; \varphi'_{A,i}, \Psi'_i) \\ &\times \int \mathcal{D}\phi_A \mathcal{D}\phi'_A \mathcal{D}\psi \mathcal{D}\psi' \\ &\times e^{i(S[\phi_A, \psi | \mathbf{J}_A] - S[\phi'_A, \psi' | \mathbf{J}'_A])} \\ &\times e^{i\tilde{S}[\psi, \psi']}, \end{aligned} \quad (3.51)$$

where the last factor above is defined as

$$\begin{aligned}
e^{i\tilde{S}[\psi, \psi']} \equiv & \prod_{\mathbf{n} \notin A} \left(\int d\varphi_{\mathbf{n},f} d\varphi_{\mathbf{n},i} d\varphi'_{\mathbf{n},i} \right. \\
& \times \langle \varphi_{\mathbf{n},i} | 0_{\mathbf{n}} \rangle \langle 0_{\mathbf{n}} | \varphi'_{\mathbf{n},i} \rangle \\
& \times \int \mathcal{D}\phi_{\mathbf{n}} \mathcal{D}\phi'_{\mathbf{n}} \\
& \times e^{i(S_{\mathbf{n}}[\phi_{\mathbf{n}}] + \lambda \int dt \psi_{\mathbf{n}}(t) \phi_{\mathbf{n}}(t))} \\
& \left. \times e^{-i(S_{\mathbf{n}}[\phi'_{\mathbf{n}}] + \lambda \int dt \psi'_{\mathbf{n}}(t) \phi'_{\mathbf{n}}(t))} \right). \tag{3.52}
\end{aligned}$$

But Equation (3.52) is nothing but a product of Schwinger-Keldysh path integrals for the modes not included in A , with $\lambda\psi_{\mathbf{n}}$ playing the role of the source for the mode $\mathbf{n}$. We remember that the modes not included in A are assumed to start in the ground state, which is why the factor $\langle \varphi_{\mathbf{n},i} | 0_{\mathbf{n}} \rangle \langle 0_{\mathbf{n}} | \varphi'_{\mathbf{n},i} \rangle$ appears in Equation (3.52). In this case, the quantity $\tilde{S}[\psi, \psi']$ (which is nothing but a vacuum influence phase, as we defined it in (3.27)) can be computed exactly [15], and the result gives us

$$\begin{aligned}
\tilde{S}[\psi, \psi'] = & i \frac{\lambda^2}{2} \sum_{\mathbf{n} \notin A} \int_{t_i}^{t_f} dt \int_{t_i}^{t_f} dt' \left(G_{\mathbf{n}}(t, t') \psi_{\mathbf{n}}(t) \psi_{\mathbf{n}}(t') \right. \\
& - W_{\mathbf{n}}(t, t') \psi'_{\mathbf{n}}(t) \psi_{\mathbf{n}}(t') \\
& + G_{\mathbf{n}}(t, t')^* \psi'_{\mathbf{n}}(t) \psi'_{\mathbf{n}}(t') \\
& \left. - W_{\mathbf{n}}(t, t')^* \psi_{\mathbf{n}}(t) \psi'_{\mathbf{n}}(t') \right), \tag{3.53}
\end{aligned}$$

where we have defined the Wightman and Feynman functions

$$W_{\mathbf{n}}(t, t') := \frac{1}{2\omega_{\mathbf{n}}} e^{-i\omega_{\mathbf{n}}(t-t')}, \tag{3.54}$$

$$\begin{aligned}
G_{\mathbf{n}}(t, t') &:= \theta(t - t') W_{\mathbf{n}}(t, t') + \theta(t' - t) W_{\mathbf{n}}(t', t) \\
&= \frac{1}{2\omega_{\mathbf{n}}} e^{-i\omega_{\mathbf{n}}|t-t'|}, \tag{3.55}
\end{aligned}$$

with their complex conjugates given by

$$W_{\mathbf{n}}(t, t')^* = \frac{1}{2\omega_{\mathbf{n}}} e^{i\omega_{\mathbf{n}}(t-t')}, \quad (3.56)$$

$$G_{\mathbf{n}}(t, t')^* = \frac{1}{2\omega_{\mathbf{n}}} e^{i\omega_{\mathbf{n}}|t-t'|}. \quad (3.57)$$

Now let us step back and contemplate what we have just obtained. If it were not for the very last factor of $e^{i\tilde{S}[\psi, \psi']}$, the Schwinger-Keldysh path integral in Equation (3.51) would be *exactly* the same as that of a finite number of harmonic-oscillator UDW detectors with amplitudes $\phi_{\mathbf{n}}$ ($\mathbf{n} \in A$), following an orbit of the Killing vector field χ^a , each coupling linearly to the target field $\psi(\mathbf{x})$ with the smearing functions $\Lambda_{\mathbf{n}}(t, \mathbf{x}) := \zeta(t, \mathbf{x})v_{\mathbf{n}}(\mathbf{x})$.⁵ This is true for any initial state $\hat{\rho}_{\psi, A}$ for the restricted system containing the modes in A and the target field. The effect of the existence of the additional modes of the probe field that have been integrated and traced out amounts to a correction that is fully encoded by the vacuum influence phase $\tilde{S}[\psi, \psi']$ for the modes not included in A . Note, however, that this correction is of higher order in the coupling strength λ , as made evident by Equation (3.53): whereas the direct coupling between the modes in A and the target field is of order λ , the net effect of the extra modes that were traced out amounts to an additional term that is of order λ^2 in the action. This means that, at lowest order in perturbation theory, if we only have access to a finite set of modes of the probe field, the physics is well-reproduced by that of a finite number of harmonic-oscillator degrees of freedom coupled to a finite set of smeared field operators.

These results reinforce the analysis of [137], where it was shown that, at leading order in the coupling constant and for the case where the probe field starts in its ground state, the final state of any individual mode of a full-blown field theory after all extra modes are traced out matches that of a harmonic-oscillator UDW detector. Therefore, in some sense, every statement about the dynamics of a harmonic-oscillator UDW detector can be understood as a statement about a mode of a localized quantum field. The analysis presented here vindicates and significantly extends the range of validity of this conclusion, by demonstrating that the statement remains true for an arbitrary number of modes included as detector degrees of freedom, and for a much more general set of initial states of the form (3.43). Furthermore, the full expression (3.51), including the factor of $e^{i\tilde{S}[\psi, \psi']}$, also

⁵In connection to the discussion after Equation (2.4) in Chapter 2 about the possibly distinct mechanisms behind the nature of the smearing, it should be clear that the effective delocalization obtained through this field-theoretic approach is more closely related to the kind of delocalization stemming from the quantum uncertainty in the location of the probe, and not any fundamental size; after all, the interactions at the level of the fields themselves are manifestly local in spacetime.

provides a way to (in principle) systematically compute the corrections that arise at higher orders in λ from the fact that the finite set of modes fundamentally emerges from a local field theory with infinitely many degrees of freedom.

It is also clear that even the assumption that the inaccessible modes start out in the vacuum state is not essential for the conclusions above. If we had instead assumed another state $\hat{\rho}_{\bar{A}}$ for these modes, the only thing that would change is that $e^{i\tilde{S}[\psi,\psi']}$ in Equation (3.51) would be the Schwinger-Keldysh path integral with sources given by $\lambda\psi_{\mathbf{n}}$ for the modes in $\bar{A}$, now calculated with a different input state. For a completely general $\hat{\rho}_{\bar{A}}$, this may not factorize as a product of Schwinger-Keldysh path integrals for each mode in $\bar{A}$ as it did in Equation (3.52); however, as long as the expectation value of the amplitudes $\hat{\phi}_{\mathbf{n}}$ for $\mathbf{n} \notin A$ vanish in the input state $\hat{\rho}_{\bar{A}}$, the lowest-order term in the perturbative expansion of $\tilde{S}[\psi,\psi']$ will be quadratic in the sources ψ,ψ' , and the lowest power of the coupling constant λ will thus be λ^2 . As such, the equivalence at leading order in perturbation theory between the dynamics of a finite number of modes of the field and the dynamics of the same number of harmonic-oscillator UDW detectors will still hold.

3.3 Generalizations to multiple trajectories

So far we have shown how to relate a description of a probe system in terms of a localized quantum field to an alternative description given by a finite number of UDW detectors following a well-defined classical trajectory in spacetime. The classical trajectory that we assign to the detectors, in turn, is directly associated to the spatial profile provided by the mode functions $v_{\mathbf{n}}(\mathbf{x})$. When $U(\mathbf{x})$ only has one minimum at each spatial slice of constant Killing time, it is very natural to interpret the “trajectory of the detector” as being simply given by the location of the minimum of the potential; after all, this will also typically be the location around which the mode functions $v_{\mathbf{n}}(\mathbf{x})$ will be mostly peaked. However, it is also possible to imagine regimes where it would be more natural to picture the potential as “localizing the field” in more than one spatial region—as would happen, for instance, if the confining potential actually has multiple minima that are sufficiently far apart. In this section we will briefly describe how to adapt the story from Secs. 3.1 and 3.2 to this case, and comment on a few simple physical setups where this generalization can be useful.

3.3.1 General strategy

For convenience, we will explicitly describe how the strategy of Section 3.2 can be carried over to the case where the potential generates localized profiles in two distinct regions; the

generalization to more than two localized regions will be evident.

The setup is then the following. We take the potential $U(\mathbf{x})$ to have two distinct minima, found in two non-overlapping spatial regions R_A and R_B in each surface of constant Killing time, with $U(\mathbf{x}) \rightarrow +\infty$ as we move away from both regions. We assume that, in R_A and R_B , the potential can be approximated by $U_A(\mathbf{x})$ and $U_B(\mathbf{x})$ respectively, where $U_{A, B}(\mathbf{x})$ are both confining potentials in their own right. We can then formally write down a mode expansion for any field configuration for the probe as

$$\phi_D(t, \mathbf{x}) = \sum_{\mathbf{n}_A} \phi_{\mathbf{n}_A}^{(A)}(t) v_{\mathbf{n}_A}^{(A)}(\mathbf{x}) + \sum_{\mathbf{n}_B} \phi_{\mathbf{n}_B}^{(B)}(t) v_{\mathbf{n}_B}^{(B)}(\mathbf{x}), \quad (3.58)$$

where $v_{\mathbf{n}_{A, B}}^{(A, B)}(\mathbf{x})$ are mode functions corresponding to spatially localized profiles associated to the potentials $U_{A, B}(\mathbf{x})$ and centered in regions $R_{A, B}$ respectively. We will also assume that the potential barrier and the spatial separation between the regions R_A and R_B under consideration are large enough for any two modes associated to two different regions to approximately not have any overlap; i.e.,

$$\int d^d \mathbf{x} \frac{\sqrt{h}}{N} v_{\mathbf{n}_A}^{(A)}(\mathbf{x}) v_{\mathbf{n}_B}^{(B)}(\mathbf{x}) \simeq 0 \quad \forall \mathbf{n}_A, \mathbf{n}_B. \quad (3.59)$$

Under this assumption, the expansion of the free action for the probe field then proceeds in exactly the same way as in the previous section. The exact same steps from Section 3.1 then lead to the probe action

$$\begin{aligned} S_D \simeq & \sum_{\mathbf{n}_A} \frac{1}{2} \int dt \left[\left(\frac{d\phi_{\mathbf{n}_A}^{(A)}}{dt} \right)^2 - \omega_{\mathbf{n}_A}^2 \phi_{\mathbf{n}_A}^{(A)2} \right] \\ & + \sum_{\mathbf{n}_B} \frac{1}{2} \int dt \left[\left(\frac{d\phi_{\mathbf{n}_B}^{(B)}}{dt} \right)^2 - \omega_{\mathbf{n}_B}^2 \phi_{\mathbf{n}_B}^{(B)2} \right]. \end{aligned} \quad (3.60)$$

In other words, the probe splits into two decoupled series of harmonic oscillators localized around two distinct regions of space. The quantization of the probe field then proceeds in a manner that is mathematically identical to what one would do in the case of two distinct fields, each under the influence of a separate confining potential.

If we do a very similar analysis of this setup using canonical quantization, we find that (3.59) can also be interpreted as the requirement that all the quadrature operators of the modes supported in each separate region satisfy canonical commutation relations

(see Appendix A.2). This is consistent with the interpretation just described, in which the no-overlap assumption (3.59) corresponds to the case where the localized quantum field can be decomposed in two independent fields that are localized around the distinct minima of the confining potential.

As a limiting case of this picture, one could consider a situation where the regions of interest R_A and R_B act like Dirichlet cavities, with the potential being set to zero in either region, and equal to infinity everywhere else. In this limiting case, the condition (3.59) holds exactly, since the spatial profiles are identically vanishing outside each of the cavities. If, on the other hand, the condition (3.59) is not met, then the approximation of two independent sets of modes localized in either R_A or R_B is not good; instead, the truly independent modes will be best described as judiciously chosen linear combinations of $v_{n_A}(\mathbf{x})$ and $v_{n_B}(\mathbf{x})$ which define a set of mode functions that do satisfy the orthogonality condition (3.10).

Now that we know that the free action of the probe field splits as in (3.60) under the no-overlap assumption (3.59), we can apply exactly the same logic from Section 3.2. In this case, we can have detectors at two fixed spatial positions $\mathbf{x}_A$ and $\mathbf{x}_B$, corresponding to the position of the minimum of the potential that each respective field mode is associated to. By tracing out all but finitely many modes in both of the regions where the field is localized, the exact same steps from Subsection 3.2.2 will then give us, at leading order in perturbation theory, the same dynamics as that of a finite number of harmonic-oscillator UDW detectors supported in each individual region. If, for instance, we keep only one mode in each region (say, the lowest-frequency mode in both R_A and R_B), the physics at leading order in the coupling constant is exactly the same as that of the following action,

$$\begin{aligned}
S[\phi_{\mathbf{0}}^{(A)}, \phi_{\mathbf{0}}^{(B)}, \psi] = & \frac{1}{2} \int dt \left[\left(\frac{d\phi_{\mathbf{0}}^{(A)}}{dt} \right)^2 - \omega_{A,0}^2 \phi_{\mathbf{0}}^{(A)2} \right] + \frac{1}{2} \int dt \left[\left(\frac{d\phi_{\mathbf{0}}^{(B)}}{dt} \right)^2 - \omega_{B,0}^2 \phi_{\mathbf{0}}^{(B)2} \right] \\
& - \frac{1}{2} \int dV (g^{ab} \nabla_a \psi \nabla_b \psi + M^2 \psi^2 + \xi R \psi^2) \\
& + \lambda \int dt \left(\phi_{\mathbf{0}}^{(A)}(t) \psi_A(t) + \phi_{\mathbf{0}}^{(B)}(t) \psi_B(t) \right)
\end{aligned} \tag{3.61}$$

where, as before,

$$\psi_A(t) = \int d^n \mathbf{x} N \sqrt{h} \zeta(t, \mathbf{x}) v_{\mathbf{0}}^{(A)}(\mathbf{x}) \psi(t, \mathbf{x}), \tag{3.62}$$

$$\psi_B(t) = \int d^n \mathbf{x} N \sqrt{h} \zeta(t, \mathbf{x}) v_{\mathbf{0}}^{(B)}(\mathbf{x}) \psi(t, \mathbf{x}). \tag{3.63}$$

With that, Equation (3.61) simply becomes the theory of two static harmonic-oscillator UDW detectors at the positions $\mathbf{x}_{A,B}$, coupled to a quantum field with the spacetime smearing functions $\zeta(\mathbf{x})v_{\mathbf{0}}^{(A,B)}(\mathbf{x})$.

A minor detail to keep in mind now is that, in general, it is not possible to normalize the Killing vector field χ^a such that $\chi^a\chi_a = -1$ in both $\mathbf{x}_A$ and $\mathbf{x}_B$. Instead, the best we can do in this case is to relate the proper time of each detector to the global Killing time t by $\tau_A = N(\mathbf{x}_A)t$ and $\tau_B = N(\mathbf{x}_B)t$, where $N(\mathbf{x})$ is the lapse function from the general expression of the metric in static coordinates (3.2). Similarly, the characteristic frequencies Ω_A and Ω_B of each oscillator in its respective proper frame are related to the Killing frequencies $\omega_{A,0}$ and $\omega_{B,0}$ in Equation (3.61) by $N(\mathbf{x}_A)\Omega_A = \omega_{A,0}$ and $N(\mathbf{x}_B)\Omega_B = \omega_{B,0}$. These modifications, plus an additional rescaling of the field amplitude as $\sqrt{N(\mathbf{x}_A)}\phi_{\mathbf{0}}^{(A)} \mapsto \phi_{\mathbf{0}}^{(A)}$ and $\sqrt{N(\mathbf{x}_B)}\phi_{\mathbf{0}}^{(B)} \mapsto \phi_{\mathbf{0}}^{(B)}$, leave the free action of the detectors in (3.61) in exactly the form of two harmonic-oscillator UDW detectors as written in Equation (2.12), when parametrized in terms of each detector's proper time.⁶

3.3.2 Examples

We now briefly comment on simple examples where the general framework described in Subsection 3.3.1 may be of physical relevance.

Entanglement harvesting

Entanglement harvesting is a very well-known protocol in RQI by which two localized systems interacting with a quantum field can become entangled even when they only couple to the field in spacelike-separated regions [150, 151, 144]. The idea is that entanglement can be shared between the detectors even before they are able to causally communicate with each other, thanks to the pre-existing entanglement between subregions in the field theory [169, 197].

The most common setup for studying entanglement harvesting consists of a minimal modification of the one described in Section 2.1. Now, instead of having just one detector, we can have two detectors (denoted from now on by A and B) which are initialized in a completely uncorrelated state,

$$\hat{\rho}_{AB,0} = \hat{\rho}_{A,0} \otimes \hat{\rho}_{B,0}, \quad (3.64)$$

⁶Note, however, that this rescaling of the amplitude depending on the lapse function on each detector's position will lead to a difference in the effective coupling constants in the interaction action in the last line of Equation (3.61).

and are then put to evolve as in Equation (2.9),

$$\hat{\rho}_{AB} = \text{Tr}_F \left[\hat{\mathcal{U}}_I (\hat{\rho}_F \otimes \hat{\rho}_{AB,0}) \hat{\mathcal{U}}_I^\dagger \right], \quad (3.65)$$

where the time evolution operator $\hat{\mathcal{U}}_I$ now contains interaction terms corresponding to both detectors. It is usual to take $\hat{\rho}_{A,0}$ and $\hat{\rho}_{B,0}$ to be pure states (typically the ground state of each detector), as we know that initial mixedness in either of the probes generally hinders harvesting [154, 86]. We then study how much entanglement the detectors have acquired, by evaluating the entanglement between the subsystems A and B in the final state $\hat{\rho}_{AB}$. This, in turn, will usually require the use of mixed-state measures of entanglement—after all, time evolution will also generically entangle the detectors with the field, and therefore the final state $\hat{\rho}_{AB}$ will typically be mixed.

It is easy to see how the general framework of Subsection 3.3.1 can be directly applied to entanglement harvesting. In this case, the two parties that we are trying to entangle are both fundamentally derived from the same underlying quantum field. This mimics a setup where the systems being used as detectors consist of identical elementary particles—for instance, two electrons probing the electromagnetic field in two distinct positions. Since the free action of the probe field splits as in (3.60), the vacuum state of the probe field can be taken as a tensor product of the ground states associated to each localization region separately. In particular, if we take only the modes of lowest frequency in both localized regions considered in Subsection 3.3.1, the initial state of the probes is simply $\hat{\rho}_{AB,0} = |0_A\rangle \langle 0_A| \otimes |0_B\rangle \langle 0_B|$. The entanglement between the two modes after interacting with the target field can then be quantified by the usual measures adopted in entanglement harvesting—most commonly, the negativity—in the final state $\hat{\rho}_{AB}$.

This shows how one can conceptualize entanglement harvesting setups where each particle detector is obtained as a different localized mode of one single field. Of course, an even more direct setting for the study of entanglement harvesting using quantum fields as probes is obtained by modeling the setup with two *different* localized quantum fields, each of them with one single minimum. In this case, the global Hilbert space of the joint system including both probes is guaranteed to factorize by construction, and the initial state of both probe fields on the vacuum is guaranteed to factorize between modes supported in either one of the local regions. In Part II of this thesis, we will revisit this point, and present more details on the effect of explicitly using modes from relativistic field theories to harvest entanglement from a quantum field.

Detector in superpositions of trajectories

Another physically interesting scenario is achieved by simply looking at the action (3.61) from a different point of view. Namely, we could imagine that (3.61) describes the dynamics of the target field and *one* single detector, the difference being that now the detector consists of two bosonic modes instead of one.

If we adopt the perspective in which the two modes in R_A and R_B are taken as a single detector with two degrees of freedom, the probe is now associated to two possible trajectories, centered around the points $\mathbf{x}_A$ and $\mathbf{x}_B$. The interaction with the field will then promote transitions between states that are localized around either one of the trajectories, and the final state of the detector will generically contain coherences that connect one trajectory to the other. In particular, one could for instance prepare an initial state of the probe which takes the form

$$\hat{\rho}_{AB,0} = |\psi\rangle\langle\psi|, \quad (3.66)$$

$$|\psi\rangle = \frac{1}{\sqrt{2}} (|1_A\rangle \otimes |0_B\rangle + |0_A\rangle \otimes |1_B\rangle), \quad (3.67)$$

where we define

$$|1_A\rangle = \hat{a}_A^\dagger |0_A\rangle, \quad |1_B\rangle = \hat{a}_B^\dagger |0_B\rangle \quad (3.68)$$

and $\hat{a}_{A,B}^\dagger$ are the creation operators for the lowest-frequency mode in regions $R_{A,B}$. This describes an initial superposition between a one-particle excitation that is localized around position $\mathbf{x}_A$ and another one-particle excitation localized around position $\mathbf{x}_B$. The final state of the probe will then involve contributions that depend on the dynamics of each degree of freedom along one single trajectory, as well as correlations that connect the degrees of freedom of the detector in both possible positions.

This shares the same spirit of other proposals found in the literature for how to treat delocalized detectors which can evolve in superpositions of trajectories [76, 12]; the details of the model, however, are different. In the future, it might be interesting to study how this model compares with other setups previously considered, and further investigate whether there are cases where a model for a delocalized center of mass based on localized quantum fields (along the lines of what was described here) is physically reasonable.

It is also interesting to note how the physical distinction between this example (considering superpositions of trajectories) and the previous one (which is naturally set up for entanglement harvesting) is, in some sense, subjective/context-dependent: roughly speaking, it depends on whether we postulate the existence of an external agent that can have global control over both modes in R_A and R_B , or we assume that only operations that act separately on R_A or R_B are allowed. Indeed, the state (3.67) is an entangled state between the parties at R_A and

R_B , so it certainly cannot be engineered by only applying local operations between the two localized regions. The framework of Section 3.3.1, however, captures both physical scenarios at once; which of them is better applicable to a given setup at hand will be determined by other constraints of the problem being treated.

3.4 Concluding remarks

We have demonstrated in this chapter that it is possible to draw a very clear connection between fully relativistic models of local probes and more usual, effectively nonrelativistic particle detector models commonly used in RQI. We explicitly verified this fact in a simple model with a linear scalar field that is effectively confined to a region of space by the action of an external potential. A series of assumptions made (namely, that the free action for the probe field was quadratic, the background metric was static, and the external potential that confined the probe was invariant under the flow of the background timelike Killing vector field) allowed us to decompose the probe field in terms of a discrete tower of decoupled modes. By tracing out all but finitely many such modes, we were able to restrict the probe field to a finite number of degrees of freedom. Doing this at the level of the Schwinger-Keldysh path integral equipped us with an analytical expression for the effective dynamics of the system containing the quantum field of interest and any finite number of modes of the probe.

The mathematical steps performed in Subsection 3.2.2 are independent of how we choose which modes of the field should be kept as UDW detectors. For the sake of concreteness, however, a physically motivated choice would be to only include those modes whose frequencies ω_n are below a certain cutoff, $\omega_n \leq \Omega$. This gives an effective-field-theory flavor to our derivation, with the procedure performed in Subsection 3.2.2 being very close in spirit to the Wilsonian approach to the renormalization group: in this case, the frequency Ω plays the role of a UV cutoff, and the action of the finite number of harmonic-oscillator UDW detectors coupled to the field, together with the additional factor from integrating out the rest of the modes, behaves as a low-energy effective action. This is in tune with the point of view that says particle detectors should be seen as effective descriptions of systems that are fundamentally relativistic, and provides a very simple toy example where the connection between this point of view and the more technical aspects of the formalism of effective field theories can be made explicit and concrete.

This comparison to effective field theories also sheds light on some of the locality issues of particle detectors which were mentioned in Subsection 2.2. From the literature on Wilsonian renormalization, we know that integrating out high energy modes can introduce

mild non-localities in the effective low-energy action, where the typical scale of the non-localities is related to the inverse of the cutoff in energy or momentum [142, 37]. As we discussed in Subsection 2.2, coupling one single degree of freedom of a detector to a smeared (non-pointlike) field operator also generally leads to violations of locality and microcausality at distance scales that are controlled by the size of the spatial smearing. Obtaining the detector model from a field theory through coarse-graining-like steps as illustrated above draws a direct connection between these two statements, and also shows how to reconcile this (only apparent) conflict between smeared particle detector models and local field theory. From this point of view, the non-locality that emerges in the detector perspective is nothing but an artifact of the coarse-graining of the probe field to a finite number of modes; by including arbitrarily many modes, of arbitrarily high frequencies, one can restore a full-fledged local field theory.

The careful reader may also have noticed that, although we started out with the target field ψ being given by a free Klein-Gordon field with action (3.14) and we coupled the probe field linearly with ψ through the interaction (3.15), none of the manipulations that led to the main result in Equation (3.51) actually depended on these two assumptions. All the steps really only depended on the free action of the *probe* being quadratic. The reason why we avoided a greater level of generality is that adding self-interactions and couplings to nonlinear observables of the target field may introduce additional complications related to renormalization and quantum corrections to the action. Taking the total action $S[\phi_D, \psi]$ to be quadratic in all the fundamental fields appearing in the path integral, on the other hand, allows us to not worry about these subtleties. Having said that, we find it perfectly reasonable to expect that, once these potential issues are carefully taken into account, the same logic presented here will still apply—namely, by tracing out the inaccessible degrees of freedom of the detector, the theory of a probe field coupled to any (possibly composite) operator, of any (possibly interacting) field theory, can be reduced at leading order to a theory of a finite number of particle detectors coupled to the target field.

In terms of future directions, we believe that this work opens the door to deeper investigations connecting UDW-like detectors and other fully field-theoretic models for local probes. It would be very interesting to see whether a more systematic use of tools from effective field theory (along the lines of the effective field theory approach to the Born-Oppenheimer approximation in molecular physics [24], for example) could be used to construct more realistic models than the one explored here, using simple solvable examples of truly interacting QFTs with features such as bound states.

Part II

Entanglement in QFT

Chapter 4

Motivation: why entanglement in QFT is hard

Part II the thesis will be devoted to investigations on *entanglement* in QFT. In order to prepare for that, in this chapter we will give a lightning review of how entanglement is usually treated in quantum mechanics, and explain why there is an important obstacle we need to circumvent when applying notions of entanglement in the case of quantum field theories. This will provide some motivation for the approach that we will adopt in Chapters 5 and 6, and also help to contextualize our work with the broader literature on the subject.

Note of clarification. This chapter (which only contains a literature review) was first written by the author as part of this thesis. However, a version of it will also appear in an upcoming paper [5]. We chose not to include [5] in the statement of contributions because that work is still in progress, and none of the original results presented there have been included here.

4.1 Entanglement in quantum mechanics

Entanglement is one of the key distinguishing features of quantum theory, playing an important role in various disparate areas of physics. In traditional quantum information theory, entanglement has been shown to be crucial for improving quantum communication protocols and quantum cryptography, and is also responsible for the computational advantage of certain quantum algorithms relative to their classical counterparts [185, 68, 67].

In theoretical physics more broadly, entanglement has also become a central aspect in the connection between quantum information, QFT, and gravity [133, 197]. Examples range from explorations of spacetime structure through entanglement in Relativistic Quantum Information [121, 126], to a microscopic origin of black hole entropy [19, 166, 163], culminating in the idea that entanglement may play a fundamental role in the emergence of gravitational dynamics both in the context of the AdS/CFT correspondence [187, 70, 113] and in quantum gravity more generally [18, 104].

A typical introduction to the concept of entanglement in quantum mechanics goes as follows.¹ Consider two physical systems A and B , with respective Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$. The Hilbert space of the joint system AB is thus defined as the tensor product $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$. Having this decomposition of the joint Hilbert space in mind, a pure quantum state $|\psi\rangle \in \mathcal{H}$ is said to be *entangled* if it is not *separable*—i.e., if it cannot be written as $|\psi\rangle = |\chi\rangle \otimes |\phi\rangle$ for some $|\chi\rangle \in \mathcal{H}_A$, $|\phi\rangle \in \mathcal{H}_B$. The entanglement between subsystems A and B can then be directly quantified by the *entanglement entropy*, which is given by the von Neumann entropy of the reduced state of either A or B ,

$$E(|\psi\rangle) := S(\hat{\rho}_A) = -\text{Tr}(\hat{\rho}_A \log \hat{\rho}_A) = -\text{Tr}(\hat{\rho}_B \log \hat{\rho}_B) = S(\hat{\rho}_B), \quad (4.1)$$

where $\hat{\rho}_A = \text{Tr}_B(|\psi\rangle\langle\psi|)$ and $\hat{\rho}_B = \text{Tr}_A(|\psi\rangle\langle\psi|)$. That the von Neumann entropies $S(\hat{\rho}_A)$ and $S(\hat{\rho}_B)$ of both reduced density matrices $\hat{\rho}_A$ and $\hat{\rho}_B$ are the same is guaranteed by the fact that the joint state $\hat{\rho}_{AB} = |\psi\rangle\langle\psi|$ in this case is pure.

The paradigmatic example of an entangled state is obtained by taking A and B to consist of two-level systems (i.e., qubits), and considering the state

$$|\text{EPR}\rangle = \frac{1}{\sqrt{2}} (|0\rangle_A \otimes |0\rangle_B + |1\rangle_A \otimes |1\rangle_B). \quad (4.2)$$

This state is usually referred to as an *EPR pair* [130]. If we trace out either one of the qubits in the state $|\text{EPR}\rangle$, the reduced state on the qubit that is left is proportional to the identity operator,

$$\hat{\rho}_B = \text{Tr}_A(|\text{EPR}\rangle\langle\text{EPR}|) = \frac{1}{2} (|0\rangle_B \langle 0|_B + |1\rangle_B \langle 1|_B), \quad (4.3)$$

which is the most mixed state—i.e., the state of maximum entropy—for a two-level system. An EPR pair thus corresponds to a maximally entangled state of two qubits. In much

¹This chapter will only focus on introducing the concepts in the theory of quantum entanglement that will be of most relevance for our later chapters. For a much more thorough review on the subject, see e.g. [99].

the same way as qubits are the “elementary unit” of information processing capability in quantum information theory, EPR pairs constitute the gold standard for an elementary unit of “perfect entanglement”, and are commonly used as references in terms of which entanglement in more complicated systems is broadly understood.

It is in general also important to have a notion of entanglement when the joint state of the quantum system AB is not pure. This is relevant, for instance, when AB is itself a subsystem of a larger subsystem ABC , one of whose components (here denoted by C) can be entangled with AB , thus resulting in a mixed state $\hat{\rho}_{AB}$ for the quantum system of interest. Entanglement in mixed states is again defined as lack of separability, but the notion of “separability” is slightly modified: now, a mixed state $\hat{\rho}$ is considered separable not only if it is given by a product $\hat{\rho}_A \otimes \hat{\rho}_B$ for a pair of states $\hat{\rho}_A$ and $\hat{\rho}_B$, but also if it can be expressed as a convex linear combination of states of this form—i.e., if there is a probability vector $\{p_i\}$ and an ensemble of states $\{\hat{\rho}_A^{(i)}\}$ and $\{\hat{\rho}_B^{(i)}\}$ such that $\hat{\rho} = \sum_i p_i \hat{\rho}_A^{(i)} \otimes \hat{\rho}_B^{(i)}$. The idea is that, once we allow for noisy (i.e., non-unitary) quantum channels acting on the system AB , it is generically possible to start in a product state $\hat{\rho}_A \otimes \hat{\rho}_B$ and then produce nontrivial convex combinations of distinct product states simply by applying local operations (i.e., channels that act separately on subsystems A and B), and sending rounds of classical communication between agents with access to A and B . But local operations and classical communication (LOCC) should not be able to create anything beyond classical correlations; therefore, if non-separability is to constitute some form of genuinely nonclassical correlation that goes beyond what can be generated classically, states of the form $\sum_i p_i \hat{\rho}_A^{(i)} \otimes \hat{\rho}_B^{(i)}$ must also be defined as separable.

The need to differentiate between classical correlations (which are generically present in mixed states) and quantum entanglement makes quantifying entanglement in mixed states much trickier. Well-known examples of mixed-state entanglement measures include the *entanglement cost* (defined as the asymptotic rate of EPR pairs that one needs to be provided in order to generate N identical copies of the state, $\hat{\rho}^{\otimes N}$, in the limit of $N \rightarrow +\infty$, using only LOCC) and the *distillable entanglement* (defined as the asymptotic rate of EPR pairs that can be *extracted* from $\hat{\rho}^{\otimes N}$ in the limit $N \rightarrow \infty$, again only using LOCC). Both of these quantities reduce to the entanglement entropy if $\hat{\rho}$ is pure, and have a clear operational meaning as measuring the nonclassical resources—namely, the amount of “pure entanglement” in the form of EPR pairs—which are needed to synthesize the state (in the case of the entanglement cost) or which can be “mined” from the state (in the case of the distillable entanglement). Unfortunately, however, the definitions of these quantities do not provide an efficient way of calculating them, and computing them explicitly is prohibitively hard in general. This, in fact, is true for most known entanglement measures; and at least in part for this reason, characterizing entanglement in full generality for mixed states is

notoriously difficult.²

Fortunately, there is a simple quantity that can be very easily computed for any mixed state $\hat{\rho}_{AB}$, and which turns out to be quite useful (even if not perfect) in the study of mixed-state entanglement. This quantity is called the *negativity* [188], and is defined as

$$\mathcal{N}(\hat{\rho}_{AB}) := \frac{\|\hat{\rho}_{AB}^{T_B}\| - 1}{2}, \quad (4.4)$$

where $\|\hat{O}\| \equiv \text{Tr} \sqrt{\hat{O}^\dagger \hat{O}}$ is the 1-norm of the operator $\hat{O}$, and $\hat{\rho}_{AB}^{T_B}$ is the *partial transpose* of the density matrix $\hat{\rho}_{AB}$ with respect to subsystem B .

Partial transposition is a linear operation whose action on a given orthonormal basis on the space of linear operators on $\mathcal{H}_A \otimes \mathcal{H}_B$ is given by

$$T_B \left((|i\rangle_A \langle k|_A) \otimes (|j\rangle_B \langle l|_B) \right) = (|i\rangle_A \langle k|_A) \otimes (|l\rangle_B \langle j|_B). \quad (4.5)$$

Therefore, $\hat{\rho}_{AB}^{T_B}$ corresponds to the operator obtained by only taking the transpose of $\hat{\rho}_{AB}$ with respect to the “indices relative to subsystem B ” in some orthonormal basis. By noting that $\text{Tr}(\hat{\rho}_{AB}^{T_B}) = \text{Tr} \hat{\rho}_{AB} = 1$, it is easy to see that the negativity (4.4) is nothing but the absolute value of the sum of all negative eigenvalues of $\hat{\rho}_{AB}^{T_B}$. The rationale behind negativity as a measure of entanglement can be explained in many levels. First, since the transpose of any positive operator is still positive, the partial transpose of any separable density matrix will still be a density matrix, and thus only have positive eigenvalues; therefore, nonzero negativity is a sufficient condition to *detect* entanglement [139]. There is also a sense in which negativity quantifies a truly nonclassical resource contained in a state, as it can be shown that the negativity can never increase under local operations and classical communication [188]; in the entanglement theory jargon, this property implies that the negativity constitutes an *entanglement monotone*. In fact, there is a closely related quantity to the negativity, known as the *logarithmic negativity* and given by

$$E_{\mathcal{N}}(\hat{\rho}_{AB}) := \log(2\mathcal{N}(\hat{\rho}_{AB}) + 1) = \log \|\hat{\rho}_{AB}^{T_B}\|, \quad (4.6)$$

which provides an *upper bound* to the distillable entanglement of the state $\hat{\rho}_{AB}$ [188].

The fact that the logarithmic negativity (4.6) provides a bound to the distillable entanglement directly implies that states with vanishing negativity—also called PPT (or

²In fact, the very problem of deciding whether a given density matrix of a bipartite quantum system is separable or entangled is known to be an NP-hard problem in general [89]!

positive-partial-transpose) states—have zero distillable entanglement, and are therefore useless as sources of EPR-type correlations. Now, it is possible to have states that are non-separable (and have, for instance, nonzero entanglement cost) and also have a positive partial transpose; when this happens, we say that the state only contains *bound* entanglement [97, 33]. The existence of bound-entangled states means that the negativity is in general not a completely infallible measure of non-separability, as it only captures the entanglement that can be distilled from a given state. However, for certain systems, it can be shown that nonzero negativity is actually not just sufficient but also *necessary* for non-separability. This happens for two-qubit and qubit-qutrit states [98], as well as for Gaussian states of continuous-variable systems when the bipartition consists of one single mode for subsystem A and an arbitrary number of modes for subsystem B [162, 194], or in the case of Gaussian states of $M \times N$ modes which are *bisymmetric* [159] (where here “bisymmetric” means that the state is invariant under arbitrary internal permutations of the modes associated to each subsystem in the bipartition separately). In all of these cases, it can be shown that bound entanglement cannot exist, and therefore the negativity provides a faithful entanglement measure.

We thus find ourselves in a position where we seem to have all the tools we need to talk about entanglement in any quantum system, both for pure and mixed states. This task, however, turns out to be quite subtle, both on a technical and a conceptual level, when we try to apply these tools to quantum field theories. In the next subsection, we will explain what makes entanglement in QFT qualitatively different.

4.2 What makes QFTs different

A QFT in Minkowski spacetime $\mathcal{M}$ can be formulated the following way.³ First, we start with the definition of an *algebra of observables* $\mathcal{A}(\mathcal{M})$ on the whole spacetime, together with a collection of subalgebras $\mathcal{A}(R)$ associated to suitable spacetime regions $R \subset \mathcal{M}$. The intuitive idea is that, for a given spacetime subregion R , the algebra $\mathcal{A}(R)$ comprises the representation of the field degrees of freedom that are supported in R . Then, the Hilbert space $\mathcal{H}_{\text{QFT}}$ can be built by starting with the vacuum state $|\Omega\rangle$ —i.e., the unique

³Note that, unlike in part I of the thesis—where we used $\mathcal{M}$ to denote the differentiable manifold associated to a general (possibly curved) spacetime—, from here on out we will assume that $\mathcal{M}$ denotes Minkowski spacetime. This is mostly for the sake of simplicity and conciseness, as the careful formulation of QFTs in an arbitrary globally hyperbolic spacetime will not be strictly needed in the rest of this thesis. However, the formalism we are about to cover can be restated and generalized in any (possibly curved) globally hyperbolic spacetime, as reviewed e.g. in [190].

Poincaré-invariant state of minimum energy according to any inertial observer—, and generating all other states by limits of sequences of vectors obtained by acting on $|\Omega\rangle$ with elements of the global algebra $\mathcal{A}(\mathcal{M})$.

In the case of a real scalar field, for instance, the algebra is generated by the identity element $\hat{1}$ and the smeared field operators of the form

$$\hat{\varphi}(f) = \int_{\mathcal{M}} dV \hat{\varphi}(\mathbf{x}) f(\mathbf{x}), \quad (4.7)$$

with $f \in C_c^\infty(\mathcal{M})$ a smooth, compactly supported function from $\mathcal{M}$ to the complex numbers. In this context, the object $\hat{\varphi}(\mathbf{x})$ is interpreted as an operator-valued distribution, with $\hat{\varphi}(f)$ denoting the output of that distribution when acting on the function $f(\mathbf{x})$. Having constructed this algebra, the Hilbert space $\mathcal{H}_{\text{QFT}}$ will be generated by states of the form

$$|\chi\rangle = \hat{\varphi}(f_1) \dots \hat{\varphi}(f_n) |\Omega\rangle, \quad (4.8)$$

for $f_1, \dots, f_n \in C_c^\infty(\mathcal{M})$.⁴

If the field in question is a Klein-Gordon field, we can also make use of the Klein-Gordon equation of motion

$$\begin{aligned} \hat{\varphi}(h) &= 0 \text{ whenever } h(\mathbf{x}) \text{ is of the form} \\ h &= (g^{ab} \nabla_a \nabla_b - m^2) f, \quad f \in C_c^\infty(\mathcal{M}). \end{aligned} \quad (4.9)$$

This will allow us to give an equivalent expression to Equation (4.7) as

$$\hat{\varphi}(f) = \int_{\Sigma} d\Sigma n^a (\hat{\varphi} \nabla_a F - F \nabla_a \hat{\varphi}), \quad (4.10)$$

where Σ is any spacelike Cauchy surface for $\mathcal{M}$, $d\Sigma$ is the induced volume form on Σ , n^a is the future-oriented unit normal to Σ , and $F(\mathbf{x})$ is a function defined by

$$F(\mathbf{x}) := (Ef)(\mathbf{x}) = \int_{\mathcal{M}} dV' E(\mathbf{x}, \mathbf{x}') f(\mathbf{x}'). \quad (4.11)$$

In particular, one can take the Cauchy surface to be defined by a constant value of the

⁴A small technical detail: for a general QFT, it is not exactly true that all the vectors in $\mathcal{H}_{\text{QFT}}$ can be obtained by acting on the vacuum with local operators. This is true in the case of a free scalar field, which is going to be our working example here. However, in more general scenarios, this prescription only describes what is known as the *vacuum sector* of the QFT, which is in general smaller—for instance, when there are globally conserved charges that are not carried by any local operator, thus leading to the emergence of superselection sectors.

time coordinate t , in some Cartesian coordinate system $(t, \mathbf{x})$. In this case, we simply have $n^a \nabla_a \hat{\varphi} = \partial_t \hat{\varphi}$, and we get

$$\hat{\varphi}(f) = \int_{\Sigma} d^d x \, (\hat{\varphi} \partial_t F - F \partial_t \hat{\varphi}). \quad (4.12)$$

This expresses the familiar fact that in Minkowski space, the full algebra of observables $\mathcal{A}(\mathcal{M})$ in the entire spacetime can be constructed out of the field and its conjugate momentum $\partial_t \hat{\varphi}$ at a constant time slice.

The construction of $\mathcal{H}_{\text{QFT}}$ we described above is global, as it starts with a global vacuum state $|\Omega\rangle$ defined on a complete Cauchy slice, and builds the rest of the states in the Hilbert space by acting with all local operators that can be defined in any subregion of spacetime. To recover something that may look more similar to a decomposition of the field in terms of a collection of subsystems, one may try to proceed as follows. Given a spatial subregion $\Sigma_A \subset \Sigma$ of a constant time slice Σ , one can talk about the subset of field observables that are “supported in Σ_A ” by restricting to field operators $\hat{\varphi}(f)$ such that both $F|_{\Sigma}$ and $n^a \nabla_a F|_{\Sigma}$ appearing in Equation (4.10) are only nonvanishing in Σ_A . It would then perhaps be natural to think that the global Hilbert space $\mathcal{H}_{\text{QFT}}$ could be decomposed as something of the form $\mathcal{H}_A \otimes \mathcal{H}_{\bar{A}}$, where $\mathcal{H}_A$ and $\mathcal{H}_{\bar{A}}$ are local Hilbert spaces associated to the spatial regions Σ_A and $\Sigma_{\bar{A}} = \Sigma \setminus \Sigma_A$ respectively; and if that were the case, one would in principle be able to talk about entanglement measures in QFT by using the same tools from the nonrelativistic setting. For instance, the entanglement between the two complementary regions Σ_A and $\Sigma_{\bar{A}}$ in the vacuum of the QFT could be obtained by directly computing the von Neumann entropy of the reduced density matrix $\hat{\rho}_A = \text{Tr}_{\bar{A}}(|\Omega\rangle\langle\Omega|)$ associated to the spatial subregion Σ_A . Similarly, if we considered two arbitrary spatial subregions Σ_A and Σ_B , we could try to define a density matrix $\hat{\rho}_{AB}$ for the joint system by tracing out the degrees of freedom associated to $\Sigma \setminus (\Sigma_A \cup \Sigma_B)$, and then compute the negativity of that density matrix as a way to quantify the portion of entanglement shared between the regions Σ_A and Σ_B which could be distilled by agents with access to the two regions.

It turns out that this intuitive picture does not quite work: rigorously speaking, the Hilbert space of a QFT does *not* factorize as a tensor product of local Hilbert spaces associated to complementary spatial regions, and it is not technically right to assign a density matrix to the restriction of the vacuum state $|\Omega\rangle$ (or any state of the QFT, for that matter) to any local subregion. The mathematical reason behind this is that the local algebra of observables on a subregion of a relativistic QFT is given by what is known as a *type III von Neumann algebra*, which does not admit any irreducible representation as an algebra of operators on a Hilbert space, and does not contain any nontrivial faithful

operation with the properties of a “trace” [164]. This means that, when dealing with a relativistic QFT, operations such as taking a partial trace over a subregion are not available, von Neumann entropies of the reduced state of a QFT on a given region are not well-defined, and one cannot even talk about the reduced state of the QFT on a local subregion in terms of a density matrix; in short, most of the technical tools that we got used to when characterizing entanglement in nonrelativistic quantum mechanics are not directly applicable to QFT.

There are very clear indications, however, that any regular state of a QFT must be highly entangled. The first hint can be seen by replacing the continuum QFT with a regularized version of it—say, a lattice theory with some minimum lattice spacing ε —, and observing what happens to entanglement measures of the system as the short-distance cutoff ε is sent to zero. In the regulated theory, the reduced state of the system on a subregion does admit a description as a density matrix, and one can compute, for instance, the entanglement entropy between two complementary regions in the ground state. As one removes the cutoff by taking the limit $\varepsilon \rightarrow 0$, it is seen that the entanglement entropy between a subregion Σ_A and its complement grows unboundedly; in a d -dimensional spacetime, the leading-order (divergent) term in the entanglement entropy is state-independent, and goes like $\text{Area}(\partial\Sigma_A)/\varepsilon^{d-2}$ [133]. This is the well-known UV divergence of entanglement entropy in QFT, with the leading-order divergence given by an area law.⁵ The fact that this divergence is state-independent, incidentally, is the physical reason why the QFT Hilbert space cannot be factorized as $\mathcal{H}_A \otimes \mathcal{H}_{\bar{A}}$: after all, if it could, then there would be states in $\mathcal{H}$ with no correlations between a subregion and its complement—but this cannot happen in any finite-energy state of the QFT, since breaking the correlations at short distances will lead to a divergent (and non-renormalizable) expectation value for the energy-momentum tensor.

Another—much more sophisticated—manifestation of the high amount of entanglement in the vacuum of a QFT is illustrated in the following theorem:

Theorem 4.2.1 (Reeh-Schlieder theorem [149, 197]). *Let $|\Omega\rangle$ be the vacuum of a relativistic QFT in Minkowski spacetime, $\mathcal{H}_{QFT}$ be the Hilbert space obtained by acting on $|\Omega\rangle$ with local field operators,⁶ and $\mathcal{A}(R)$ be the algebra of field observables restricted to any local subregion $R \subset \mathcal{M}$. Then, every state in $\mathcal{H}_{QFT}$ can be approximated arbitrarily well by states of the form $\hat{A}|\Omega\rangle$ for $\hat{A} \in \mathcal{A}(R)$. In other words, the set $\mathcal{A}(R)|\Omega\rangle$ is dense in $\mathcal{H}_{QFT}$.*

⁵For $d = 2$, the polynomial divergence becomes logarithmic and is given by $\log(\xi/\varepsilon)$, where ξ is the minimum between the length of the spatial segment under consideration and the correlation length of the system [39, 95].

⁶In a general theory with superselection sectors, this same statement is true, but now replacing $\mathcal{H}_{QFT}$ with the vacuum sector $\mathcal{H}_0$; see footnote 4.

When a state $|\psi\rangle$ in a Hilbert space $\mathcal{H}$ is such that the set $\mathcal{A}|\psi\rangle$ is dense in $\mathcal{H}$ for some algebra $\mathcal{A}$ acting on $\mathcal{H}$, we say that $|\psi\rangle$ is a *cyclic* vector for $\mathcal{A}$. The Reeh-Schlieder theorem can thus be stated as the observation that the vacuum of a relativistic QFT in Minkowski space is cyclic for every local algebra $\mathcal{A}(R)$. This is a truly remarkable result: it tells us that, although the local algebra $\mathcal{A}(R)$ only contains a subset of the degrees of freedom of the QFT, the action of elements of $\mathcal{A}(R)$ is enough to approximate any state of the QFT to arbitrary accuracy. The only way for this to be true is if the state Ω contains a very high amount of entanglement between degrees of freedom in a region R and its causal complement R' .

To see why that is the case, it is useful to first think about what cyclicity of a state with respect to an algebra means in the case of factorizable Hilbert spaces. In the more familiar picture of a bipartite Hilbert space of the form $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$, a vector $|\psi\rangle \in \mathcal{H}$ will be cyclic for the algebra $\mathcal{B}(\mathcal{H}_1) \otimes \mathbb{1}_{\mathcal{H}_2}$ (which only acts nontrivially on one of the subsystems) if and only if $|\psi\rangle$ is *fully entangled*—i.e., if the reduced states $\hat{\rho}_1 = \text{Tr}_2(|\psi\rangle\langle\psi|)$ and $\hat{\rho}_2 = \text{Tr}_1(|\psi\rangle\langle\psi|)$ are max-rank density operators on $\mathcal{H}_1$ and $\mathcal{H}_2$ respectively. Therefore, the entanglement in any state that is cyclic with respect to a subalgebra in a Hilbert space of the form $\mathcal{H}_1 \otimes \mathcal{H}_2$ is “maximal”, in the sense that the reduced density matrices on either side of the bipartition span their entire respective Hilbert spaces.

Of course, in the case of subregions in QFT, it does not make sense to talk about “max-rank” density matrices, because (as we have emphasized) local subregions do not admit descriptions in terms of such objects. There is, however, a generalization of the notion of a max-rank density matrix which *does* apply to type III von Neumann algebras, which is the concept of a *separating state*. We say that a state $|\psi\rangle \in \mathcal{H}$ is *separating* for an algebra $\mathcal{A}$ if there is no nontrivial operator in $\mathcal{A}$ which annihilates $|\psi\rangle$; in other words, $|\psi\rangle$ is separating for $\mathcal{A}$ if the only way for us to have $\hat{A}|\psi\rangle = 0$ with $\hat{A} \in \mathcal{A}$ is if $\hat{A} = 0$. It is easy to see that a state that is cyclic with respect to an algebra $\mathcal{A}$ is separating for its commutant algebra $\mathcal{A}'$, where we define the commutant $\mathcal{A}'$ as the set of all operators that commute with $\mathcal{A}$. To see that, just note that for any $\hat{A} \in \mathcal{A}$ and for every $\hat{A}' \in \mathcal{A}'$, we have

$$\hat{A}'(\hat{A}|\psi\rangle) = \hat{A}(\hat{A}'|\psi\rangle). \quad (4.13)$$

If $|\psi\rangle$ is annihilated by $\hat{A}'$, then the right-hand side above vanishes for every $\hat{A} \in \mathcal{A}$; therefore, $\hat{A}'$ annihilates the entire set $\mathcal{A}|\psi\rangle$. But if $|\psi\rangle$ is cyclic for $\mathcal{A}$, the set $\mathcal{A}|\psi\rangle$ is dense on the Hilbert space $\mathcal{H}$, so the only way for $\hat{A}'$ to annihilate every state on $\mathcal{A}|\psi\rangle$ is if $\hat{A}'$ is identically zero—in short, $|\psi\rangle$ must be separating for $\mathcal{A}'$.

For simplicity, consider now that the region R is the domain of dependence of some

spatial subregion of a Cauchy surface; i.e., $R = D(\Sigma_A)$ for some $\Sigma_A \subset \Sigma$. This algebra is always the commutant of another algebra—namely, the one associated to the domain of dependence of $\Sigma_{\bar{A}} \equiv \Sigma \setminus \Sigma_A$. Putting these two facts together, plus the fact that the vacuum is also separating for $\mathcal{A}(D(\Sigma_{\bar{A}}))$, we thus conclude that the vacuum state $|\Omega\rangle$ is both cyclic and separating for every local algebra $\mathcal{A}(D(\Sigma_A))$.

It is therefore clear that, although entanglement in QFT is hard to quantify in a mathematically rigorous way, QFTs store a formally infinite amount of entanglement; in fact, a more physical understanding of why the standard tools from entanglement theory in quantum mechanics do not directly apply to QFTs is that, in a way, physical states of a QFT are just “too entangled” for the tools briefly reviewed in Section 4.1 to be useful. In this context, it becomes natural to ask whether this fact can be put to use by agents with limited resources that can interact with the quantum field.

Our goal in the next two chapters will thus be to explore two attempts to tame this divergent amount of entanglement and turn it into something less abstract. The idea in both cases is to try to make sense of this formally infinite amount of entanglement in QFT (which ultimately boils down to the fact that the QFT contains infinitely many degrees of freedom at arbitrarily short distances) into something more operationally meaningful, by translating the relevant physics to subsystems with finitely many degrees of freedom.

In our first attempt, we will explore a very well-known protocol in RQI known as *entanglement harvesting*. This is a process by which two localized probes that couple to a quantum field can get entangled with each other, even if the regions where they are coupled to the field are spacelike separated. In Chapter 5, we will show that, even at leading order in perturbation theory in the coupling between probes and target field, the simple model of probe that was explained in Part I of this thesis is indeed capable of harvesting entanglement at spacelike separation. These results address earlier claims in the literature [154] according to which, due to the fact that every local mode of a relativistic field has to be in a mixed state as a consequence of the Reeh-Schlieder theorem, they would require a finite (i.e., not arbitrarily small) coupling constant in order to harvest entanglement—the reason being that the initial mixedness of the probes is detrimental to entanglement harvesting at weak coupling. We show that, although the Reeh-Schlieder theorem implies that, on a technical level, no local mode of a field is ever perfectly pure, there is no nontrivial lower bound on the mixedness of modes constructed in a local subregion; in particular, one can always make it as pure as one wants, by choosing the correct spatial profile of the mode. We illustrate this concretely by constructing an explicit sequence of modes of the field that get arbitrarily close to being pure. We also show that the mixedness of local modes can effectively be made to vanish by making use of an external potential in the form of what was discussed in Part I, and reveal that when properly interpreted as statements about modes of a quantum

field, the well-known predictions of entanglement harvesting using regular UDW models actually provide a lower bound to the entanglement that can be extracted to field-theoretic probes.

Subsequently, in Chapter 6, we look more deeply into the entanglement structure of quantum fields in order to see how to optimize the extraction of entanglement by local probes. Our main result is an upper bound on the entanglement that can be stored between any finite number of local modes in complementary subregions, which we then use to argue for the optimal coupling between detectors and field which would harvest as much entanglement as possible. The main technical result is rather general, and applies to any globally pure bipartite Gaussian state. When applied to a harmonic system with local interactions (which provides a regularized version of a QFT), we see—as expected—that the spatial profile of the optimal coupling strongly concentrates near the boundary separating the region of interest from its complement. In general, this represents a prescription that establishes the limits of entanglement harvesting, and points to a promising direction in the pursuit of a more general detector-based description of entanglement between arbitrary regions of a QFT.

Chapter 5

Local particle detectors and entanglement harvesting

It has recently been pointed out [154] that local probes consisting of localized modes of quantum fields, which are obtained by restricting a vacuum state of a quantum field to a localized region, will be unavoidably mixed. This places a challenge for the task of entanglement harvesting with field-theoretic probes, since any initial level of mixedness in the probes decreases their ability to harvest entanglement. We will show, however, that this fundamental mixedness (which is a consequence of the Reeh-Schlieder theorem) does not actually pose a restriction to one's ability to perform entanglement harvesting; although every mode of a quantum field restricted to a local subregion will be somewhat mixed, one can always make this mixedness as low as one wishes, especially when considering the model presented in Part I of a field under the action of an external potential that effectively confines it. Therefore, in practice, the limitations to entanglement harvesting with modes of quantum fields are not more severe than the limitations to preparing any other localized quantum system in its ground state.

5.1 Review of entanglement harvesting

We have already said a few words about entanglement harvesting in Section 3.3.2 of Part I, as an illustration of a setup where a simple extension to our formalism of localized quantum fields as local probes can be useful. In this chapter, we will revisit this idea, and see how entanglement can indeed be harvested by modes of a localized quantum field. Before doing

that, however, we should review in more detail how entanglement harvesting is commonly introduced in RQI.

The setup follows the general idea explained in Section 2.1, now with two probe systems. Each probe is again described by a localized quantum system coupled to the quantum field, which we take to be described by a real scalar field. The interaction between probes and systems is thus defined by an action of the form in (2.11),

$$S_I = - \sum_{i \in \{A,B\}} \lambda_i \int dV \Lambda_i(\mathbf{x}) \hat{\mu}_i(\tau_i) \hat{\psi}(\mathbf{x}), \quad (5.1)$$

where we recall that the $\Lambda_i(\mathbf{x})$ are the two spacetime smearing functions that localize the support of the interaction region both in space and in time, and the sub-index $i = A, B$ labels each of the detectors.

We then start the field-detectors system in a fully uncorrelated state,

$$\hat{\rho}_0 = \hat{\rho}_F \otimes \hat{\rho}_{0,AB} \quad (5.2)$$

with $\hat{\rho}_{0,AB}$ being the initial state of the probes. After the interaction, the state of the probes will be given by,

$$\hat{\rho}_{AB} = \text{Tr}_F \left(\hat{U} \hat{\rho}_0 \hat{U}^\dagger \right). \quad (5.3)$$

Note that, although subregions in QFT do not admit traces or density matrices (because they are associated to type III von Neumann algebras, as we mentioned in Section 4.2), both $\hat{\rho}_F$ and Tr_F are defined in the *global* Hilbert space of the field theory, which is type I. Therefore, it is perfectly legal and rigorous to define Tr_F and $\hat{\rho}_F$ for global QFT, just as we do in ordinary quantum mechanics.

To compute quantities associated to this final state (5.3), we resort to perturbation theory. First, we expand U via the Dyson expansion:

$$\begin{aligned} \hat{U} &= \mathcal{T} \exp \left(-i \int dt \hat{H}_I(t) \right) \\ &= \sum_{n=0}^{\infty} (-i)^n \int_{-\infty}^{+\infty} dt_1 \cdots \int_{-\infty}^{t_{n-1}} dt_n \hat{H}_I(t_1) \cdots \hat{H}_I(t_n), \end{aligned} \quad (5.4)$$

where we have implicitly picked a time coordinate t with respect to which spacetime is

being foliated, and the interaction Hamiltonian is thus given by

$$\hat{H}_I(t) = \sum_{i \in \{A, B\}} \lambda_i \int d^d \mathbf{x} \sqrt{-g} \Lambda_i(t, \mathbf{x}) \hat{\mu}_i(\tau_i) \hat{\psi}(t, \mathbf{x}). \quad (5.5)$$

The time evolution operator can thus be expanded as

$$\hat{U} = \mathbb{1} + \hat{U}^{(1)} + \hat{U}^{(2)} + \dots, \quad (5.6)$$

where the term $\hat{U}^{(k)}$ collects all contributions of order λ^k in the coupling constants. With this, the final state can be written as

$$\begin{aligned} \hat{\rho} = \hat{\rho}_0 &+ \overbrace{\hat{U}^{(1)} \hat{\rho}_0 + \hat{\rho}_0 \hat{U}^{(1)\dagger}}^{\propto \lambda} \\ &+ \underbrace{\hat{U}^{(1)} \hat{\rho}_0 \hat{U}^{(1)\dagger} + \hat{U}^{(2)} \hat{\rho}_0 + \hat{\rho}_0 \hat{U}^{(2)\dagger}}_{\propto \lambda^2} + \dots \end{aligned} \quad (5.7)$$

After taking the trace over the field, assuming the initial state of the field to be a Gaussian (quasi-free) state with vanishing one-point function (which implies that all n -point functions of the field with n odd vanish), the leading order contributions to the final state of the probes will be

$$\begin{aligned} \hat{\rho}_{AB} &= \hat{\rho}_{0,AB} \\ &+ \text{Tr}_F \left(\hat{U}^{(1)} \hat{\rho}_0 \hat{U}^{(1)\dagger} + \hat{U}^{(2)} \hat{\rho}_0 + \hat{\rho}_0 \hat{U}^{(2)\dagger} \right) + \mathcal{O}(\lambda^4). \end{aligned} \quad (5.8)$$

Now, denoting the two-point function of the field by $\text{tr}_F \left(\hat{\rho}_F \hat{\psi}(\mathbf{x}) \hat{\psi}(\mathbf{x}') \right) \equiv W(\mathbf{x}, \mathbf{x}')$, and implementing the time-ordering explicitly by inserting Heaviside functions in $t - t'$, we have

$$\begin{aligned} \hat{\rho}_{AB} &= \hat{\rho}_{0,AB} + \sum_{i,j} \lambda_i \lambda_j \int dV dV' \Lambda_i(\mathbf{x}) \Lambda_j(\mathbf{x}') \\ &\times \left[\hat{\mu}_i(t) \hat{\rho}_{0,AB} \hat{\mu}_j(t') W(\mathbf{x}', \mathbf{x}) \right. \\ &\quad - \theta(t - t') \left(\hat{\mu}_i(t) \hat{\mu}_j(t') \hat{\rho}_{0,AB} W(\mathbf{x}, \mathbf{x}') \right. \\ &\quad \left. \left. + \hat{\rho}_{0,AB} \hat{\mu}_i(t') \hat{\mu}_j(t) W(\mathbf{x}', \mathbf{x}) \right) \right] + \mathcal{O}(\lambda^4). \end{aligned} \quad (5.9)$$

To make further progress now we will have to provide some further details about the probes

and their initial state. Most studies of entanglement harvesting in RQI employ detectors given by two-level systems (i.e, qubits), in which case the probe observable $\hat{\mu}_i(\tau)$ is the Pauli $\hat{\sigma}_x^{(i)}(\tau)$ on the Hilbert space of each qubit. In the interaction picture, these evolve as $\hat{\sigma}_x^{(i)}(\tau) = \hat{\sigma}_i^+ e^{i\Omega_i\tau} + \hat{\sigma}_i^- e^{-i\Omega_i\tau}$, where $\hat{\sigma}_i^\pm$ are the raising and lowering operators on the two-dimensional Hilbert space of each detector. In the case of a harmonic oscillator detector, as discussed in Section 2.1, we can take the probe observable to simply be the position operator $\hat{q}(\tau)$, in which case $\hat{\mu}_i(\tau)$ will be proportional to $\hat{a}_i e^{-i\Omega_i\tau} + \hat{a}_i^\dagger e^{i\Omega_i\tau}$, with $\hat{a}_i$ and $\hat{a}_i^\dagger$ being the annihilation and creation operators of the harmonic oscillator. In both cases, Ω_i is the characteristic frequency of each detector, which is given by the energy gap between the ground and excited states of the qubit, or the gap between any two consecutive energy eigenstates of the harmonic oscillator (both measured in the rest frame of each detector). As for the initial state of the probes, it is customary to assume that both start in their ground states; in other words,

$$\hat{\rho}_{0,AB} = |00\rangle\langle 00|, \quad (5.10)$$

where we denote $|nm\rangle \equiv |n\rangle_A \otimes |m\rangle_B$, and $|n\rangle_i$ labels the basis eigenstates of each detector. Naturally, in the qubit case, the labels n, m only take the values 0 or 1, whereas for the harmonic oscillator n and m can take any nonnegative integer value.

After all of this is said and done, Equation (5.9) yields the following density matrix for the two-qubit case, represented in the basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$:

$$\hat{\rho}_{AB}^{\text{qubit}} = \begin{pmatrix} 1 - \mathcal{L}_{AA} - \mathcal{L}_{BB} & 0 & 0 & \mathcal{M}^* \\ 0 & \mathcal{L}_{BB} & \mathcal{L}_{BA} & 0 \\ 0 & \mathcal{L}_{AB} & \mathcal{L}_{AA} & 0 \\ \mathcal{M} & 0 & 0 & 0 \end{pmatrix} + \mathcal{O}(\lambda^4). \quad (5.11)$$

In the case of harmonic oscillator detectors, there will be extra matrix elements coming from two excitations of each probe. In the basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle, |02\rangle, |20\rangle\}$, the final state becomes

$$\hat{\rho}_{AB}^{\text{HO}} = \begin{pmatrix} 1 - \mathcal{L}_{AA} - \mathcal{L}_{BB} & 0 & 0 & \mathcal{M}^* & \mathcal{K}_B^* & \mathcal{K}_A^* \\ 0 & \mathcal{L}_{BB} & \mathcal{L}_{BA} & 0 & 0 & 0 \\ 0 & \mathcal{L}_{AB} & \mathcal{L}_{AA} & 0 & 0 & 0 \\ \mathcal{M} & 0 & 0 & 0 & 0 & 0 \\ \mathcal{K}_B & 0 & 0 & 0 & 0 & 0 \\ \mathcal{K}_A & 0 & 0 & 0 & 0 & 0 \end{pmatrix} + \mathcal{O}(\lambda^4). \quad (5.12)$$

The various terms found appearing above are given by¹

$$\mathcal{L}_{ij} = \lambda_i \lambda_j \int dV dV' e^{-i\Omega_i t} e^{i\Omega_j t'} \Lambda_i(\mathbf{x}) \Lambda_j(\mathbf{x}') W(\mathbf{x}, \mathbf{x}'), \quad (5.13)$$

$$\mathcal{M} = -\lambda_A \lambda_B \int dV dV' \Theta(t - t') W(\mathbf{x}, \mathbf{x}') [e^{i(\Omega_A t + \Omega_B t')} \Lambda_A(x) \Lambda_B(\mathbf{x}') + e^{i(\Omega_A t' + \Omega_B t)} \Lambda_A(\mathbf{x}') \Lambda_B(x)], \quad (5.14)$$

$$\mathcal{K}_i = -\lambda_i^2 \int dV dV' \Theta(t - t') e^{i\Omega_i(t+t')} \Lambda_i(\mathbf{x}) \Lambda_i(\mathbf{x}') W(\mathbf{x}, \mathbf{x}'). \quad (5.15)$$

Now that we have explicit formulas for the final density matrix describing the two detectors, we can easily compute its negativity. At leading order in perturbation theory, the partial transpose of both density matrices (5.32) and (5.33) can have at most one negative eigenvalue, given by [144]

$$E = \frac{1}{2} \left[\mathcal{L}_{AA} + \mathcal{L}_{BB} - \sqrt{(\mathcal{L}_{AA} - \mathcal{L}_{BB})^2 + 4|\mathcal{M}|^2} \right] + \mathcal{O}(\lambda^4). \quad (5.16)$$

The negativity for both the two-qubit and the two-harmonic oscillator cases at leading order in perturbation theory reduces to

$$\mathcal{N} = \max(0, -E) + \mathcal{O}(\lambda^4). \quad (5.17)$$

The problem now reduces to evaluating the quantities $\mathcal{L}_{ii}$ and $\mathcal{M}$, for explicit choices of spacetime smearings $\Lambda_i(\mathbf{x})$ associated to each probe. There is a vast body of literature in RQI devoted to studying this problem in a variety of cases, involving different trajectories for the detectors and different features of the background spacetime (which indirectly affect the detectors' ability to harvest entanglement due to their effects on the Wightman function $W(\mathbf{x}, \mathbf{x}')$ of the field). For several different examples, it is found that there is a range of parameters for which the negativity from Equation (5.16) is nonzero [156, 144, 145, 96]. This can be true even in cases where the supports of the spacetime smearings $\Lambda_A(\mathbf{x})$ and $\Lambda_B(\mathbf{x})$ of both detectors are spacelike separated. When this happens, the detectors cannot signal to each other, but they can still get entangled due to the pre-existing entanglement in the initial state of the field. This is the trademark of entanglement harvesting.

¹To follow standard convention in entanglement harvesting, we have used $\mathcal{M}$ for the nonlocal matrix element of the detectors' density matrix, even though the same character was employed for the spacetime manifold in most of this thesis. We hope that this will not be a source of confusion, and that the meaning of $\mathcal{M}$ will be clear from context whenever it appears.

5.2 Can local particle detectors really harvest entanglement?

The detectors used in the previous section for the purpose of entanglement harvesting were assumed to be effectively nonrelativistic. Based on our comments from the introduction of Part I, however, it would also be good to have an explicit field-theoretic model of the probes which is also capable of harvesting entanglement. A first step in this direction was taken by Maximilian Ruez in [154], which used the Fewster-Verch framework to devise a simple, explicit model of local probes based on modes of a fully relativistic quantum field, and applied this model to the problem of entanglement harvesting.

The model consisted of the following. First, we define the probe theory in terms of two independent probe fields $\mathcal{P}_A \otimes \mathcal{P}_B$, where $\mathcal{P}_A$ and $\mathcal{P}_B$ correspond to algebras of observables associated to the two probes. This system of two probes will then couple to the target field in two separate regions K_A and K_B . The probe fields are initialized in a separable state $\sigma_{AB} = \sigma_A \otimes \sigma_B$, with the initial state of the target being given by ω_S ; the full system is therefore initialized in the state $\omega_S \otimes \sigma_A \otimes \sigma_B$, and is then sent through the scattering map Θ , which is a slightly more formal way of modeling the interaction between probe and field using the formalism of [154]. This then gives rise to an updated state σ'_{AB} of the uncoupled probe theory $\mathcal{P}_A \otimes \mathcal{P}_B$, which will encode the effect of the interaction between probe and field. Finally, agents Alice and Bob, with access to the probe fields A and B in *processing* regions² $N_A, N_B \subseteq \mathcal{M}$ respectively (See Fig. 5.1), will test the entanglement on the joint updated state σ'_{AB} of the two probes by measuring the statistics of the correlations of probe field observables located in their respective processing regions.

In a concrete example with a scalar field, the probe degrees of freedom chosen by Ruez in [154] can be characterized by a pair of field operators

$$\hat{Q}_A = \hat{\varphi}(f_A), \tag{5.18}$$

$$\hat{P}_A = \hat{\varphi}(g_A), \tag{5.19}$$

$$\hat{Q}_B = \hat{\varphi}(f_B), \tag{5.20}$$

$$\hat{P}_B = \hat{\varphi}(g_B), \tag{5.21}$$

where the functions f_i and g_i are compactly supported inside the processing region N_i of interest, and we can choose them such that the commutation relations between the

²In the causal future of K_A and K_B , respectively.




Figure 5.1: Schematic view of a possible configuration of the coupling zones K_A , K_B and processing regions N_A and N_B in a spacetime diagram.

quadratures satisfy

$$[\hat{Q}_i, \hat{P}_j] = i\delta_{ij}\mathbb{1}, \quad (5.22)$$

and thus define a pair of modes of each probe field. The scattering map Θ can then be seen as an effecting evolution of the reduced state of the two modes $(\hat{Q}_A, \hat{P}_A)$ and $(\hat{Q}_B, \hat{P}_B)$, which can in principle lead to entanglement between the two modes if the target field the couple to displays entanglement between the subregions N_A and N_B .

Ruep's conclusion in [154], however, was that, for sufficiently weak couplings between probe and target field, no entanglement can ever be obtained in the final state of the two modes used as detectors in the setup above; the coupling constant between probe and target field must be large enough before the two modes start to become entangled.

This seems to contrast with the well-established results of entanglement harvesting cited above, since the parameter space for which the negativity (5.16) is nonzero does not usually depend on the coupling constant. This observation then led to questions about whether something about the nonrelativistic nature of the probes typically used in RQI was somehow responsible for their ability to harvest entanglement, with other probe models based on quantum fields thus being incapable of reproducing the same results.

From part I we already know that this is not true: since it is in general possible to construct modes of a QFT that behave exactly like harmonic oscillator UDW detectors at leading order in perturbation theory, the results for entanglement harvesting obtained with the usual UDW detector can actually be reinterpreted as statements about modes of a QFT. The inability of the Ruep model to harvest entanglement at weak coupling, therefore, has nothing to do with whether the model is fully relativistic or not. Instead, the difference is that, while in the usual UDW setting the probes were both initialized in their ground states, the initial state of the modes used as detectors by Ruep is always *mixed*, and this initial mixedness is detrimental to the probes' ability to become entangled perturbatively.

The reason why any mode of a QFT whose support is contained in a finite region must be mixed can be traced back to the Reeh-Schlieder theorem. As we have reviewed in Chapter 4, the Reeh-Schlieder theorem tells us that the vacuum state is both cyclic and separating for the algebra of any local subregion. For the Ruep model, since the field vacuum (and therefore its restriction to any local mode) is a Gaussian state with vanishing one-point function, the only way for the mode with quadrature operators $(\hat{Q}, \hat{P})$ to be in a pure state is if there is some value of $\alpha > 0$ for which

$$\hat{a} |\Omega\rangle = \frac{1}{\sqrt{2}} \left(\alpha \hat{Q} + i \frac{\hat{P}}{\alpha} \right) |\Omega\rangle = 0. \quad (5.23)$$

This is because, if a mode is in a pure Gaussian state with vanishing one-point function, then there is a squeezing of its quadratures which will take the mode to its ground state, such that the state is annihilated by the annihilation operator $(\alpha \hat{Q} + i \hat{P}/\alpha)/\sqrt{2}$. But $\hat{a}$ is an element of the algebra $\mathcal{A}(K)$, where K is the local region containing the support of the mode. Since $|\Omega\rangle$ is separating for $\mathcal{A}(K)$, it then follows immediately that $\hat{a} |\Omega\rangle$ cannot vanish—and therefore, any local mode of the QFT must be mixed.

The effect of mixedness can be very easily included in our earlier calculations of entanglement harvesting. The simplest way to see this is to initialize each of the states in a thermal state of nonzero temperature, such that

$$\hat{\rho}_{0,AB} = \hat{\rho}_A \otimes \hat{\rho}_B, \quad (5.24)$$

with the density matrix of each detector being given by

$$\hat{\rho}_i = \frac{1}{Z_i} e^{-\beta_i \hat{H}_{0,i}}. \quad (5.25)$$

Here, $\hat{H}_{0,i}$ naturally corresponds to the free Hamiltonian of the i -th detector which generates time translations with respect to its proper time, and $Z_i := \text{tr } e^{-\beta \hat{H}_{0,i}}$ is its partition function. For the case of a two-level UDW detector, this reduces to,

$$\hat{\rho}_i^{\text{qubit}} = \frac{1}{1 + z_i} (|0\rangle \langle 0| + z_i |1\rangle \langle 1|), \quad (5.26)$$

whereas for a harmonic oscillator UDW detector, we have,

$$\hat{\rho}_i^{\text{HO}} = (1 - z_i) \sum_{n=0}^{\infty} z_i^n |n\rangle_i \langle n|_i \quad (5.27)$$

where $\{|n\rangle_i\}$ naturally corresponds to the eigenbasis of the free Hamiltonian of the i -th harmonic oscillator. In both cases, we have denoted the Boltzmann factor as

$$z_i = e^{-\beta_i \Omega_i} \quad (5.28)$$

where β_i is the inverse temperature of each probe.

It is clear that $z_i = 0$ corresponds to the pure case (detectors at $T = 0$ or $\beta \rightarrow \infty$), and $z_i = 1$ yields the maximally mixed state (with $T \rightarrow \infty$ or $\beta = 0$). This is a useful parametrization because it facilitates a perturbative expansion in the low-mixedness regime³—which, as made explicit by the definition (5.28), corresponds to the regime where the energy scale given by the temperature is much smaller than the detector’s characteristic frequency.

The purity of both the two-level system and the harmonic oscillator can be given simple expressions in terms of z : for the qubit case, we have,

$$\text{Tr } \hat{\rho}_{\text{qubit}}^2 = \frac{1 + z^2}{(1 + z)^2} = 1 - 2z + \mathcal{O}(z^2), \quad (5.29)$$

whereas for the harmonic oscillator,

$$\text{Tr } \hat{\rho}_{\text{HO}}^2 = \frac{1 - z}{1 + z} = 1 - 2z + \mathcal{O}(z^2). \quad (5.30)$$

Therefore, in both cases, the dependence on the Boltzmann factor is exactly the same as long as we are only interested in a low-mixedness regime, where we only keep the leading order contributions in the Boltzmann factor.

By looking at a low-mixedness regime and only keeping terms that are at most linear in z_i on the initial state of the probes, both for the case with two qubit detectors and with two harmonic oscillators, the joint initial state of the detectors can be written as

$$\begin{aligned} \hat{\rho}_{0,\text{AB}} = & (1 - z_A - z_B) |00\rangle \langle 00| + z_A |10\rangle \langle 10| + z_B |01\rangle \langle 01| \\ & + \mathcal{O}(z_A z_B). \end{aligned} \quad (5.31)$$

³Notice small mixedness may not mean small temperatures for human scales. For example, for a hydrogen atom ($\hbar\Omega = 13.6$ eV) $z = 0.001$ implies $T \sim 2.3 \times 10^4$ K. For reference, this is about 4 times larger than the temperature at the surface of the Sun.

This yields the following density matrix for the qubit (in the basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$):

$$\hat{\rho}_{\text{AB}}^{\text{qubit}} = \begin{pmatrix} 1 - (\mathcal{L}_{\text{AA}} + z_{\text{A}}) - (\mathcal{L}_{\text{BB}} + z_{\text{B}}) & 0 & 0 & \mathcal{M}^* \\ 0 & \mathcal{L}_{\text{BB}} + z_{\text{B}} & \mathcal{L}_{\text{BA}} & 0 \\ 0 & \mathcal{L}_{\text{AB}} & \mathcal{L}_{\text{AA}} + z_{\text{A}} & 0 \\ \mathcal{M} & 0 & 0 & 0 \end{pmatrix} + \mathcal{O}(z^2) + \mathcal{O}(z\lambda^2) + \mathcal{O}(\lambda^4). \quad (5.32)$$

In the case of harmonic oscillator detectors, we obtain

$$\hat{\rho}_{\text{AB}}^{\text{HO}} = \begin{pmatrix} 1 - (\mathcal{L}_{\text{AA}} + z_{\text{A}}) - (\mathcal{L}_{\text{BB}} + z_{\text{B}}) & 0 & 0 & \mathcal{M}^* & \mathcal{K}_{\text{B}}^* & \mathcal{K}_{\text{A}}^* \\ 0 & \mathcal{L}_{\text{BB}} + z_{\text{B}} & \mathcal{L}_{\text{BA}} & 0 & 0 & 0 \\ 0 & \mathcal{L}_{\text{AB}} & \mathcal{L}_{\text{AA}} + z_{\text{A}} & 0 & 0 & 0 \\ \mathcal{M} & 0 & 0 & 0 & 0 & 0 \\ \mathcal{K}_{\text{B}} & 0 & 0 & 0 & 0 & 0 \\ \mathcal{K}_{\text{A}} & 0 & 0 & 0 & 0 & 0 \end{pmatrix} + \mathcal{O}(z^2) + \mathcal{O}(z\lambda^2) + \mathcal{O}(\lambda^4). \quad (5.33)$$

It is thus clear that at leading order in both the coupling constant and the Boltzmann factor, the net effect of the nonzero temperature simply amounts to a correction to the populations of the first excited states of each detector. One can therefore trivially obtain the leading order corrections to entanglement harvesting due to mixedness by referencing the known results for initially pure states of the probes, and making the replacement $\mathcal{L}_{ii} \rightarrow \mathcal{L}_{ii} + z_i$. This only affects the negativity between the two detectors by

$$\mathcal{N} = \max(-E_1, 0). \quad (5.34)$$

where now

$$E_1 = \frac{1}{2} \left[P_{\text{A}} + P_{\text{B}} - \sqrt{(P_{\text{A}} - P_{\text{B}})^2 + 4|\mathcal{M}|^2} \right] + \mathcal{O}(z^2, z\lambda^2, \lambda^4), \quad (5.35)$$

with $P_i := \mathcal{L}_{ii} + z_i$. If the Boltzmann factor of both detectors is the same (that is, $z_{\text{A}} = z_{\text{B}} \equiv z$), this expression is further simplified to

$$\mathcal{N} = \max(\mathcal{N}_{z=0} - z, 0), \quad (5.36)$$

where $\mathcal{N}_{z=0}$ corresponds to the negativity computed at second order in the coupling for the case of perfectly pure initial probes. This highlights the fact that the presence of any initial mixedness systematically decreases the entanglement harvested. One can clearly identify the existence of a threshold mixedness, $z_c = \mathcal{N}_{z=0}^{(2)}$, above which one no longer has any entanglement in the final state of the probes at lowest order in perturbation theory.

The result for the negativity simplifies even further when we consider the particular case with two inertial detectors in flat spacetime with identical smearing functions and the same energy gap. In that configuration, we will have $\mathcal{L}_{\text{AA}} = \mathcal{L}_{\text{BB}} \equiv \mathcal{L}$, and therefore the negativity reduces to,

$$\mathcal{N}^{(2)} = \max(|\mathcal{M}| - \mathcal{L} - z, 0). \quad (5.37)$$

This is a well-known expression that has often been described as a competition between the non-local term $|\mathcal{M}|$ and the local noise term $\mathcal{L}$ [151]. We thus see, very intuitively, that entanglement harvesting is only possible at this level in perturbation theory if the nonlocal term $|\mathcal{M}|$, which involves the two-point function of the field evaluated along the trajectories of both detectors, overcomes the local “noise” terms. For the case of initially pure probes, the local noise only includes the vacuum excitation probability term (given by $\mathcal{L}_{ii}$), but when the initial state is already thermally populated due to nonzero temperature of the detector, one also has to take the Boltzmann factor z into account.

In summary, we have shown that, in the usual framework of UDW-like particle detector models commonly employed in RQI, the initial mixedness of the probes does indeed hinder their ability to harvest entanglement perturbatively. This vindicates the qualitative claim made in [154], where probe mixedness was identified as the culprit for their detector model not being able to extract entanglement from the vacuum of a quantum field for low enough coupling. Here, we have provided evidence that this fact is indeed a generic feature of initially mixed probes, that can also be made manifest in UDW-like detectors such as the qubit or the harmonic oscillator.

We also see that, within the context of UDW-like detector models, the threshold mixedness that allows for harvesting (or equivalently, the critical coupling strength λ_c that allows for harvesting at a given mixedness) can be readily computed at lowest order with a perturbative calculation. Since $\mathcal{N}_{z=0}^{(2)} \propto \lambda^2$, for any sufficiently high purity $1 - 2z$, the threshold coupling strength (i.e., the minimum value for the coupling constant necessary in order for entanglement to be harvested at lowest order) scales as $\lambda_c \propto \sqrt{z}$. In particular, if the mixedness of the probe comes from the fact that it is at a small but finite temperature then $\sqrt{z} = \exp(-\hbar\Omega/2k_{\text{B}}T)$. This means that the critical coupling strength is exponentially suppressed with $1/T$, showing that the well-known results on perturbative entanglement harvesting are effectively protected from the problem pointed out in [154] due to the mixedness of the probe, for sufficiently low temperatures.

5.3 Estimating realistic levels of probe mixedness

As discussed in the previous section, one of the central ideas supporting the conclusions of [154] is the fact that quantum fields (in particular, Reeh-Schlieder states) generically have entanglement between space-like separated compact regions. A direct consequence of this entanglement is that when we consider the reduced field state associated with any compact region, we will find that it is a mixed state.

However, [154] provides no quantitative analysis of how mixed local modes must be due to the Reeh-Schlieder property. A quantitative analysis of this mixedness is important because it factors directly into the magnitude of coupling strength needed for entanglement harvesting: if local modes can be arbitrarily pure, then we can potentially extract entanglement using arbitrarily small coupling strengths. This section will provide a quantitative discussion of mixedness due to Reeh-Schlieder-like localization effects. Moreover, we will compare this to the level of mixedness that one can expect due to the ambient temperature of the field to assess how relevant the fundamental *Reeh-Schlieder mixedness* is in a realistic scenario.

This quantitative analysis will be carried out in two stages. Firstly, we will search for the minimally mixed local mode compactly supported in some region and given some ambient temperature. As we will see, we can find local modes which are arbitrarily pure in every case. Secondly, we will investigate the balance between mixedness due to localization and mixedness due to temperature, for a number of fixed local mode profiles. As we will see, in practical scenarios, mixedness due to temperature dominates over mixedness due to localization. This is, of course, what one would expect from experimental physics. Temperature is often the limiting factor in producing pure initial states for quantum computing and quantum information processing.

This section will conclude with some comments comparing the levels of mixedness that one can expect from local modes in a free theory, with the typical mixedness of more realistic local probes constructed from QED (namely, atomic probes or other similar probes typically modeled with UDW-like interactions).

5.3.1 Minimally mixed local field modes

In this subsection, we are interested in identifying the minimum mixedness that a field mode can have while being supported over some compact spacetime region, $x \in R \subset \mathcal{M}$. In particular, we are interested in how the mixedness of this mode varies as we change both the size of the region and the ambient temperature, β , of the field. We will argue that for

any region and any temperature, there are local field modes with non-zero but arbitrarily low mixedness, i.e., arbitrarily high purity.

Dimension-independent setup

Recall that, as we discussed before Equation (5.22), a generic field mode is specified by two dimensionless self-adjoint operators $\hat{V}$ and $\hat{W}$ —both in the span of $\{\hat{\varphi}(\mathbf{x}) \mid \mathbf{x} \in \mathcal{M}\}$ —which fail to commute with each other as $[\hat{V}, \hat{W}] = i\hat{1}$. Since we are interested here in finding the field mode whose reduced state is a) the least mixed possible, while at the same time, b) having support over only a compact region, $R \subset \mathcal{M}$, we will find these modes by a minimization search of $\hat{V}$ and $\hat{W}$ satisfying $\hat{V}, \hat{W} \in \text{span}(\{\hat{\varphi}(\mathbf{x}) \mid \mathbf{x} \in R\})$. Thus, for some yet-to-be-specified measure of mixedness, S , we can define,

$$\begin{aligned} S_{\min}(R, \beta) &:= \inf_{\hat{V}, \hat{W}} S(\hat{V}, \hat{W}, \hat{\rho}) \\ \text{s.t.} \quad &[\hat{V}, \hat{W}] = i\hat{1} \\ &\hat{V} \in \text{span}(\{\hat{\varphi}(\mathbf{x}) \mid \mathbf{x} \in R\}) \\ &\hat{W} \in \text{span}(\{\hat{\varphi}(\mathbf{x}) \mid \mathbf{x} \in R\}). \end{aligned} \tag{5.38}$$

We will call a pair of operators, $\hat{V}$ and $\hat{W}$, feasible if they satisfy the above three constraints.

For any unitarily invariant measure of mixedness (e.g., one minus purity, entropy, etc.), the pair of operators $\hat{V}, \hat{W}$ which achieve this minimum⁴ will not be unique. Indeed, even a linear transformation mapping the quadratures $(\hat{V}, \hat{W})$ to another pair of operators $(\hat{V}', \hat{W}')$ that satisfy the same commutation relations will preserve their mixedness. Let $\hat{X}$ and $\hat{Y}$ be a pair of dimensionless self-adjoint operators that are related to the generic $\hat{V}$ and $\hat{W}$ in this way. Using this freedom, we can impose further constraints on $\hat{X}$ and $\hat{Y}$, namely: $\langle \hat{X} \rangle_{\hat{\rho}} = 0$, $\langle \hat{Y} \rangle_{\hat{\rho}} = 0$, $\langle \{\hat{X}, \hat{Y}\} \rangle_{\hat{\rho}} = 0$ and $\langle \hat{X}^2 \rangle_{\hat{\rho}} = \langle \hat{Y}^2 \rangle_{\hat{\rho}}$. For any $\hat{V}$ and $\hat{W}$, such $\hat{X}$ and $\hat{Y}$ exist.

Let us now take our measure of mixedness, $S(\hat{V}, \hat{W}, \hat{\rho})$, to be the symplectic eigenvalue,⁵ denoted by ν . In Chapter 6, we will come back to more detailed definitions and properties of the symplectic eigenvalue; for now, all we need to know is that the symplectic eigenvalue

⁴Technically speaking we are taking the infimum in the above expression. As such, there may not be any pair of operators $\hat{V}, \hat{W}$ which achieve $S_{\min}$. More carefully we should speak only of sequences of $(\hat{V}, \hat{W})$ which in the limit achieve $S_{\min}$.

⁵Roughly speaking, the symplectic eigenvalue corresponds to the “occupation number” of the mode $(\hat{V}, \hat{W})$ in state $\hat{\rho}$. For more details about its general definition and properties, see the discussion after Equation (6.23) in the next chapter.

of any bosonic mode takes values in the interval $[1, \infty)$ (which 1 corresponding to perfect purity, and ∞ corresponding to maximum mixedness), and for $\hat{X}$ and $\hat{Y}$ obeying the constraints given above, the symplectic eigenvalue takes the very simple form

$$\nu(\hat{X}, \hat{Y}, \hat{\rho}) = \frac{1}{2} \langle \hat{X}^2 + \hat{Y}^2 \rangle_{\hat{\rho}}. \quad (5.39)$$

For the generic $\hat{V}$ and $\hat{W}$ that is obtained by a local symplectic transformation on $(\hat{X}, \hat{Y})$, the symplectic eigenvalue attains the same value as that of $\hat{X}$ and $\hat{Y}$, but the right-hand side of the above equation may be larger. That is,

$$\nu(\hat{X}, \hat{Y}, \hat{\rho}) = \nu(\hat{V}, \hat{W}, \hat{\rho}) \leq \frac{1}{2} \langle \hat{V}^2 + \hat{W}^2 \rangle_{\hat{\rho}}. \quad (5.40)$$

It is important to note that for any feasible $\hat{V}$ and $\hat{W}$ this inequality is saturated by some other feasible pair of operators (namely, $\hat{X}$ and $\hat{Y}$). Thus, identifying the minimally mixed mode with compact support over $R \subset \mathcal{M}$ reduces to the following minimization problem:

$$\begin{aligned} \nu_{\min}(R, \beta) &= \inf_{\hat{V}, \hat{W}} \frac{1}{2} \langle \hat{V}^2 + \hat{W}^2 \rangle_{\hat{\rho}} \\ \text{s.t.} \quad & [\hat{V}, \hat{W}] = i\mathbb{1} \\ & \hat{V}, \hat{W} \in \text{span}(\{\hat{\varphi}(\mathbf{x}) \mid \mathbf{x} \in R\}). \end{aligned} \quad (5.41)$$

We are thus in a sense looking for the minimum uncertainty mode with compact support over R . For notational convenience let us define this uncertainty measure as,

$$u(\hat{V}, \hat{W}, \hat{\rho}) := \frac{1}{2} \langle \hat{V}^2 + \hat{W}^2 \rangle_{\hat{\rho}}. \quad (5.42)$$

We will call this u the objective function.

We claim that for any compact spacetime region, R , and any ambient field temperature, β , we have⁶ $\nu_{\min}(R, \beta) = 1$. That is, while (as was proved in [154]) every local mode of a free quantum field in a thermal state is mixed ($\nu > 1$), there are local modes in any compact region R which are arbitrarily pure ($\nu \rightarrow 1$).

We first note that, since growing the size of a region will always allow for the definition of more modes, decreasing the size of R never decreases $\nu_{\min}$ because then strictly less $\hat{V}$ and $\hat{W}$ will be feasible. That is, $\nu_{\min}(R', \beta) \geq \nu_{\min}(R, \beta)$ when $R' \subset R$. Since we know

⁶Note that we are taking an infimum here such that $\nu_{\min} = 1$ does not imply that there exists a mode which itself is completely pure, i.e., $\nu = 1$.

$\nu_{\min} \geq 1$, if we can prove that $\nu_{\min}(R', \beta) = 1$ for some $R' \subset R$ then we will have proved that $\nu_{\min}(R, \beta) = 1$. This will be our proof strategy.

Assuming that the spacetime is flat, within any compact region, R , we can construct a subregion R' as follows: Consider a point $\mathbf{x}_0$ in the interior of R . Take an inertial frame $(t, \mathbf{x})$ with $\mathbf{x}_0$ as its origin. Choose the $t = 0$ Cauchy surface $\Sigma(t=0)$. Let A_ℓ be the set which is the intersection of $\Sigma(t=0)$ with an ℓ -sized ball centered at $\mathbf{x}_0$. We take $R' = R_\ell := D(A_\ell)$ to be the domain of dependence of A_ℓ . Because $\mathbf{x}_0$ is in the interior of R , we will have $R_\ell \subset R$ for small enough ℓ .

We can then identify any $\hat{V}$ and $\hat{W}$ within the algebra of $\hat{\varphi}$ on any thickened Cauchy surface. Equivalently, we can identify $\hat{V}$ and $\hat{W}$ within the algebra of $\hat{\varphi}$ and its conjugate momentum $\hat{\pi}$ on any (non-thickened) Cauchy surface. For the case of interest $\hat{V}$ and $\hat{W}$ are localized on some compact spacetime region, R_ℓ . By construction when we move $\hat{V}$ and $\hat{W}$ onto the $\Sigma(t=0)$ Cauchy surface, we will find that they are compactly supported over A_ℓ . Thus we have,

$$\begin{aligned} \nu_{\min}(R_\ell, \beta) = \inf_{\hat{V}, \hat{W}} \quad & \frac{1}{2} \langle \hat{V}^2 + \hat{W}^2 \rangle_{\hat{\rho}} \\ \text{s.t.} \quad & [\hat{V}, \hat{W}] = i\hat{1} \\ & \hat{V} \in \text{span}(\{\hat{\varphi}(\mathbf{x}), \hat{\pi}(\mathbf{x}) \mid |\mathbf{x}| < \ell\}) \\ & \hat{W} \in \text{span}(\{\hat{\varphi}(\mathbf{x}), \hat{\pi}(\mathbf{x}) \mid |\mathbf{x}| < \ell\}). \end{aligned} \tag{5.43}$$

where $\mathbf{x}$ is a vector in $\mathbb{R}^d$. Note that for notational convenience we have dropped the time dependence of the field operators: $\hat{\varphi}(\mathbf{x}) = \hat{\varphi}(0, \mathbf{x})$ and $\hat{\pi}(\mathbf{x}) = \hat{\pi}(0, \mathbf{x})$. We will now prove that $\nu_{\min}(R_\ell, \beta) = 1$ for every lengthscale ℓ and every ambient field temperature β . As we have discussed above, since we can find some R_ℓ within any region R this is sufficient to prove $\nu_{\min}(R, \beta) = 1$ for every R and β .

Our proof method will be as follows: We will provide a sequence of feasible operators, $\hat{V}$ and $\hat{W}$, which achieve the limit $u = \frac{1}{2} \langle \hat{V}^2 + \hat{W}^2 \rangle_{\hat{\rho}} \rightarrow 1$, where we dropped the arguments of the objective function $u(\hat{V}, \hat{W}, \hat{\rho})$ for notational convenience. In particular, our sequence will have $\hat{V}$ supported only by $\hat{\varphi}(\mathbf{x})$ and $\hat{W}$ supported only by $\hat{\pi}(\mathbf{x})$ as,

$$\hat{V} = \int d^d \mathbf{x} v(\mathbf{x}) \hat{\varphi}(\mathbf{x}), \quad \hat{W} = \int d^d \mathbf{x} w(\mathbf{x}) \hat{\pi}(\mathbf{x}), \tag{5.44}$$

for some real functions $v(\mathbf{x})$ and $w(\mathbf{x})$. In these terms, the compact support constraint is just that both $v(\mathbf{x})$ and $w(\mathbf{x})$ have support only over $\mathbf{x}$ with $|\mathbf{x}| < \ell$. The commutation

constraint for $\hat{V}$ and $\hat{W}$ can be written in terms of $v(\mathbf{x})$ and $w(\mathbf{x})$ as,

$$[\hat{V}, \hat{W}] = i\hat{1} \Leftrightarrow \int d^d \mathbf{x} v(\mathbf{x}) w(\mathbf{x}) = 1. \quad (5.45)$$

Finally, the objective function in terms of $v(\mathbf{x})$ and $w(\mathbf{x})$ is,

$$\begin{aligned} u = & \frac{1}{2} \iint d^d \mathbf{x} d^d \mathbf{y} \langle \hat{\varphi}(\mathbf{x}) \hat{\varphi}(\mathbf{y}) \rangle_{\hat{\rho}} v(\mathbf{x}) v(\mathbf{y}) \\ & + \frac{1}{2} \iint d^d \mathbf{x} d^d \mathbf{y} \langle \hat{\pi}(\mathbf{x}) \hat{\pi}(\mathbf{y}) \rangle_{\hat{\rho}} w(\mathbf{x}) w(\mathbf{y}), \end{aligned} \quad (5.46)$$

where the thermal two-point functions are well-known for free scalar fields [193]. Written in terms of the Fourier transforms of $v(\mathbf{x})$ and $w(\mathbf{x})$, the objective function then takes the form

$$\begin{aligned} u = & \frac{1}{4} \left[\int d^d k \frac{1}{\omega_{\mathbf{k}}} \coth(\beta \omega_{\mathbf{k}}/2) |\tilde{v}(\mathbf{k})|^2 \right. \\ & \left. + \int d^d k \omega_{\mathbf{k}} \coth(\beta \omega_{\mathbf{k}}/2) |\tilde{w}(\mathbf{k})|^2 \right]. \end{aligned} \quad (5.47)$$

For $d \geq 2$ spatial dimensions, we will take our sequence to have $v(\mathbf{x})$ and $w(\mathbf{x})$ spherically symmetric. In this case, the objective function (5.47) simplifies greatly, since the Fourier transforms $\tilde{v}(\mathbf{k})$ and $\tilde{w}(\mathbf{k})$ will only depend on $|\mathbf{k}|$, and the angular part of the integral in (5.47) can be readily computed. In this case, Equation (5.47) then becomes

$$\begin{aligned} u = & \frac{\pi^{d/2}}{2\Gamma(\frac{d}{2})} \left[\int_0^\infty dk \frac{k^{d-1}}{\omega_{\mathbf{k}}} \coth\left(\frac{\beta \omega_{\mathbf{k}}}{2}\right) |\tilde{v}(k)|^2 \right. \\ & \left. + \int_0^\infty dk k^{d-1} \omega_{\mathbf{k}} \coth\left(\frac{\beta \omega_{\mathbf{k}}}{2}\right) |\tilde{w}(k)|^2 \right], \end{aligned} \quad (5.48)$$

where $\omega_{\mathbf{k}} = |\mathbf{k}|$ is the field's dispersion relation and $k \equiv |\mathbf{k}|$ is just shorthand for the modulus of the wave vector. It should be noted that the factor of k^{d-1} above makes both integrals convergent in the IR for $d \geq 3$. For $d = 1$ or $d = 2$, we will need to choose $\tilde{v}(k)$ and $\tilde{w}(k)$ to go to zero sufficiently quickly as $k \rightarrow 0$ to regularize the IR. For all d we need $\tilde{v}(k)$ and $\tilde{w}(k)$ to go to zero sufficiently quickly as $k \rightarrow \infty$ to regularize the UV.

In the following subsection, we consider the flat 1+1 dimensional case and will produce a sequence of $v(\mathbf{x})$ and $w(\mathbf{x})$ with $u \rightarrow 1$ in the limit. At the end of the next subsection, we will also briefly discuss the prospects of generalizing our proof to higher dimensions and

curved spacetimes.

1+1 dimensional proof

For simplicity, we will now restrict our attention to the 1+1 dimensional case. In this case the objective function (5.47) simplifies to,

$$u = \frac{1}{2} \int_{-\infty}^{\infty} dk \frac{1}{\omega_k} \coth(\beta\omega_k/2) |\tilde{v}(k)|^2 + \frac{1}{2} \int_{-\infty}^{\infty} dk \omega_k \coth(\beta\omega_k/2) |\tilde{w}(k)|^2, \quad (5.49)$$

In order to show that $\nu_{\min}(R_\ell, \beta) = 1$ we will now present a sequence of functions $v_{m,\kappa}(x)$ and $w_{m,\kappa}(x)$ which in the limit have $u \rightarrow 1$ while satisfying the above discussed constraints. Namely, each pair of $v_{m,\kappa}(x)$ and $w_{m,\kappa}(x)$ will obey the commutation condition (5.45) with $n = 1$ and will be supported only over $x \in [-\ell, \ell]$.

Consider the two-parameter family of functions,

$$v_{m,\kappa}(x) = \sqrt{\omega_\kappa} z_{m,\kappa}(x), \quad (5.50)$$

$$w_{m,\kappa}(x) = \frac{1}{\sqrt{\omega_\kappa}} z_{m,\kappa}(x), \quad (5.51)$$

where

$$z_{m,\kappa}(x) = A_{m,\kappa} \partial_x^2 [\cos(\kappa x) B_m(x/\ell)] \quad (5.52)$$

is parametrized by a real number $\kappa > 0$ and an integer $m \geq 3$. The prefactor $A_{m,\kappa}$ ensures that $z_{m,\kappa}(x)$ is L^2 normalized. $B_m(s)$ are the B-spline functions [103] which result from repeated self-convolution of the rectangle function,

$$B_m(s) := \Pi^{*m}(m s) \quad (5.53)$$

where $\Pi(s)$ is the rectangle function over the domain $s \in [-1, 1]$, the operator $*$ indicates convolution, and the superscript $*m$ indicates repeated self-convolution. For instance, $B_1(s)$ is the rectangle function over $s \in [-1, 1]$, and $B_2(s)$ is the triangle function over $s \in [-1, 1]$. In general, $B_m(s)$ is supported over $s \in [-1, 1]$ and composed piecewise of m polynomials of degree $m - 1$.

These are convenient functions to investigate because a) they ensure that $\hat{V}$ and $\hat{W}$




Figure 5.2: The power spectrum, $|\tilde{z}_{m,\kappa}(k)|^2$, given by (5.54) is shown for several values of m and κ . The peaks are paired mirror-wise. From inside to outside we have $(m = 3, \kappa = 5\pi)$, $(m = 4, \kappa = 15\pi)$, $(m = 5, \kappa = 25\pi)$, $(m = 6, \kappa = 35\pi)$, and $(m = 7, \kappa = 45\pi)$ with $\ell = 1$ throughout. As κ increases the peaks are displaced farther from the origin. As m is increased the polynomial tails of the distributions are suppressed. Away from the peaks $|\tilde{z}_{m,\kappa}(k)|^2$ decays polynomially as $|k|^{4-2m}$. Increasing m also widens the peaks. The peak width is proportional to m .

are automatically feasible, and b) their Fourier transforms take a convenient form. To see the feasibility of $\hat{V}$ and $\hat{W}$, first note that since $z_{m,\kappa}(x)$ is L^2 normalized, $v_{m,\kappa}(x)$ and $w_{m,\kappa}(x)$ satisfy the commutation constraint, Equation (5.45). Secondly note that $B_m(x/\ell)$ is compactly supported over $x \in [-\ell, \ell]$. Moreover, multiplying by $\cos(\kappa x)$ and taking two derivatives does not change the function's support. As such $z_{m,\kappa}(x)$ —and therefore also $v_{m,\kappa}(x)$ and $w_{m,\kappa}(x)$ —are compactly supported over $x \in [-\ell, \ell]$.

The Fourier transform of $z_{m,\kappa}(x)$ is given by

$$\begin{aligned} \tilde{z}_{m,\kappa}(k) &= k^2 C_{m,\kappa} [(\delta(k - \kappa) + \delta(k + \kappa)) * S_m(k\ell)], \\ S_m(r) &:= \left[\frac{\sin(r/m)}{r/m} \right]^m. \end{aligned} \quad (5.54)$$

where the L^2 normalizing prefactor, $A_{m,\kappa} \rightarrow C_{m,\kappa}$, has absorbed some stray constants. $S_m(k\ell)$ is proportional to the Fourier transform of $B_m(x/\ell)$. The convolution, $*$, with the two delta functions accounts for the multiplication by $\cos(\kappa x)$. Finally, the k^2 appearing above expression is a consequence of the two derivatives taken in Equation (5.52). To anticipate what is to come, this factor of k^2 helps us regulate the IR divergences present in (5.47) for $n = 1$. Fig. 5.2 shows $|\tilde{z}_{m,\kappa}(k)|^2$ for values of increasing m and κ .

As Fig. 5.2 shows, $|\tilde{z}_{m,\kappa}(k)|^2$ has two peaks located roughly at $k = \pm\kappa$. By increasing κ we can push these peaks as far apart as we like. Away from these peaks $|\tilde{z}_{m,\kappa}(k)|^2$ decays polynomially as $|k|^{4-2m}$. By increasing m we can suppress these tails as much as we like. However, increasing m does have the effect of widening both peaks: note $S_m(k\ell)$ has its first zeroes at $k = \pm m\pi/\ell$. As we will now discuss, by increasing κ and m we can make the objective function, $u = \frac{1}{2}\langle \hat{V}^2 + \hat{W}^2 \rangle_{\hat{\rho}}$, as near to one as we wish.

Written in terms of the Fourier transform of $z_{m,\kappa}(x)$, the objective function is

$$\begin{aligned} u &= \frac{1}{2} \int_{-\infty}^{\infty} dk \left(\frac{\omega_{\kappa}}{\omega_k} + \frac{\omega_k}{\omega_{\kappa}} \right) \coth(\beta\omega_k/2) |\tilde{z}_{m,\kappa}(k)|^2 \\ &=: \left\langle \frac{1}{2} \left(\frac{\omega_{\kappa}}{\omega_k} + \frac{\omega_k}{\omega_{\kappa}} \right) \coth(\beta\omega_k/2) \right\rangle_{|\tilde{z}_{m,\kappa}(k)|^2}. \end{aligned} \quad (5.55)$$

First, it should be noted that this integral is finite. The IR divergences coming from $1/\omega_k \rightarrow \infty$ and $\coth(\beta\omega_k/2) \rightarrow \infty$ as $k \rightarrow 0$ are controlled by k^2 appearing in $\tilde{z}_{m,\kappa}(k)$. Moreover, the UV divergence coming from $\omega_k \rightarrow \infty$ as $k \rightarrow \infty$ is controlled by the polynomial tails of $\tilde{z}_{m,\kappa}(k)$ so long as $m \geq 3$.

Let us discuss the asymptotic behavior of the sequence of functions $v_{m,\kappa}(x)$ and $w_{m,\kappa}(x)$: As we will now show, by increasing both κ and m we can bring the objective function as near to 1 as we like. First, we note that by taking κ large enough we can move an arbitrary amount of the support of $|\tilde{z}_{m,\kappa}(k)|^2$ into a regime where $\coth(\beta\omega_k/2) \approx 1$. Moreover, this approximation can be made arbitrarily accurate by increasing κ . Making this approximation the objective function becomes,

$$u \approx \frac{1}{2} \left\langle \frac{\omega_{\kappa}}{\omega_k} + \frac{\omega_k}{\omega_{\kappa}} \right\rangle_{|\tilde{z}_{m,\kappa}(k)|^2}. \quad (5.56)$$

Note that for large enough κ we can also approximate both ω_k/ω_{κ} and ω_{κ}/ω_k as linear functions of k in the neighborhoods around $k = \pm\kappa$. For any fixed width around our two peaks, we can also make these linear approximations arbitrarily accurate within those regions by increasing κ . Thus, even if m is large such that the peaks around $k = \pm\kappa$ are wide, we can always move to larger κ to justify this approximation over the peak's width.

Finally, note that the average of any linear function, $L(x)$, with respect to any probability distribution $p(x)$, is simply given by $\langle L(x) \rangle_p = L(\langle x \rangle_p)$. Thus recalling that we can

approximate both ω_k and $1/\omega_k$ linearly, we have that for large enough κ ,

$$\langle \omega_k \rangle_{|\tilde{z}_{m,\kappa}(k)|^2} \approx \omega_\kappa \quad (5.57)$$

$$\left\langle \frac{1}{\omega_k} \right\rangle_{|\tilde{z}_{m,\kappa}(k)|^2} \approx \frac{1}{\omega_\kappa} \quad (5.58)$$

to arbitrary accuracy. Therefore in the limit $\kappa \rightarrow \infty$ we have $u \approx 1$ to arbitrary accuracy. This establishes that $u \rightarrow 1$ as claimed above.

We have thus shown for a free 1+1 dimensional thermal quantum field in a flat spacetime that within any compact spacetime region $R \subset \mathcal{M}$ we can find a local mode that is arbitrarily pure. To extend this proof to higher dimensions, one simply needs to design $v(x)$ and $w(x)$ with $u \rightarrow 1$ for each n . Concerning curved spacetimes, since this proof works primarily in the UV of the theory (i.e., large k and small ℓ) it should also be true in curved spacetimes at least at scales where the spacetime looks flat enough and as long as the spacetime has enough symmetry for the definitions of a vacuum and thermal field states to be unambiguous.

5.3.2 Realistic levels of mixedness

In the previous subsection, we saw that local modes can be arbitrarily pure regardless of the region size ℓ and the ambient field temperature β . However, constructing these arbitrarily pure modes required us to go into the deep UV of our theory. What level of mixedness can we expect for reasonable choices of parameters describing the given modes? In this subsection, we will provide a quantitative analysis of the mixedness of a few local mode profiles in 1 + 1 and 3 + 1 dimensions. As we will see the dependence of the local mode's purity on its shape, size, and the ambient temperature all indicate that local field modes of a free theory are not good models for experimentally realistic localized probes.

In particular, we will consider local modes specified by $\hat{V}$ and $\hat{W}$ given by (5.44) for some $v(\mathbf{x})$ and $w(\mathbf{x})$. We take $v(\mathbf{x})$ and $w(\mathbf{x})$ to be spherically symmetric functions which are supported only over $\mathbf{x}$ with $|\mathbf{x}| < \ell$ and satisfy the commutation constraint, Equation (5.45). For a thermal state we know that $\langle \{\hat{V}, \hat{W}\} \rangle_{\hat{\rho}} = 0$ as well as $\langle \hat{V} \rangle_{\hat{\rho}} = 0$ and $\langle \hat{W} \rangle_{\hat{\rho}} = 0$. For such $\hat{V}$ and $\hat{W}$, the symplectic eigenvalue associated with this mode is given by,

$$\nu = 2\sqrt{\langle \hat{V}^2 \rangle_{\hat{\rho}} \langle \hat{W}^2 \rangle_{\hat{\rho}}} \quad (5.59)$$

where for spatial dimension $d = 1$ we have, formally,

$$\begin{aligned}\langle \hat{V}^2 \rangle_{\hat{\rho}} &= \frac{1}{2} \int_{-\infty}^{\infty} dk \frac{1}{\omega_k} \coth(\beta\omega_k/2) |\tilde{v}(k)|^2 \\ \langle \hat{W}^2 \rangle_{\hat{\rho}} &= \frac{1}{2} \int_{-\infty}^{\infty} dk \omega_k \coth(\beta\omega_k/2) |\tilde{w}(k)|^2,\end{aligned}\tag{5.60}$$

and for spatial dimension $d \geq 2$ we have,

$$\begin{aligned}\langle \hat{V}^2 \rangle_{\hat{\rho}} &= \frac{\pi^{d/2}}{2\Gamma\left(\frac{d}{2}\right)} \int_0^{\infty} dk \frac{k^{d-1}}{\omega_{\mathbf{k}}} \coth\left(\frac{\beta\omega_{\mathbf{k}}}{2}\right) |\tilde{v}(k)|^2 \\ \langle \hat{W}^2 \rangle_{\hat{\rho}} &= \frac{\pi^{d/2}}{2\Gamma\left(\frac{d}{2}\right)} \int_0^{\infty} dk k^{d-1} \omega_{\mathbf{k}} \coth\left(\frac{\beta\omega_{\mathbf{k}}}{2}\right) |\tilde{w}(k)|^2,\end{aligned}\tag{5.61}$$

where $\omega_{\mathbf{k}}$ is the field's dispersion relation and $k \equiv |\mathbf{k}|$ is shorthand for the modulus of the wavevector $\mathbf{k}$.

As we noted following Equation (5.48), for $d = 1$ and $d = 2$ there may be an IR divergence in $\langle \hat{V}^2 \rangle_{\hat{\rho}}$ if $|\tilde{v}(k)|^2$ does not go to zero quickly enough as $k \rightarrow 0$. When needed, we will consider two kinds of IR regularization: 1) we can take the field to be massive with a mass M . In this case the dispersion relation becomes $\omega_{\mathbf{k}} = \sqrt{|\mathbf{k}|^2 + M^2}$. Alternatively, 2) we can take the field to be in a Dirichlet cavity of size L . For $n = 1$ spatial dimensions this discretizes the allowed wavenumbers as $k_j = j\pi/L$ for $j = 1, 2, 3, \dots$. In either case, our primary interest will be in the $M \rightarrow 0$ and $L \rightarrow \infty$ limits of the regulators, when they exist.

Mixedness of select 1+1 and 3+1 dimensional profiles

As a first example, let us consider the local modes in a 1+1 dimensional flat spacetime. Specifically, let us consider profiles 1 through 6 shown in Figure 5.3. Profiles 1, 2, and 3 are given by

$$v(x) = w(x) = A_m B_m(x/\ell)\tag{5.62}$$

with $m = 1, 2$, and 3 where $B_m(s)$ are the B-spline functions defined in Equation (5.53). The constants $A_m \in \mathbb{R}$ are chosen so that the functions are L^2 -norm one. As we will see, the mixedness of these modes with these profiles presents an IR divergence as we remove the IR regulator (i.e., taking either $m \rightarrow 0$ or $L \rightarrow \infty$). Profiles 4, 5, and 6 in Figure 5.3




Figure 5.3: The profiles of the local modes considered in this chapter. Specifically, profiles 1, 2, and 3 show Equation (5.62) with $m = 1, 2, 3$. Likewise, profiles 4, 5, and 6 show Equation (5.63) with $m = 4, 5, 6$. Unlike the previous profiles, the final three profiles are 3 dimensional. Specifically, profiles 7, 8, and 9 show Equation (5.64) with $m = 2, 3, 4$.

are given by

$$v(x) = w(x) = C_m \partial_x^2 B_{m-1}(x/\ell) \quad (5.63)$$

with $m = 4, 5$, and 6 where the $C_m \in \mathbb{R}$ are chosen such that the functions have L^2 -norm of one. As we will see, the modes with these ‘derivative’ profiles do not have an IR divergence as we take either $M \rightarrow 0$ or $L \rightarrow \infty$.

Let us first consider profiles 1, 2, and 3. The purity (i.e, $P := \text{Tr}(\hat{\rho}_{\text{local}}^2) = 1/\nu$) of these modes is shown in Fig. 5.4. On the top row the IR regularization is set by the field mass M and on the bottom row it is set by a cavity length scale L . Let us consider the massive case first.

As the top row of Fig. 5.4 shows, for these modes we have high purity only when both the temperature is much less than the mass scale (when $k_B T \ll M c^2$, on the left side of each figure) and when the mode’s support is much larger than the mass scale (when $\ell \gg \hbar/Mc$, on the top of each figure). In the massless limit $M \rightarrow 0$ we move downward and rightward on these figures. Due to the IR divergence in $\langle \hat{V}^2 \rangle_\rho$ we have $P \rightarrow 0$ in the $M \rightarrow 0$ limit for these modes.

Note that for temperatures at or above the mass scale, $kT \gtrsim M c^2$, the contour lines become nearly vertical, indicating that these modes’ purities are nearly independent of the support size in this regime. Conversely, for temperatures below the mass scale, $kT \lesssim M c^2$, the contour lines are nearly horizontal such that the modes’ purities are nearly independent




Figure 5.4: The purity of select local modes as we vary the size of their support (vertical axis) and the ambient temperature of the field (horizontal axis). The top row of subfigures has an IR regulator introduced by means of a field mass M to which the support size ℓ and ambient temperature T are compared. The bottom row of subfigures has an IR regulator introduced by means of cavity walls a distance L apart from each other. The columns show the purity of the local mode given by profiles 1, 2, and 3 in Figure 5.3.

of the ambient temperature in this regime. Thus (at least for these modes, and with this IR regularization) we have an idea of when the mode's mixedness is dominated by temperature effects and when it is dominated by localization effects.

Where can we expect a typical laboratory situation to lie on these figures? As a quasi-realistic⁷ example, consider a local mode of atomic size, $\ell = 53$ pm (Bohr radius), in a field with the mass of an electron, $M = 511$ keV/ c^2 (electron mass) and a temperature of $T = 4.2$ K (liquid Helium temperature). In this scenario we have $\ell Mc/\hbar = 137$ and $k_B T/Mc^2 = 7 \times 10^{-10}$. This places us in the top left corner of these figures, where we have very high purity and near independence from temperature. Indeed, one might expect that in practice we will always be on the far-left sides of these figures where localization dominates; one rarely deals with temperatures larger than the mass scale in practice. However it is worth recalling that as we remove the IR regularization $M \rightarrow 0$, we move both downward and rightward in these figures. That is, we move into the regime where mixedness due to temperature is dominant (on the right side) and into lower purities (downwards). We will hold the interpretation of these results for the moment: before extracting conclusions from this study, we will first see that things actually change substantially when we consider other IR regularizations and other mode profiles.

Let us consider the same three mode profiles, but now massless, $M = 0$, and an IR regularization introduced by placing the field in a cavity of size L . The purity of the first three local modes with this IR regularization is shown on the bottom row of Fig. 5.4. As the figure shows, for these modes we have high purity only when both the temperature is much less than the cavity scale (when $k_B T \ll \hbar c/L$, on the left side of the figure) and when the mode's support is comparable to the cavity size (when $\ell \sim L$ at the top of the figure). In the limit $L \rightarrow 0$ we move downward and rightward on this figure. As before due to the IR divergence in $\langle \hat{V}^2 \rangle_{\hat{\rho}}$ we have in the $L \rightarrow 0$ limit $\mu \rightarrow 0$ for these modes.

We can again identify in which regimes the mixedness is dominated by localization and when it is dominated by temperature. For this IR regularization, the mixedness is localization-dominated (i.e., the purity is nearly independent of temperature) when the ambient temperature is less than the scale set by the cavity size. When the temperature is larger than this scale the mixedness is temperature-dominated (i.e., the purity is nearly independent of the support size).

We can again ask where can we expect a typical situation to lie in these figures. As a

⁷Of course, pretending to accurately model any kind of atomic probe by using a (temporarily) localized free field theory mode is problematic and provides misleading conclusions about purity. A slightly more appropriate version of this is to consider a quadratic field that is confined by some external potential along the lines of what was presented in Part I of this thesis. The external potential, as we will see later, typically decreases the mixedness of modes with a given spatial support.




Figure 5.5: The purity of select local modes as we vary the size of their support (vertical axis) and the ambient temperature of the field (horizontal axis). The top row of subfigures has an IR regulator introduced by means of a field mass M to which the support size ℓ and ambient temperature T are compared. The bottom row of subfigures has an IR regulator introduced by means of cavity walls a distance L apart from each other. The columns show the purity of the local mode given by profiles 4, 5, and 6 in Figure 5.3.




Figure 5.6: The purity of select local modes as we vary the size of their support (vertical axis) and the ambient temperature of the field (horizontal axis). Each subfigure has an IR regulator introduced by means of a field mass M to which the support size ℓ and ambient temperature T are compared. The subfigures show the purity of the local mode given by profiles 7, 8, and 9 in Figure 5.3.

quasi-realistic⁸ example, consider a local mode of atomic size, $\ell = 53$ pm, in a room-sized cavity, $L = 1$ m and a temperature of $T = 4.2$ K. In this scenario we have $\ell/L = 5 \times 10^{-11}$ and $k_B T L / \hbar c = 1834$. This places us in the bottom right corner of these figures, where we have very low purity and near independence from the support size ℓ . Indeed, one might expect that in practice we will always be on the far-right temperature-dominated sides of these figures; as $L \rightarrow \infty$ every temperature will eventually have $k_B T L / \hbar c \gg 1$.

It should be stressed that the qualitative features of the above discussion (regularization via M or L) are limited to the particular modes we are studying here—namely, Equation (5.62) with $m = 1, 2$ and 3. As will now discuss, the situation changes qualitatively when we consider the modes given by profiles 4, 5, and 6 in Figure 5.3 which do not need IR regularization.

The purity of these modes is shown in Fig. 5.5. On the top row, we consider a field of mass M for comparison (although no IR regularization is needed in this case). On the bottom row, this IR regularization is set by the cavity length L . As before, $M \rightarrow 0$ or $L \rightarrow \infty$ moves us downwards and rightwards in each figure. The key qualitative difference between the purity of these three modes and those previously considered is that now the purity is well-defined as we take the massless or infinite length limits. The behavior of these modes purities when $M = 0$ and $L \rightarrow \infty$ is shown in Fig. 5.7a.

⁸Again, we warn strongly against using a local mode of a free field as a realistic model of an atomic probe. We will discuss this in more detail later.

The intuition from both Fig. 5.5 and Fig. 5.7a is that the purity of these modes increases as we lower the ambient field temperature, moving leftward within each figure. However, perhaps counter-intuitively, we have increasing purity as we make the support of the mode smaller in some regimes (moving downward in each figure). To see why we have this behavior note that as we decrease ℓ we have $\tilde{v}(k)$ and $\tilde{w}(k)$ supported over increasingly large k . For any given temperature the modes with high wavenumbers are themselves purer than modes with lower k . Intuitively, while the temperature β is the same for each k the number of excitations $\coth(\beta\omega_k/2)$ decreases for increasing k .

Indeed, the shape of the mode (and not just the size) is important to determine the purity of the partial state of those modes when the total field state is pure. This points out that purity is extremely sensitive to the shape of the chosen mode: even a mode that occupies all the space in an optical cavity can be entangled with other modes—of different shapes—also supported in the whole cavity.

We thus have very different qualitative features depending on whether we are considering the 1+1 dimensional local modes. We now turn to the question of which of these behaviors we should expect in higher dimensions. Consider the 3D analogues of the B-spline functions defined in Equation (5.53), namely

$$B_m^{(3D)}(x, y, z) := \Pi_{(3D)}^{*m}(m x, m y, m z) \quad (5.64)$$

where $\Pi_{(3D)}(x, y, z)$ is the characteristic function for the unit ball, and the exponent $*m$ indicates repeated self-convolution. Equation (5.64) with⁹ $m = 2, 3, 4$ are profiles 7, 8, and 9 in Fig. 5.3. The purities of these modes with a field mass M are shown in Fig. 5.6. As this figure shows, the purity of these 3+1 dimensional local modes is well defined in the massless limit and it is qualitatively similar to 1+1 dimensional modes considered in Fig. 5.5 (e.g., profiles 4, 5, and 6 in Fig. 5.3). The behavior of these modes' purities in the $M = 0$ case is shown in Fig. 5.7b.

Same as we did before, we can set a quasi-realistic atomic model by taking scales $\ell = 53$ pm and $T = 4.2$ K. For these scales, we have $k_B T \ell / \hbar c = 9.7 \times 10^{-8}$ placing us on the far left of Fig. 5.7a and Fig. 5.7b. Note that the purity of these modes does not increase to unity as we decrease the ambient field temperature. Rather, the purity seems to max out at at most $P \approx 0.95$. For these mode profiles, there appears to be a limit to how pure the local modes can be made by decreasing the ambient temperature.

This is a particular illustration of the claim in [154]: These modes are spatially localized and therefore are unavoidably entangled with other modes when the state of the field is

⁹The case $m = 1$ presents UV divergences associated with that particular shape.




Figure 5.7: The purity of three different local modes is shown without an IR regulator. In subfigure a) the field is 1+1 dimensional and in b) the field is 3+1 dimensional. In figure a) the three lines show the purity of the local mode given by profiles 4, 5, and 6 in Figure 5.3. b) The three lines show the purity of the local mode given by profiles 7, 8, and 9 in Figure 5.3.

the vacuum [149, 157]. However, this is certainly not what we would expect from a good model of an atomic probe. Indeed, atomic probes could be prepared in states as pure as > 0.999 experimentally with technology now more than a decade old [102]. Furthermore, the amount of purity does depend strongly on the shape of the mode. Just specifying a localization scale is not enough to fully account for the mixedness behavior of these modes. Indeed, as discussed in the previous subsection we could always find modes with a shape such that they are arbitrarily pure. What this points out is that understanding these localized modes as a model for localized detectors (as for example atomic probes) is not a good idea.

There is also another reason why modeling a realistic localized probe as a mode of a free scalar field of a relativistic QFT may be inappropriate. In general terms, this has to do with the relevant scales that are at play when one models the excitations and de-excitations of the probe. When considering the ratio $k_B T/Mc^2$ (i.e., comparing the temperature with the rest energy of each excitation of a field with mass M), we are implicitly assuming that the source of mixedness arising from a nonzero temperature would be due to the creation of additional particle excitations: in other words, the mixedness would come from sectors of the field's Hilbert space with different numbers of particles, whose energy gap would have to be at least of the order of the rest energy of a single excitation. For realistic scenarios, however, there is a much more relevant energy scale—which is expected to be the one that actually prevails in experimental setups with atomic probes—coming from the energy gap within the internal levels of an effectively non-relativistic detector following *single*-particle dynamics. This would correspond, for instance, to the energy gap between the atomic levels of a hydrogen-like atom—which is of the same order of magnitude as the typical energy of a free *non*-relativistic particle in a box, if we let the size of the box be of the order of the Bohr

radius. This strongly suggests that the relevant effective degrees of freedom of the probe that will account for its dynamics (at least at the scales we are mainly interested in) are, for all practical purposes, non-relativistic—which reinforces our confidence in UDW-like particle detector models as effective descriptions for realistic probes.

5.4 Entanglement harvesting with modes of localized quantum fields

5.4.1 Typical mixedness of modes in localized quantum fields

In the previous section we argued that, upon further inspection, using local modes of a free scalar field is not a very promising idea to define a detector. This is, in large part, due to the fact that the probe system in this case will naturally diffuse and delocalize over time, which is not what we would like a truly localized detector to do. Just as we discussed in Part I, however, a minimal addition that we can include in this model is the external potential that effectively confines the probe field in a finite region of space. Using an external potential along the lines of what was introduced in Part I, we can also evaluate how fast the mixedness of a given localized mode varies with the size of the mode’s support, at a fixed strength of the potential. This is what we will address in what follows.

Let us consider a probe field ϕ_D in Minkowski space, and take the field mode defined at an instant of time $t = 0$ by the quadratures

$$\hat{Q} = \int d^d \mathbf{x} \hat{\phi}_D(0, \mathbf{x}) f(\mathbf{x}), \quad (5.65)$$

$$\hat{P} = \int d^d \mathbf{x} \hat{\pi}_D(0, \mathbf{x}) f(\mathbf{x}) \quad (5.66)$$

where $\hat{\pi}_D(t, \mathbf{x}) = \partial_t \hat{\phi}_D(t, \mathbf{x})$ is the conjugate momentum to $\hat{\phi}_D$, and $f(\mathbf{x})$ is a strongly localized function satisfying

$$\int d^d \mathbf{x} f(\mathbf{x})^2 = 1. \quad (5.67)$$

The $L^2(\mathbb{R}^d)$ normalization is of course chosen in order to guarantee $[\hat{Q}, \hat{P}] = i\mathbb{1}$. Once again, the mixedness of this mode is fully characterized by its symplectic eigenvalue, given by

$$\nu = 2\sqrt{\langle \hat{Q}^2 \rangle \langle \hat{P}^2 \rangle} \quad (5.68)$$

where the expectation values above are taken with respect to the global vacuum state of the probe field. By expanding the probe field on its basis of normal modes, we can write

$$\hat{Q} = \sum_{\mathbf{n}} c_{\mathbf{n}} \hat{\phi}_{\mathbf{n}}, \quad (5.69)$$

$$\hat{P} = \sum_{\mathbf{n}} c_{\mathbf{n}} \hat{\pi}_{\mathbf{n}} \quad (5.70)$$

where

$$c_{\mathbf{n}} = \int d^d \mathbf{x} f(\mathbf{x}) v_{\mathbf{n}}(\mathbf{x}), \quad (5.71)$$

and $v_{\mathbf{n}}(\mathbf{x})$ form a basis of real functions built from the spatial profiles $\Phi_{\mathbf{n}}(\mathbf{x})$, that can be chosen to be orthonormal. The normalization condition for $f(\mathbf{x})$ and the fact that the basis $\{v_{\mathbf{n}}(\mathbf{x})\}$ is orthonormal then implies that

$$\sum_{\mathbf{n}} c_{\mathbf{n}}^2 = 1. \quad (5.72)$$

The global state of the field is just the tensor product of the ground states of each of the basis modes

$$|0_{\text{D}}\rangle = \bigotimes_{\mathbf{n}} |0_{\mathbf{n}}\rangle, \quad (5.73)$$

and we know that for each mode we have

$$\langle \hat{\phi}_{\mathbf{n}}^2 \rangle = \frac{1}{2\omega_{\mathbf{n}}}, \quad (5.74)$$

$$\langle \hat{\pi}_{\mathbf{n}}^2 \rangle = \frac{\omega_{\mathbf{n}}}{2}. \quad (5.75)$$

Therefore, the expectation values $\langle \hat{Q}^2 \rangle$ and $\langle \hat{P}^2 \rangle$ can be very simply written as

$$\langle \hat{Q}^2 \rangle = \frac{1}{2} \sum_{\mathbf{n}} \frac{c_{\mathbf{n}}^2}{\omega_{\mathbf{n}}}, \quad (5.76)$$

$$\langle \hat{P}^2 \rangle = \frac{1}{2} \sum_{\mathbf{n}} \omega_{\mathbf{n}} c_{\mathbf{n}}^2. \quad (5.77)$$

By plugging this into (5.68), we obtain

$$\nu = \left[\left(\sum_{\mathbf{n}} \frac{c_{\mathbf{n}}^2}{\omega_{\mathbf{n}}} \right) \left(\sum_{\mathbf{m}} \omega_{\mathbf{m}} c_{\mathbf{m}}^2 \right) \right]^{1/2}. \quad (5.78)$$

The expression above makes it clear that ν satisfies $\nu \geq 1$ (as it must) thanks to the Cauchy-Schwarz inequality,

$$||u|| \cdot ||v|| \geq |\langle u, v \rangle| \quad (5.79)$$

where we take u and v to be vectors with components $u_{\mathbf{n}} = c_{\mathbf{n}}/\sqrt{\omega_{\mathbf{n}}}$ and $v_{\mathbf{n}} = \sqrt{\omega_{\mathbf{n}}}c_{\mathbf{n}}$ respectively, $\langle u, v \rangle$ is the standard Euclidean inner product, and we note that

$$\langle u, v \rangle = \sum_{\mathbf{n}} c_{\mathbf{n}}^2 = 1. \quad (5.80)$$

Furthermore, by recalling that the Cauchy-Schwarz inequality is saturated only when u and v are proportional to each other, we see that $\nu = 1$ (i.e., the chosen mode of the field is pure) only when its spatial profile $f(\mathbf{x})$ has nonzero overlap exclusively with basis modes of the same eigenfrequency. If $f(\mathbf{x})$ includes eigenfunctions associated with more than one frequency, the reduced state of the mode $(\hat{Q}, \hat{P})$ is mixed.

For a free field in Minkowski space, one can choose the basis of normal modes of the field to consist of plane waves (or sines/cosines if we want to select real spatial profiles), in which case $c_{\mathbf{n}}$ is directly related to the Fourier transform of $f(\mathbf{x})$. Since any compactly supported function in position space cannot be compactly supported in Fourier space, it follows that any localized spatial profile will necessarily include nonzero coefficients of arbitrarily high frequencies—which, according to the last comment made in the previous paragraph, means that the symplectic eigenvalue for the reduced state of the local mode $(\hat{Q}, \hat{P})$ must be strictly larger than 1. Once again, this is yet another way of seeing that any local mode of a free scalar field is strictly mixed, as we have pointed out earlier.

By adiabatically turning on an external confining potential, however, one can systematically take the localized mode closer to a true normal mode of the field, and thus effectively bring it closer to its ground state. For concreteness, let us now consider a localized mode whose spatial profile is given by

$$f(\mathbf{x}) = \frac{1}{\pi^{d/4} \sigma^{d/2}} e^{-|\mathbf{x}|^2/2\sigma^2}, \quad (5.81)$$

where σ is a quantity with dimensions of length which determines the effective size of the spatial profile of the mode (i.e., the radius of the region where the mode is strongly

supported). The probe field will be placed in an external potential $V(\mathbf{x})$ given by

$$V(\mathbf{x}) = \frac{|\mathbf{x}|^2}{2\ell^4}. \quad (5.82)$$

The spatial profile for the lowest-frequency mode of the field in this case is simply given by

$$v_0(\mathbf{x}) = \frac{1}{\pi^{d/4}\ell^{d/2}} e^{-|\mathbf{x}|^2/2\ell^2}. \quad (5.83)$$

Now, it should be clear that the spatial profile $f(\mathbf{x})$ simply corresponds to the wavefunction for the ground state of a harmonic oscillator of different natural frequency—i.e., a *squeezed* ground state. From our knowledge of the quantum harmonic oscillator, this immediately allows us to write down what the overlap coefficients $c_{\mathbf{n}}$ are. In terms of the quantum numbers in Cartesian coordinates $\mathbf{n} = (n_1, \dots, n_d) \in \mathbb{N}^d$ in d spatial dimensions, we have that only the all-even ones contribute:

$$c_{2\mathbf{n}}^2 = \frac{1}{(\cosh r)^d} \left(\frac{\tanh^2 r}{4} \right)^{\sum_i n_i} \prod_{i=1}^d \left(\frac{(2n_i)!}{(n_i!)^2} \right), \quad (5.84)$$

and the squeezing parameter r is related to the length scales σ and ℓ by

$$r = \log \left(\frac{\sigma}{\ell} \right). \quad (5.85)$$

Finally, we recall that the normal frequencies of the probe field are determined by

$$\omega_{\mathbf{n}} = \sqrt{m^2 + \frac{2}{\ell^2} \left(\sum_{i=1}^d n_i + \frac{d}{2} \right)}. \quad (5.86)$$

Putting this all together, we can write down the symplectic eigenvalue of the local mode

with spatial profile (5.81) as

$$\begin{aligned} \nu^2 = & \frac{1}{(\cosh^2 r)^d} \left(\sum_{\mathbf{n} \in \mathbb{N}^d} \sqrt{m^2 + \frac{2}{\ell^2} \left(2 \sum_{i=1}^d n_i + \frac{d}{2} \right)} \left(\frac{\tanh^2 r}{4} \right)^{\sum_i n_i} \prod_{i=1}^d \left(\frac{(2n_i)!}{(n_i!)^2} \right) \right) \\ & \times \left(\sum_{\mathbf{n} \in \mathbb{N}^d} \frac{1}{\sqrt{m^2 + \frac{2}{\ell^2} \left(2 \sum_{i=1}^d n_i + \frac{d}{2} \right)}} \left(\frac{\tanh^2 r}{4} \right)^{\sum_i n_i} \prod_{i=1}^d \left(\frac{(2n_i)!}{(n_i!)^2} \right) \right). \end{aligned} \quad (5.87)$$

By replacing $\sum n_i = n$, we can turn each expectation value into a single sum, and write

$$\begin{aligned} \nu^2 = & \frac{1}{(\cosh^2 r)^d} \left(\sum_{n=0}^{\infty} \sqrt{m^2 \ell^2 + 4n + d} \left(\frac{\tanh^2 r}{4} \right)^n F_d(n) \right) \\ & \times \left(\sum_{n=0}^{\infty} \frac{1}{\sqrt{m^2 \ell^2 + 4n + d}} \left(\frac{\tanh^2 r}{4} \right)^n F_d(n) \right), \end{aligned} \quad (5.88)$$

where we define

$$F_d(n) = \sum_{\substack{\mathbf{n} \in \mathbb{N}^d, \\ \sum n_i = n}} \prod_{i=1}^d \binom{2n_i}{n_i}. \quad (5.89)$$

The series can then be evaluated numerically, with the result in 1, 2, and 3 dimensions being displayed in Fig. 5.8. We see that, for a fixed value of the ratio σ/ℓ , the mode deviates from purity faster in higher dimensions, but for a reasonably large interval—with σ and ℓ differing by almost a factor of 10—the symplectic eigenvalue ν stays within 5% of perfect purity.

5.4.2 Results for entanglement harvesting

For concreteness, we now present specific examples of entanglement harvesting using two localized quantum field theories. Specifically, we consider the lowest energy modes of 1) fields under the influence of quadratic potentials and 2) fields in cubic boxes with Dirichlet boundary conditions. We will see that indeed it is possible for these localized quantum fields




Figure 5.8: Symplectic eigenvalue of the mode defined with the spatial profile (5.81) in 1, 2, and 3 dimensions, as a function of the ratio σ/ℓ . For concreteness, the mass of the field was set so that $m\ell = 10$ in all cases.

to extract entanglement from the vacuum of a free Klein-Gordon field in 3 + 1 Minkowski spacetime.

As our first example of fully relativistic entanglement harvesting, we consider two localized quantum fields in Minkowski spacetime under the influence of potentials $V_A(\mathbf{x}) = |\mathbf{x}|^2/2\ell^4$ and $V_B(\mathbf{x}) = V_A(\mathbf{x} - \mathbf{L})$, where $L = |\mathbf{L}|$ denotes the proper distance between the centers of the trapping potentials. In essence, the two quantum fields are identical, apart from a spatial shift in the potentials that confine them.

Both fields will interact with a free scalar field $\hat{\psi}(\mathbf{x})$ according to the interaction Hamiltonian of Equation (3.15), where the functions $\zeta_A(\mathbf{x})$ and $\zeta_B(\mathbf{x})$ will conveniently be prescribed as

$$\zeta_A(\mathbf{x}) = \zeta_B(\mathbf{x}) = e^{-\frac{\pi t^2}{2T^2}}. \quad (5.90)$$

This corresponds to interactions that are adiabatically switched on, and peak at $t = 0$. The effective time of the switching is controlled by the timescale T . The reason that we consider $\zeta_A(\mathbf{x}) = \zeta_B(\mathbf{x})$ independent of the spatial coordinates is that the effective spacetime region where the localized fields interact with $\hat{\psi}(\mathbf{x})$ is defined by the product of $\zeta_A(\mathbf{x})$ with the mode localization of the fields. The spatial localization of the modes together with the time localization of $\zeta_A(\mathbf{x})$ then gives an overall interaction that is localized in spacetime for each mode.

We consider the three fields to start in their respective ground states, $|0_A\rangle \otimes |0_B\rangle \otimes |0\rangle$, and we assume that we only have access to the localized fields' mode excitations with the lowest energy, $\omega_{0_A} = \omega_{0_B} \equiv \Omega$. Due to the general result of Part I, the final state of the two detectors will take the same form at leading order in perturbation theory as the usual

UDW case with two harmonic oscillator detectors, and therefore its negativity will take the same form of Equation (5.17),

$$\mathcal{N}(\hat{\rho}_D) = \max(0, |\mathcal{M}| - \mathcal{L}) + \mathcal{O}(\lambda^4), \quad (5.91)$$

where the $\mathcal{L}$ and $\mathcal{M}$ terms are given by

$$\begin{aligned} \mathcal{L} &= \lambda^2 \int dV dV' \Lambda_A(\mathbf{x}) \Lambda_A(\mathbf{x}') e^{-i\Omega(t-t')} W(\mathbf{x}, \mathbf{x}'), \\ &= \lambda^2 \int dV dV' \Lambda_B(\mathbf{x}) \Lambda_B(\mathbf{x}') e^{-i\Omega(t-t')} W(\mathbf{x}, \mathbf{x}'), \\ \mathcal{M} &= -\lambda^2 \int dV dV' \Lambda_A(\mathbf{x}) \Lambda_B(\mathbf{x}') e^{i\Omega(t+t')} G_F(\mathbf{x}, \mathbf{x}'), \end{aligned} \quad (5.92)$$

where we have also denoted $G_F(\mathbf{x}, \mathbf{x}') \equiv \langle \mathcal{T}(\hat{\psi}(\mathbf{x}) \hat{\psi}(\mathbf{x}')) \rangle$. With this choice, the eigenfrequency of the ground state can be simply obtained by the spectrum of the three-dimensional harmonic oscillator, and gives us $\Omega = \sqrt{m^2 + 3/\ell^2}$. Moreover, the spacetime smearing functions are given by

$$\begin{aligned} \Lambda_A(\mathbf{x}) &= \zeta_A(\mathbf{x}) \Phi_{\mathbf{0}_A}^A(\mathbf{x}) = e^{-\frac{\pi t^2}{2T^2}} \left(\frac{1}{\pi \ell^2} \right)^{\frac{3}{4}} \frac{e^{-\frac{|\mathbf{x}|^2}{2\ell^2}}}{(m^2 + \frac{3}{\ell^2})^{1/4}}, \\ \Lambda_B(\mathbf{x}) &= \zeta_B(\mathbf{x}) \Phi_{\mathbf{0}_B}^B(\mathbf{x}) = \Lambda_A(t, \mathbf{x} - \mathbf{L}). \end{aligned} \quad (5.93)$$

We then see that the effective size of the interaction region can be estimated by looking at the standard deviation of the space-dependent Gaussian function in Equation (5.93). In this case, the spatial size of the interaction region can be estimated to be $\sigma \sim \ell$, so that smaller values of the parameter ℓ that defines the confining potential correspond to more localized detectors.

Motivated by the discussion of entanglement harvesting, we focus on the case where the interaction regions are approximately spacelike separated¹⁰, so that entanglement acquired by the localized modes via communication can be neglected. For this reason, we consider $L = 5T$ in the specific example that we explore here, where we verified that the effect of communication is 6 orders of magnitude smaller than the effects arising from vacuum entanglement harvesting [176, 135]. In Fig. 5.9 we plot the entanglement acquired by the

¹⁰Even though the Gaussian tails of the switching are technically infinitely long, using these switchings is effectively equivalent to considering compactly supported switchings. For more details about effective communication from Gaussian tails and how the use of Gaussians is justified, we refer the reader to [176, 58].




Figure 5.9: The negativity of the state of two localized quantum fields confined by a quadratic potential when restricted to their lowest energy after interacting with a massless scalar field. The negativity is plotted as a function of the energy of the modes $\Omega = \sqrt{m^2 + 3/\ell^2}$. The time duration of the interaction T is used as a scale. The separation between the interaction regions of each detector for these plots is $L = 5T$.

localized modes as a function of their energy gap. The plot is what is expected for the behavior of entanglement harvesting in the Minkowski vacuum, where there is a threshold in the energy gap below which no entanglement can be extracted. For ΩT above this threshold, the entanglement peaks and quickly decays. The behavior seen in Fig. 5.9 is identical to that of a UDW detector with a Gaussian spacetime smearing function, as expected (for comparison, see e.g. [144]).




Figure 5.10: The negativity of the state of two localized quantum fields in boxes of sides d when restricted to their lowest energy mode after spacelike interaction with a massless scalar field. The negativity is plotted as a function of the energy of the localized mode, $\Omega = \sqrt{m^2 + 3\pi/d^2}$. The time duration of the interaction T is used as a scale. The separation between the interaction regions of each detector for these plots is $L = 4.5T$.

Finally, in Fig. 5.10 we consider the case where two massive fields in cubic cavities of size length d with Dirichlet boundary conditions interact with a free massless scalar

field. This leads to mode functions for the field which have the same spatial profile as of a nonrelativistic quantum particle in a box. We consider the same choices of $\zeta_A(\mathbf{x})$ and $\zeta_B(\mathbf{x})$ as in Equation (5.90). We also restrict the localized fields to the lowest energy mode $\mathbf{1}_A = \mathbf{1}_B = (1, 1, 1)$, with energy $\omega_{\mathbf{1}_A}^A = \omega_{\mathbf{1}_B}^B = \sqrt{m^2 + 3\pi^2/d^2}$. To ensure that communication between the detectors is negligible, we pick $d \sim 0.5T$ and consider the distance between the cavities to be given by $L = 5T$. The negativity in this case can be seen in Fig. 5.10 as a function of the energy gap of the $\mathbf{1}_A, \mathbf{1}_B$ modes of the fields. The behavior of the negativity is similar to most cases of entanglement harvesting in spacelike separated regions. We see more entanglement in this setup due to the smaller choice of L , which can be taken in this case because the communication between the detectors is naturally smaller due to the compact support of the modes in space.

Figs. 5.9 and 5.10 showcase examples where two completely relativistic (microcausal) localized quantum field theories can extract entanglement from the vacuum of a free field. This proves that entanglement harvesting is not a consequence of any features of effective non-relativistic theories, and it is indeed a prediction of quantum field theory.

Chapter 6

Optimal entanglement extraction from a QFT

The protocol of entanglement harvesting discussed in the previous chapter can certainly be a useful tool to establish the existence of field entanglement in a way that is concretely accessible by local probes. However, it must be acknowledged there are limits to the detectors' ability to fully capture the entanglement structure originally present in the field. For starters, the correlations acquired by the detectors after coupling to the field will generally be affected by local noise, which invariably makes the entanglement between the probes less robust than what the field contains [154, 86]. Moreover, the entanglement in the QFT may be distributed nontrivially among the degrees of freedom the detector is coupled to, and it is in principle not clear that a simple/naive choice of coupling between probe and field will be enough to reveal this structure. If we want to use the entanglement acquired by the detectors as an operational way of accurately characterizing the entanglement in the field theory, it is thus important to know to what extent the finite degrees of freedom of the probe system can optimally encode the entanglement originally present in the field they coupled to.

In this chapter, we will analyze how to optimally transfer features of the entanglement structure of a quantum field to particle detectors that couple to the field in localized regions of spacetime. We will first provide upper bounds to how much entanglement a fixed number of detectors can extract from the field in a given region. We will then proceed to show how the spatial distribution of the extracted entanglement behaves, and use this to prescribe the coupling between detectors and field (for any local subregion) which will best reflect the entanglement that the given spatial region had with its complement.

This chapter is organized as follows. Section 6.1 reviews the basic elements of a measurement framework for quantum fields using particle detectors, and lays out the setup for the detector and field that we will employ in this chapter. Section 6.2 describes a discretized version of the field theory as a chain of harmonic oscillators, which will be convenient for the use of numerical methods in the description of the system. We also review basic features of the entanglement between subregions in the harmonic chain which we aim to capture with the aid of particle detectors. Section 6.3 establishes the bound on how much entanglement can be obtained from a given region of the field, given the number of degrees of freedom accessible in the probes. This will lead to a prescription for the coupling between the field and a given number of detectors, in order for the final state of the probes to reflect as much as it can of the entanglement between the coupling region and the rest of space. As a proof of principle, we show in Section 6.4 that, with the preferred choice of coupling outlined in Section 6.3, the entanglement between the detector degrees of freedom and the rest of space displays an area-law behavior, as expected for entanglement entropies between complementary regions in ground states of systems with short-range interactions. Section 6.5 addresses how the results obtained in Sections 6.3 and 6.4 can be adapted to reflect the case of a continuum field theory, even though they were mostly explored in the context of a discretized version of the field as described in Section 6.2.

6.1 Extracting modes of the field with particle detectors

For concreteness, we will once again focus on probing a free real scalar field in Minkowski space. We will choose to fix an inertial coordinate system $(t, \mathbf{x})$, such that the free dynamics of the field in $d + 1$ spacetime dimensions can be described in a given inertial coordinate system $(t, \mathbf{x})$ by the Hamiltonian

$$\hat{H}_\phi = \frac{1}{2} \int d^d \mathbf{x} \left(\hat{\pi}^2 + (\nabla \hat{\phi})^2 + m^2 \hat{\phi}^2 \right) \quad (6.1)$$

where m is the field mass and $\hat{\pi} := \partial_t \hat{\phi}$ is the field's canonical momentum. $\hat{\phi}$ and $\hat{\pi}$ satisfy the equal-time canonical commutation relations

$$\left[\hat{\phi}(t, \mathbf{x}), \hat{\pi}(t, \mathbf{x}') \right] = i \delta^{(d)}(\mathbf{x} - \mathbf{x}'). \quad (6.2)$$

The local probes (i.e., the detectors) will once again be taken to be harmonic oscillators, which we will assume to be comoving with the frame $(t, \mathbf{x})$. For each detector we have a pair

of dimensionless quadratures $(\hat{q}_d, \hat{p}_d)$ satisfying $[\hat{q}_d, \hat{p}_d] = i\mathbb{1}$ and free detector Hamiltonian

$$\hat{H}_d = \frac{\Omega}{2} (\hat{p}_d^2 + \hat{q}_d^2), \quad (6.3)$$

which will now couple to the field through an interaction that will be of the form

$$\begin{aligned} \hat{H}_I = \lambda \chi(t) & \left(\lambda_q \hat{q}_d \int d^d \mathbf{x} f(\mathbf{x}) \hat{\pi}(t, \mathbf{x}) \right. \\ & \left. - \lambda_p \hat{p}_d \int d^d \mathbf{x} g(\mathbf{x}) \hat{\phi}(t, \mathbf{x}) \right). \end{aligned} \quad (6.4)$$

Here, λ is a dimensionless coupling constant, λ_q and λ_p are scales adjusting the dimensions of each term¹, and $\chi(t)$ is a switching function which dictates the time dependence of the coupling between the detector and the field. $f(\mathbf{x})$ and $g(\mathbf{x})$ correspond to the (spatial) smearing functions—the spatial profile of the interaction—for the coupling to the field and its conjugate momentum, respectively. Naturally, if there are multiple detectors, the free detector Hamiltonian contains a sum of terms like (6.3), and the interaction Hamiltonian will contain a sum of terms like (6.4), with possibly different choices of f and g for each detector.

If the smearing functions $f(\mathbf{x})$ and $g(\mathbf{x})$ satisfy

$$\lambda_q \lambda_p \int d^d \mathbf{x} f(\mathbf{x}) g(\mathbf{x}) = 1, \quad (6.5)$$

then we can interpret the operators

$$\hat{Q} := \lambda_p \int d^d \mathbf{x} g(\mathbf{x}) \hat{\phi}(t, \mathbf{x}), \quad (6.6)$$

$$\hat{P} := \lambda_q \int d^d \mathbf{x} f(\mathbf{x}) \hat{\pi}(t, \mathbf{x}) \quad (6.7)$$

as the quadratures that define one single mode in the field, since they satisfy $[\hat{Q}, \hat{P}] = i\mathbb{1}$. With this, the interaction Hamiltonian (6.4) can be seen as a coupling between the detector and one particular mode of the field.

¹Since we consider that the switching and smearing functions are L^1 -normalized functions, $[\chi(t)] = L^{-1}$, $[f(\mathbf{x})] = [g(\mathbf{x})] = L^{-d}$, this means that $[\lambda_q] = [\pi]^{-1} = L^{\frac{d+1}{2}}$ and $[\lambda_p] = [\phi]^{-1} = L^{\frac{d-1}{2}}$. Some natural parameters that could be used to adjust the units would be the field mass, m , or the detector's frequency, Ω , with $[m] = [\Omega] = L^{-1}$.

The time evolution operator for the dynamics between detector and field in this setup simplifies drastically when we consider the interaction Hamiltonian to be very intensely peaked around a very short period of time, with the final state being examined immediately after the interaction takes place. Mathematically, this amounts to taking $\chi(t) = \delta(t)$ in (6.4), where the internal dynamics generated by the free Hamiltonians can be neglected, and the full unitary time evolution operator will simply be given by

$$\hat{U} = \exp \left[-i\lambda \left(\hat{q}_d \hat{P} - \hat{p}_d \hat{Q} \right) \right]. \quad (6.8)$$

This will simplify even further if we take $\lambda = -\pi/2$, in which case the dynamics implemented by

$$\hat{U}_{\text{SWAP}} = \exp \left[i \frac{\pi}{2} (\hat{q}_d \hat{P} - \hat{p}_d \hat{Q}) \right] \quad (6.9)$$

will act on the quadratures as

$$\begin{aligned} \hat{q}_d &\mapsto \hat{Q} \\ \hat{Q} &\mapsto -\hat{q}_d \\ \hat{p}_d &\mapsto \hat{P} \\ \hat{P} &\mapsto -\hat{p}_d \end{aligned} \quad (6.10)$$

and will therefore swap the phase spaces between the mode $(\hat{q}_d, \hat{p}_d)$ and the mode $(\hat{Q}, \hat{P})$. This is why we refer to the Gaussian unitary (6.9) as a *swap operator*. With the dynamics reduced to a swap between the detector and some mode of the field, we guarantee that the final state of the detector will be exactly the initial state of the field mode that is probed. In other words, the final state of the detector can be immediately identified as the reduced state of the field in the selected mode $(\hat{Q}, \hat{P})$. Therefore, the swap operation allows us to speak about what a particle detector can measure in terms of specific field modes—or, in short, it will allow one to effectively extract the state of a mode of the field to the state of a detector. This idea will guide how we will approach the rest of the problem in this chapter, as we will now focus on talking about entanglement properties of select modes of the field theory.

6.2 Entanglement in the ground state of a harmonic chain

In order to employ explicit analytical and numerical methods, it will be useful for our purposes to consider first a setup where we can use the tools of Gaussian quantum mechanics

in full force. For that, we replace the field theory (which contains infinitely many degrees of freedom) with a lattice of harmonic oscillators, equivalent to introducing an IR and a UV cutoff in the theory. We will address how the results can be adapted to the approach continuum case in Section 6.5. The Hamiltonian (6.1) then takes the discrete form

$$\hat{H}_\phi = \frac{1}{2} \sum_i \omega (\hat{q}_i^2 + \hat{p}_i^2) - \frac{\alpha}{2} \sum_{\langle i,j \rangle} \hat{q}_i \hat{q}_j. \quad (6.11)$$

The first term is a sum over all the sites of the n -dimensional lattice, and the notation $\langle i, j \rangle$ in the second term denotes a sum over nearest neighbors. Each pair $(\hat{q}_i, \hat{p}_i)$ corresponds to dimensionless quadratures, and ω represents the decoupled (free) frequency of each local oscillator, which is the same for all sites. One can see the lattice Hamiltonian (6.11) as a discrete approximation² to the Hamiltonian (6.1) through the replacements

$$\begin{aligned} \text{spatial coordinate } \mathbf{x} &\mapsto \text{discrete label } i, \\ \int d^d x &\mapsto \varepsilon^d \sum_i, \\ \delta^{(d)}(\mathbf{x} - \mathbf{y}) &\mapsto \frac{1}{\varepsilon^d} \delta_{ij}, \\ (\nabla \hat{\phi}(\mathbf{x}))^2 &\mapsto \frac{1}{2\varepsilon^2} \sum_j \left(\hat{\phi}_i - \hat{\phi}_j \right)^2 \end{aligned} \quad (6.12)$$

where ε is the short-distance cutoff length scale, and in the last expression the sum in j runs over all nearest neighbors of the site with label i . With these replacements, the free frequency ω is related to the field mass m and the cutoff ε by

$$\omega^2 = m^2 + \frac{2d}{\varepsilon^2}, \quad (6.13)$$

²Note that (6.11) is not exactly the result of applying UV and IR cutoffs to (6.1): there is an approximation in taking the derivative of the field to be given by a first-neighbor discrete derivative. This is not a problem because the continuum limit of this approximation is still the continuous derivative and (6.11) indeed becomes (6.1) in the limit $\varepsilon \rightarrow 0$. For a deeper discussion on this subtlety see, e.g., [83].

the quadratures $(\hat{q}_i, \hat{p}_i)$ at each lattice site are related to the discretized version of $\hat{\phi}$ and $\hat{\pi}$ by

$$\hat{\phi}_i = \frac{1}{\sqrt{\omega \varepsilon^d}} \hat{q}_i, \quad (6.14)$$

$$\hat{\pi}_i = \sqrt{\frac{\omega}{\varepsilon^d}} \hat{p}_i, \quad (6.15)$$

and the effective lattice coupling constant α is given by

$$\alpha = \frac{1}{\omega \varepsilon^2}. \quad (6.16)$$

Since $m^2 \geq 0$, Equations (6.13) and (6.16) imply that, for a given value of ω , the coupling α is constrained by $\alpha \leq \alpha_c \equiv \omega/2d$, with the inequality being saturated in the case of a massless field. If α were greater than α_c , the Hamiltonian (6.11) would not be bounded below.

Imposing a short-distance cutoff places a finite number of modes per unit volume in the system. By then adding an IR cutoff (say, by placing the field in a box with a large but finite side length) the total number of oscillators is made finite, making the problem more amenable to Gaussian methods. Since the interactions are short-range (only nearest neighbors coupling), the local features of the system will not be affected by the presence of the boundary for subregions of sufficiently small size compared to the full size of the system.

Since the Hamiltonian is quadratic, the ground state of the lattice is Gaussian and hence fully characterized by its *covariance matrix*. For a general quantum state $\hat{\rho}$ of a system of N bosonic degrees of freedom, the covariance matrix is a $2N \times 2N$ matrix given by³

$$\sigma^{ab} = \langle \hat{\xi}^a \hat{\xi}^b + \hat{\xi}^b \hat{\xi}^a \rangle_{\hat{\rho}}, \quad (6.17)$$

where $\hat{\xi}^a$ is an operator-valued on phase space that encompasses all the quadratures of the system,

$$\hat{\xi}^a = (\hat{q}^1, \dots, \hat{q}^N, \hat{p}_1, \dots, \hat{p}_N)^\top \equiv (\hat{\mathbf{q}}, \hat{\mathbf{p}})^\top. \quad (6.18)$$

³This expression for the covariance matrix assumes that the vector of means $z^a \equiv \langle \hat{\xi}^a \rangle_{\hat{\rho}}$ vanishes, which is a condition that will hold for the entirety of this chapter. The slightly more general expression of the covariance matrix will be used later in Part III of this thesis.

By writing $\hat{H}_\phi$ as

$$\hat{H}_\phi = \frac{\omega}{2} \sum_i \hat{p}_i^2 + \frac{1}{2} \hat{\mathbf{q}}^\top \mathbf{V} \hat{\mathbf{q}}, \quad (6.19)$$

where the symmetric matrix $\mathbf{V}$ has its elements defined by

$$V_{ij} = \begin{cases} \omega, & i = j \\ -\alpha, & i, j \text{ nearest neighbors} \\ 0 & \text{otherwise,} \end{cases} \quad (6.20)$$

the ground state covariance matrix σ_ϕ can be written in the ordered basis $(\hat{\mathbf{q}}, \hat{\mathbf{p}})$ as [11]

$$\sigma_\phi = \begin{pmatrix} \sqrt{\omega} \mathbf{V}^{-1} & \mathbf{0} \\ \mathbf{0} & \sqrt{\omega^{-1}} \mathbf{V} \end{pmatrix}. \quad (6.21)$$

Thanks to Williamson's theorem [195], we know that every covariance matrix can be written as

$$\sigma = S \sigma_D S^\top \quad (6.22)$$

where σ_D is of the form

$$\sigma_D = \bigoplus_{i=1}^n \begin{pmatrix} \nu_i & 0 \\ 0 & \nu_i \end{pmatrix} \quad (6.23)$$

and S is a *symplectic matrix*—i.e., a matrix satisfying

$$S \tilde{\Omega} S^\top = \tilde{\Omega} \quad (6.24)$$

with $\tilde{\Omega}$ being the (inverse of the) symplectic matrix,

$$\tilde{\Omega} = \begin{pmatrix} 0_{n \times n} & \mathbf{1}_{n \times n} \\ -\mathbf{1}_{n \times n} & 0_{n \times n} \end{pmatrix}. \quad (6.25)$$

The set $\{\nu_i\}_{i=1}^n$ is called the *symplectic spectrum* of σ , and the elements of the symplectic spectrum are known as the *symplectic eigenvalues*. It can be shown that for Gaussian states, the condition that the density matrix $\hat{\rho}$ is positive-definite reduces to the condition $\nu_i \geq 1$ for every i , with the state being pure when all symplectic eigenvalues are equal to 1, and mixed otherwise.

The symplectic matrix (6.25) encodes the canonical commutation relations between the quadratures—i.e., the positions and canonical momenta—that constitute the degrees of

freedom of the system. A symplectic transformation can thus be interpreted as a linear transformation that preserves the canonical commutation relations of the quadratures, and therefore maps the original set of modes of the system to a new set of modes that are equally valid in the description of the dynamics. Williamson's theorem then implies that one can always find a basis of canonical coordinates on phase space (i.e., a choice of modes) such that the covariance matrix in that basis is fully decoupled. We refer to this set of modes as the *normal mode basis*, since they are mathematically analogous to the set of normal modes in terms of which a given time-independent quadratic Hamiltonian becomes just a sum of decoupled single-mode Hamiltonians. In fact, for the covariance matrix (6.21), the normal modes of the system in the sense above are also the normal modes of the Hamiltonian.

Once the covariance matrix of the ground state is known, characterizing entanglement measures is in principle straightforward [162, 188, 1]. Since the ground state is pure, the entanglement between a set of modes and its complement can be quantified by the entanglement entropy, which is the von Neumann entropy of either of the subsystems in the bipartition. To obtain the covariance matrix of any subsystem (which will also fully determine its state, since any subset of modes of a system in a Gaussian state will also be in a Gaussian state), we simply remove all the lines and columns of the global covariance matrix (6.21) which are not associated to the modes in question. The von Neumann entropy of the resulting density matrix $\hat{\rho}$ associated to the subsystem's covariance matrix σ is then given by

$$S(\hat{\rho}) = \sum_i \left[\frac{\nu_i + 1}{2} \log \left(\frac{\nu_i + 1}{2} \right) - \frac{\nu_i - 1}{2} \log \left(\frac{\nu_i - 1}{2} \right) \right], \quad (6.26)$$

which only depends on the symplectic eigenvalues ν_i of the reduced covariance matrix σ .

In the case where we want to characterize the entanglement between two subsets of modes (let us call them A and B) that do not form the global system, we will adopt the logarithmic negativity, which attains a convenient form in the case of Gaussian states. On phase space, transposition amounts to time reversal, which acts on the quadrature operators by reversing the sign of $\hat{P}_i$ and leaving $\hat{Q}^i$ unchanged. To represent this transformation more succinctly, it is more convenient to reorder the rows and columns of the covariance matrix by taking permutations from the $(\hat{\mathbf{q}}, \hat{\mathbf{p}})$ basis to the ordered basis $(\hat{q}_1, \hat{p}_1, \dots, \hat{q}_{2N}, \hat{p}_{2N})$, where the first N_A pairs refer to quadratures of subsystem A , and the last N_B pairs correspond to subsystem B . The transformation to the covariance matrix resulting from partial transposition with respect to subsystem B in this case can thus be written as

$$\tilde{\sigma}_{AB} = (\mathbb{1}_{2N_A} \oplus T_B) \sigma_{AB} (\mathbb{1}_{2N_A} \oplus T_B) \quad (6.27)$$

where $\mathbb{1}_{2N_A}$ is the $2N_A \times 2N_A$ identity matrix, and T_B is given by

$$T_B = \bigoplus_{i=1}^{N_B} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (6.28)$$

For a bipartite Gaussian state, the logarithmic negativity between subsystems A and B is fully determined by the symplectic spectrum of $\tilde{\sigma}_{AB}$. The statement that $\hat{\rho}^{T_B}$ is not positive semidefinite then turns into the statement that some of the symplectic eigenvalues of $\tilde{\sigma}_{AB}$ are smaller than 1. In this case, the logarithmic negativity is given by

$$E_N(\hat{\rho}) = \sum_i F(\tilde{\nu}_i), \quad (6.29)$$

where $\{\tilde{\nu}_i\}$ is the symplectic spectrum of $\tilde{\sigma}_{AB}$, and we define $F(x) = -\log(x)$ for $x \in (0, 1]$ and $F(x) = 0$ for $x > 1$.

6.3 Upper bound to entanglement extraction with a finite set of modes

In this section, we will prescribe how to select, within any localized region A of a harmonic lattice and for a given integer N , the set of N modes supported in A which are most entangled with the region's complement $\bar{A}$. This will naturally lead to a prescription for the set of modes that local probes with N bosonic degrees of freedom must couple to in order to be able to extract the most entanglement with $\bar{A}$. As we will see in Section 6.5, since a preferred choice of modes can be reinterpreted as a preferred smearing function for the detector's coupling to the field, these results can be directly rephrased as a recipe for the structure of the coupling between probes and field that will optimally transfer entanglement from the field to the detectors.

The key to the recipe we are about to describe comes from the following result, which in turn derives from how the entanglement contained in a globally pure Gaussian state is distributed across (Williamson) normal modes⁴ supported in complementary regions.

⁴(Williamson) normal modes should not be confused with the mechanical normal modes of a quadratic interaction Hamiltonian (which are decoupled modes of free oscillation). In this manuscript, following common convention, we will use the name normal modes always referring to the modes obtained under a canonical change of variables for which the covariance matrix is diagonal (i.e. the modes are not correlated).

Theorem 6.3.1. *Let A be a local subregion containing N_A modes of an overall pure Gaussian state, and let $\bar{A}$ be the region that is complementary to A . Then, the entanglement (as quantified by the logarithmic negativity) between the region $\bar{A}$ and any collection of $m \leq N_A$ modes supported entirely within A is upper-bounded by the entanglement between $\bar{A}$ and the set of m most mixed normal modes of the reduced covariance matrix of subregion A .*

Proof. First we note that, by definition, the covariance matrix of subregion A in the basis of normal modes takes the form of a direct sum, with all modes in that basis being completely uncorrelated with each other. Because the joint state of $A\bar{A}$ is pure, however, the normal modes of A must be purified by another set of modes that are fully supported on $\bar{A}$. The purification, in fact, works mode-by-mode: for each normal mode of A , there is a separate normal mode of the complementary region $\bar{A}$ that shares all of its entanglement with the given mode in A , such that the two modes paired together are in a product state with the rest of the system.⁵

In order to analyze entanglement between normal modes in A and $\bar{A}$, we first build the covariance matrix $\tilde{\sigma}_{A\bar{A}}$ associated to the partially transposed density matrix $\hat{\rho}_{A\bar{A}}^{T_{\bar{A}}}$, in the basis of normal modes of A and $\bar{A}$. In this basis, it is particularly easy to see that the symplectic spectrum of $\tilde{\sigma}_{A\bar{A}}$ —from which the logarithmic negativity between A and $\bar{A}$ is computed—can be fully expressed in terms of the symplectic eigenvalues of σ_A . More specifically, the i -th lowest symplectic eigenvalue $\tilde{\nu}_i$ of $\tilde{\sigma}_{A\bar{A}}$ is related to the i -th *largest* symplectic eigenvalue ν_i of σ_A through

$$\tilde{\nu}_i = \nu_i - \sqrt{\nu_i^2 - 1}. \quad (6.30)$$

A proof of this statement can be found in Appendix B.1. One readily sees that $\nu_i \geq 1$ automatically guarantees that $\tilde{\nu}_i \leq 1$, and that $\tilde{\nu}_i$ is a monotonically decreasing function of ν_i .

A change of basis on the modes supported in A will transform the covariance matrix of $A\bar{A}$ as

$$\sigma'_{A\bar{A}} = (S_A \oplus \mathbb{1}_{2N_{\bar{A}}}) \sigma_{A\bar{A}} (S_A^T \oplus \mathbb{1}_{2N_{\bar{A}}}), \quad (6.31)$$

⁵This is actually a special case of a more general fact: in a pure Gaussian state, for *any* single mode (not necessarily a normal mode of a subregion’s covariance matrix), one can define a second mode with the property that the joint state of the two modes fully decouples from the rest. This pair of modes that completely purify each other is known as a *partner mode* pair [100, 180, 181]. What is special about partner mode pairs comprising normal modes of complementary subregions is that, in this case in particular, the spatial supports of each mode in the pair do not overlap, and each mode is completely uncorrelated from the rest of the modes in the region where it is supported [22, 23, 198].




Figure 6.1: Spatial profile of the most mixed normal mode of a given region (in red) and its complement (in blue) as given by the occupation function F_i . Subfigures a) and b) show this for a lattice in 1 and 2 spatial dimensions respectively. Subfigures c) and d) are both in 3D, with c) being a full three-dimensional visualization of the mode and its complement, and d) showing two-dimensional slices at different heights of the plot at the top.

where S_A is some symplectic transformation acting only on A . The logarithmic negativity between regions A and $\bar{A}$ can also be computed from $\sigma'_{A\bar{A}}$ according to the symplectic spectrum of $\tilde{\sigma}'_{A\bar{A}}$, obtained by

$$\begin{aligned}\tilde{\sigma}'_{A\bar{A}} &= (\mathbb{1}_{2N_A} \oplus T_{\bar{A}}) \sigma'_{A\bar{A}} (\mathbb{1}_{2N_A} \oplus T_{\bar{A}}) \\ &= (S_A \oplus \mathbb{1}_{2N_{\bar{A}}}) \tilde{\sigma}_{A\bar{A}} (S_A^T \oplus \mathbb{1}_{2N_{\bar{A}}}).\end{aligned}\tag{6.32}$$

Here, $T_{\bar{A}}$ is the operator that performs a local time reversal on the modes in $\bar{A}$ —i.e., it sends $\hat{P}$ to $-\hat{P}$ for every momentum operator supported in region $\bar{A}$, and does nothing to the position operators. Since symplectic transformations do not change the symplectic spectrum, we know that the symplectic eigenvalues of $\tilde{\sigma}'_{A\bar{A}}$ are the same as those of $\tilde{\sigma}_{A\bar{A}}$, as they should be.

Now consider any collection of m modes supported within region A . These can be

defined by applying an appropriate symplectic transformation that sends the basis of normal modes of subsystem A to a basis that contains the list of desired modes, after which we discard the extra $N_A - m$ degrees of freedom. The logarithmic negativity between the chosen collection of modes in A and the entire complementary region $\bar{A}$ is similarly extracted from the symplectic spectrum of the covariance matrix obtained after discarding the rows and columns of $\tilde{\sigma}'_{A\bar{A}}$ that correspond to the modes we threw away.

We can now invoke the interlacing theorem for symplectic eigenvalues [112, 16]. This theorem (which we state in more detail as Theorem B.1.1 in Appendix B.1) guarantees that the j -th lowest symplectic eigenvalue of $\tilde{\sigma}'_{A\bar{A}}$ lower-bounds the j -th lowest symplectic eigenvalue of any matrix obtained from $\tilde{\sigma}'_{A\bar{A}}$ by discarding the rows and columns associated with one of its modes. By applying this result $N_A - m$ times, and using that the logarithmic negativity for a Gaussian state is given by Equation (6.29) with $F(x)$ being a monotonically decreasing function in the interval $(0, 1)$, it then follows that the maximum logarithmic negativity between m modes of region A and the complementary region $\bar{A}$ cannot be larger than the logarithmic negativity between $\bar{A}$ and the m most mixed normal modes of A . This completes the proof that, indeed, the largest amount of entanglement between a collection of m modes in A and the complementary region $\bar{A}$ is attained by choosing the m most mixed normal modes of region A . $\square$

Since the reduced covariance matrix on any subregion of the lattice does not contain correlations between the position and momentum variables, the quadratures for the normal modes $(\hat{Q}_i, \hat{P}_i)$ can be expressed as linear combinations of the quadratures of the local harmonic oscillators that build the subregion A being probed. That is, for any normal mode $(\hat{Q}, \hat{P})$, we can write

$$\hat{P} = \sum_i f_i \hat{p}_i, \quad \hat{Q} = \sum_i g_i \hat{q}_i \quad (6.33)$$

where f_i and g_i are the spatial profiles of the normal mode in momentum and position, respectively. The spatial profile of the most mixed normal mode of a typical subregion in 1, 2, and 3 spatial dimensions are shown in Fig. 6.1. In each case, the parameters ω and α from Equation (6.11) were chosen such that the ratio α/ω is within 0.1% of the critical value $1/2n$, with n being the number of spatial dimensions of the lattice. This corresponds to the (almost) massless limit of the field theory⁶. The quantity being plotted in Fig. 6.1 is the *mode occupation function*, which is given at each site i by $F_i := f_i g_i$, with g_i and f_i defined

⁶Notice that if we increase the mass of the field, since the correlations of the region with its complement decay exponentially with the field mass, we would expect the most mixed modes to quickly become more localized close to the boundary.

as in (6.33). It is clear that the support of the normal modes is strongly concentrated near the boundary of the region. This is to be expected, since the Hamiltonian only contains nearest-neighbor interactions, and therefore the entanglement present in the ground state between a region and its complement is mostly shared across the region's boundary.

Once we have established the list of m modes within region A that are most entangled with $\bar{A}$, it is then straightforward to prescribe an idealized coupling that transfers their entanglement to a series of m probes whose coupling to the lattice is restricted to lie strictly on A . For a detector with m bosonic degrees of freedom—or, equivalently, for m single-mode detectors—given by $(\hat{q}_i, \hat{p}_i)$, we take the m most mixed normal modes of region A , which we will denote by $(\hat{Q}_i, \hat{P}_i)$, where $i = 1, \dots, m$. Then, we can apply an m -mode swap operator,

$$\hat{U} := \prod_{i=1}^m \hat{U}_{\text{SWAP}}^{(i)} = \exp \left[i \frac{\pi}{2} \sum_{i=1}^m (\hat{q}_i \hat{P}_i - \hat{Q}_i \hat{p}_i) \right], \quad (6.34)$$

which acts by mapping

$$\begin{aligned} \hat{q}_i &\mapsto \hat{Q}_i \\ \hat{Q}_i &\mapsto -\hat{q}_i \\ \hat{p}_i &\mapsto \hat{P}_i \\ \hat{P}_i &\mapsto -\hat{p}_i \end{aligned} \quad (6.35)$$

and can be seen as generated by an interaction Hamiltonian of the form

$$\hat{H}_I = \frac{\pi}{2} \delta(t) \sum_{i=1}^m \left(\hat{Q}_i \hat{p}_i - \hat{q}_i \hat{P}_i \right). \quad (6.36)$$

Since this operation swaps the phase space of each detector degree of freedom with that of one of the normal modes of the region, all of the correlations between the chosen normal modes of A and its complement $\bar{A}$ will be transmitted to the detectors.

Recall that the quadratures of the normal modes of A can be rewritten as linear combinations of the quadratures of the local oscillators in A . The unitary in Equation (6.34) can therefore be reinterpreted as a coupling between each detector and a privileged linear combination of the local degrees of freedom of the harmonic lattice, in the form

$$\hat{H}_I = \lambda \chi(t) \left(\hat{q}_d \sum_i f_i \hat{p}_i - \hat{p}_d \sum_i g_i \hat{q}_i \right), \quad (6.37)$$

which is the discrete version of Equation (6.4). In the continuum limit, the linear combination

of the different spatial modes of A will translate into a series of smearing functions that define the continuum limit of each normal mode (see Section 6.5).

6.4 Entanglement area laws from finite set of modes

In the previous section, we found bounds on how much entanglement a set of m detectors can acquire with the complement of the region that they couple to. The question remains of how many detectors are necessary in order to capture enough entanglement to faithfully represent the entanglement structure of the field state. In particular, we are going to analyze whether it is possible to increase the number of detectors so that they capture the entanglement area laws that the lattice theory should display, and determine the smallest number of detectors that one needs for this.

It is well-known that subregions of the lattice theory display area law growth of entanglement with increasing size when all of the region's degrees of freedom are taken into account at each step [66]. However, from our discussion in Section 6.3 we know that (i) the entanglement between a region and its complement can be fully characterized by the mixedness of its normal modes, and (ii) the most mixed normal modes of a region have strong support only around the region's boundary. This suggests that it may be possible to extract a similar pattern of entanglement growth by considering only a subset of the degrees of freedom for a given region.

This is indeed the case: if the boundary of a given region A consists of N sites, then one only needs the first N normal modes of subregion A to display area law behavior, thus encompassing nearly all of the entanglement between A and its complement. Once again, we can verify this explicitly in 1, 2, and 3 dimensions as shown in Fig. 6.2, which uses the same parameters as Fig. 6.1. In 1 spatial dimension, for any subregion size, the 2 most mixed normal modes already encode all of the region's entanglement with the rest of the system. In higher dimensions, the number of modes has to grow at each step, but it does so more slowly than the subsystem's size: whereas the total number of degrees of freedom in a given region scales with the region's volume, the number of degrees of freedom supported at the boundary only scales with the area. Therefore, the fraction of modes needed decreases with $1/L$ as the region's side length grows.

This fact has a direct implication on our use of detectors. Namely, it shows that, in order to capture an area-law behavior for the entanglement between a given region and its complement, the number of detector degrees of freedom that is needed scales with the region's boundary area. Furthermore, the optimal choice for the detector couplings would




Figure 6.2: Entanglement entropy between a subregion and its complement as a function of the ratio between the subregion's side length and the total side length of the lattice. Subfigures a), b), and c) show this in 1, 2, and 3 spatial dimensions, respectively. If we take the number of boundary sites to be N_B , then each colored line corresponds to the entanglement captured by the given fraction of N_B modes. The total entanglement entropy is saturated, and we observe an area-law behavior in each case, when we count the entanglement coming from N_B modes.

be one that can effectively be simplified to a swap of the detectors' state and the state of the most mixed normal modes of the given region.

Notice that it is important to distinguish the entropy of a single mode M within A and the entropy of the most mixed normal mode in subregion A . Indeed, the entropy of a single mode in subregion A could be larger than the entropy of the most mixed normal mode in A . However, this does not mean that M has more entanglement with the complement of A than the most mixed normal mode. Instead what happens is that M is entangled both with other modes in A , as well as possibly modes in the complement of A . A normal mode of A does not have any entanglement with other normal modes in A , and therefore its entropy quantifies entanglement with modes that live entirely in $\bar{A}$.

This is important when working with particle detectors, since the interaction of the detector with the lattice itself can create extra entanglement between modes in the interaction region. For instance, if one does not work with normal modes, it is always possible to choose one single mode supported on a given region such that the von Neumann entropy of this mode equals the full entanglement entropy between the region and its complement. One simple way to do this is to pick one of the modes obtained from a two-mode squeezing between the two most mixed normal modes of subregion A . One of the modes thus generated will always have a higher von Neumann entropy than the most mixed

normal mode of A , and by making the squeezing parameter high enough, such von Neumann entropy can be made equal to the total region's von Neumann entropy. By then swapping this mode with just one single-mode detector, the detector would acquire a von Neumann entropy that would also scale according to an area law. It is clear, however, that the resulting mixedness of the detector in the final state would not correspond to entanglement with the complement of the coupling region. On the contrary, the entanglement in this case would be mostly shared with a mode fully supported in the region the detector couples to, and would be almost entirely generated by the direct interaction of the detector and the lattice.

Similarly, a naive swap between N detector degrees of freedom and the N local oscillators supported on the boundary of the coupling region would not serve our purposes. This too would lead to a von Neumann entropy that scales with the area of the subregion of interest. However, each boundary site is equally entangled with its nearest neighbors to the inside and the outside of the coupling region. Therefore, the resulting entanglement transmitted to the detectors in this case would not be representative of the entanglement between the coupling region and its complement either.

6.5 The continuum limit

The results of the previous sections have been obtained using harmonic lattices as a discrete approximation of a quantum field theory. In this section, we will show how to adapt these results in a way that allows us to recover the continuum limit in which the lattice theory becomes a quantum field theory in a box of finite fixed length. We do so by revisiting the discretization procedure that took us from the continuum to the lattice more carefully. As we shall see, this will reveal how to implement the limit somewhat more rigorously, and help to understand the subtleties behind the discrete approximations made in Section 6.2.

Consider a real scalar field theory constrained to a volume $V = [0, L]^d \in \mathbb{R}^d$ with periodic boundary conditions, and with Hamiltonian

$$\hat{H}_\phi = \frac{1}{2} \int_V d^d \mathbf{x} \left(\hat{\pi}^2 + (\nabla \hat{\phi})^2 + m^2 \hat{\phi}^2 \right). \quad (6.38)$$

The periodic boundary conditions allow us to expand the field in a Fourier series,

$$\hat{\phi}(t, \mathbf{x}) = \frac{1}{L^d} \sum_{\mathbf{m} \in \mathbb{Z}^d} \hat{\varphi}(t, \mathbf{k}_\mathbf{m}) e^{-i\mathbf{k}_\mathbf{m} \cdot \mathbf{x}}, \quad (6.39)$$

where $\mathbf{k}_m = \frac{2\pi}{L}\mathbf{m}$. We can then impose a UV cutoff by setting that $\hat{\varphi}(t, \mathbf{k}_m) = 0$ for any $\mathbf{m}$ such that its largest component in absolute value, $\|\mathbf{k}_m\|_\infty$, is bigger than K , with $K = \pi(2N+1)/L$, for a positive integer $N \geq 1$. This allows us to write

$$\hat{\phi}(t, \mathbf{x}) = \frac{1}{L^d} \sum_{\mathbf{m}} \hat{\varphi}(t, \mathbf{k}_m) e^{-i\mathbf{k}_m \cdot \mathbf{x}}, \quad (6.40)$$

where the sum runs over $\mathbf{m} \in \mathbb{Z}^d \cap [-N, N]^d$. In particular, for $\mathbf{l} \in \mathbb{Z}^d \cap [0, 2N]^d$, we have that

$$\hat{\phi}\left(t, \frac{L}{2N+1}\mathbf{l}\right) = \frac{1}{L^d} \sum_{\mathbf{m}} \hat{\varphi}(t, \mathbf{k}_m) e^{-\frac{2\pi i}{2N+1}\mathbf{m} \cdot \mathbf{l}}, \quad (6.41)$$

which means that

$$\left\{ \hat{\phi}\left(t, \frac{L}{2N+1}\mathbf{l}\right) : \mathbf{l} \in \mathbb{Z}^d \cap [0, 2N]^d \right\} \quad (6.42)$$

and

$$\left\{ \frac{1}{L^d} \hat{\varphi}(t, \mathbf{k}_m) : \mathbf{m} \in \mathbb{Z}^d \cap [-N, N]^d \right\} \quad (6.43)$$

are related by a discrete Fourier transform. Thus,

$$\hat{\varphi}(t, \mathbf{k}_m) = a^d \sum_{\mathbf{l}} \hat{\phi}(t, a\mathbf{l}) e^{\frac{2\pi i}{2N+1}\mathbf{m} \cdot \mathbf{l}}, \quad (6.44)$$

where $a = \pi/K = L/(2N+1)$. By Equation (6.39), the field can be expressed as

$$\hat{\phi}(t, \mathbf{x}) = \frac{1}{(2N+1)^d} \sum_{\mathbf{m}} \sum_{\mathbf{l}} \hat{\phi}(t, a\mathbf{l}) e^{-2\pi i \mathbf{m} \cdot (\mathbf{x} - a\mathbf{l})/L}, \quad (6.45)$$

where the sums are running over $\mathbf{m} \in [-N, N]^d$ and $\mathbf{l} \in [0, 2N]^d$. The sum in $\mathbf{m}$ can be evaluated to yield

$$\hat{\phi}(t, \mathbf{x}) = \sum_{\mathbf{l}} \hat{\phi}(t, a\mathbf{l}) \mathcal{S}_{\mathbf{l}}^{(n)}\left(N; \frac{\mathbf{x}}{a}\right), \quad (6.46)$$

where we have denoted

$$\mathcal{S}_{\mathbf{l}}^{(d)}\left(N; \frac{\mathbf{x}}{a}\right) := \prod_{r=1}^d \mathcal{S}_{l_r}\left(N; \frac{x_r}{a}\right), \quad (6.47)$$

with l_r and x_r denoting the r -th component of $\mathbf{l}$ and $\mathbf{x}$, respectively, and

$$\mathcal{S}_j\left(N; \frac{y}{a}\right) := \frac{\sin\left(\pi\left[\frac{y}{a} - j\right]\right)}{(2N+1)\sin\left(\frac{\pi}{2N+1}\left[\frac{y}{a} - j\right]\right)}. \quad (6.48)$$

The same procedure can be applied to the momentum $\hat{\pi}(t, \mathbf{x})$ and the gradient $\nabla\hat{\phi}(t, \mathbf{x})$ yielding

$$\hat{\pi}(t, \mathbf{x}) = \sum_{\mathbf{l}} \hat{\pi}(t, a\mathbf{l}) \mathcal{S}_{\mathbf{l}}^{(d)}\left(N; \frac{\mathbf{x}}{a}\right), \quad (6.49)$$

and

$$\nabla\hat{\phi}(t, \mathbf{x}) = \sum_{\mathbf{l}} \nabla\hat{\phi}(t, a\mathbf{l}) \mathcal{S}_{\mathbf{l}}^{(d)}\left(N; \frac{\mathbf{x}}{a}\right). \quad (6.50)$$

The conjunction of the IR cutoff imposed by the box constraint and the UV cutoff allows us to reduce the description of the field in the continuum to a finite number of degrees of freedom, namely, those of the field at the sampling positions $\{a\mathbf{l}, \mathbf{l} \in [0, 2N]^d\}$. This is akin to a wavelet discretization of the QFT [34]; in fact, similar techniques have been used in the past in related contexts (see for instance [26, 79]).

Once we have reduced the system to a finite number of degrees of freedom, we can use Equations (6.46), (6.49), (6.50) to obtain a discrete version of the Hamiltonian (6.38):

$$\begin{aligned} \hat{H}_\phi = & \frac{1}{2} \sum_{\mathbf{m}, \mathbf{l}} (\hat{\pi}(t, a\mathbf{m}) \hat{\pi}(t, a\mathbf{l}) + \nabla\hat{\phi}(t, a\mathbf{m}) \nabla\hat{\phi}(t, a\mathbf{l}) \\ & + m^2 \hat{\phi}(t, a\mathbf{m}) \hat{\phi}(t, a\mathbf{l})) \mathcal{I}_{\mathbf{ml}}, \end{aligned} \quad (6.51)$$

where

$$\mathcal{I}_{\mathbf{ml}} = \int_V d^d\mathbf{x} \mathcal{S}_{\mathbf{m}}^{(d)}\left(N; \frac{\mathbf{x}}{a}\right) \mathcal{S}_{\mathbf{l}}^{(d)}\left(N; \frac{\mathbf{x}}{a}\right) \quad (6.52)$$

$$= \prod_{r=1}^d \int_0^L dx_r \mathcal{S}_{m_r}\left(N; \frac{x_r}{a}\right) \mathcal{S}_{l_r}\left(N; \frac{x_r}{a}\right) \quad (6.53)$$

$$= a^n \prod_{r=1}^d \delta_{m_r l_r} = a^n \delta_{\mathbf{ml}}^{(d)}. \quad (6.54)$$

Since these integrals cancel out when $\mathbf{m} \neq \mathbf{l}$, the Hamiltonian reduces to

$$\hat{H}_\phi = \frac{a^d}{2} \sum_{\mathbf{m}} (\hat{\pi}(t, a\mathbf{m})^2 + \nabla \hat{\phi}(t, a\mathbf{m})^2 + m^2 \hat{\phi}(t, a\mathbf{m})^2). \quad (6.55)$$

It is important to note that in Equation (6.55) we still have derivatives of the field, so even though we reduced the description of the system to a finite number of degrees of freedom associated to the sampling points, we still rely on the field being defined in a continuum. To complete the discretization and reduce the QFT to a lattice of harmonic oscillators, we need to discretize the derivative. Let us consider the partial derivative in the direction of the s -th coordinate axis:

$$\partial_s \hat{\phi}(t, \mathbf{x}) = \sum_{\mathbf{m}} \partial_s \left[\hat{\phi}(t, a\mathbf{m}) \mathcal{S}_{\mathbf{m}}^{(d)} \left(N; \frac{\mathbf{x}}{a} \right) \right] \quad (6.56)$$

$$= \sum_{\mathbf{m}} \hat{\phi}(t, a\mathbf{m}) \partial_s \mathcal{S}_{\mathbf{m}}^{(d)} \left(N; \frac{\mathbf{x}}{a} \right). \quad (6.57)$$

This allows us to write

$$\sum_{\mathbf{m}} \partial_s \left[\hat{\phi}(t, a\mathbf{m}) \mathcal{I}_{\mathbf{m}\mathbf{l}} \right] = \sum_{\mathbf{m}} \hat{\phi}(t, a\mathbf{m}) \mathcal{J}_{\mathbf{m}\mathbf{l}}^s, \quad (6.58)$$

where

$$\mathcal{J}_{\mathbf{m}\mathbf{l}}^s = \int_V d^d \mathbf{x} \partial_s \mathcal{S}_{\mathbf{m}}^{(n)} \left(N; \frac{\mathbf{x}}{a} \right) \mathcal{S}_{\mathbf{l}}^{(d)} \left(N; \frac{\mathbf{x}}{a} \right). \quad (6.59)$$

It can be shown that

$$\mathcal{J}_{\mathbf{m}\mathbf{l}}^s = (1 - \delta_{m_s l_s}) \left[\frac{(-1)^{l_s - m_s}}{l_s - m_s} a^{d-1} \prod_{r \neq s} \delta_{m_r l_r} + \mathcal{O}(N^{-2}) \right]. \quad (6.60)$$

Thus, when $N \gg 1$,

$$\begin{aligned} \partial_s \hat{\phi}(t, a\mathbf{l}) &= \frac{1}{a} \sum_{\mathbf{m}} \frac{(-1)^{l_s - m_s}}{l_s - m_s} (1 - \delta_{l_s m_s}) \\ &\quad \times \prod_{r \neq s} \delta_{l_r m_r} \hat{\phi}(t, a\mathbf{m}) + \mathcal{O}(N^{-2}). \end{aligned} \quad (6.61)$$

With this, we obtain an approximation to the gradient $\nabla \hat{\phi}(t, a\mathbf{l})$ in terms of the field

amplitude at the sampling points, finally completing the discretization, i.e., the reduction of the field system to a lattice of harmonic oscillators.

Now a few comments are in order. Note that, as one would expect, the expression for the derivative at one point depends on the field amplitude at different sites, so the spatial derivative in the continuum is responsible for the coupling between distinct oscillators in the lattice. We also see, however, that resulting interaction terms in Equation (6.55) include not just nearest-neighbor couplings; therefore, for a generic UV cutoff, the lattice representation of the local QFT will appear to be a lattice of harmonic oscillators with long-range interactions. However, in the limit of very large N , it is still legal to approximate the derivative to just nearest-neighbor differences (see, for instance, [83]), so that the field Hamiltonian reduces to

$$\begin{aligned} \hat{H}_\phi \approx & \frac{a^d}{2} \sum_{\mathbf{l}} \left(\hat{\pi}(t, a\mathbf{l})^2 + m^2 \hat{\phi}(t, a\mathbf{l})^2 \right. \\ & \left. + \sum_{r=1}^d \left[\frac{\hat{\phi}(t, a(\mathbf{l} + \mathbf{1}_r)) - \hat{\phi}(t, a\mathbf{l})}{a} \right]^2 \right), \end{aligned} \quad (6.62)$$

where $\mathbf{1}_r$ is the unit vector with all components equal to zero except for the r -th component. Finally, defining

$$\omega^2 = m^2 + \frac{2d}{a^2}, \quad \alpha = \frac{1}{\omega a^2}, \quad (6.63)$$

and the dimensionless quadratures

$$\hat{q}_{\mathbf{m}}(t) = \sqrt{\omega a^d} \hat{\phi}(t, a\mathbf{m}), \quad \hat{p}_{\mathbf{m}}(t) = \sqrt{\frac{a^d}{\omega}} \hat{\pi}(t, a\mathbf{m}), \quad (6.64)$$

we recover Equation (6.11). Notice that by setting $a = \varepsilon$ in Equations (6.63) and (6.64) we recover Equations (6.13)–(6.16) from Section 6.2.

The analysis we have presented here can be regarded as a more careful examination of the discretization procedure laid out in Section 6.2, which also hints at how to revert this process and go from a discrete lattice to the continuum. If we consider a n -dimensional hypercubic lattice of $(2N+1)^n$ sites with dimensionless quadratures $(q_{\mathbf{l}}, p_{\mathbf{l}})$ and Hamiltonian given by Equation (6.11), given a fixed length L , we can define

$$\hat{\phi}_N(t, \mathbf{x}) = \sum_{\mathbf{l}} \frac{1}{\sqrt{\omega a^d}} \hat{q}_{\mathbf{l}} \mathcal{S}_{\mathbf{l}}^{(d)} \left(N; \frac{\mathbf{x}}{a} \right), \quad (6.65)$$




Figure 6.3: Entanglement entropy between one third of a 1D chain and its complement as a function of the total number of sites in the lattice, i.e. the inverse UV cutoff, as here we keep the length of the chain fixed. As the density of lattice sites increases, we see that the dependence of the entanglement entropy on the effective short-distance cutoff quickly converges to the behavior that is expected from other methods in field theory [133].

$$\hat{\pi}_N(t, \mathbf{x}) = \sum_l \sqrt{\frac{\omega}{a^d}} \hat{p}_l \mathcal{S}_l^{(d)}\left(N; \frac{\mathbf{x}}{a}\right), \quad (6.66)$$

for $\mathbf{x} \in [0, L]^d$, and

$$\hat{H}_{\phi_N} = \frac{1}{2} \int_V d^d \mathbf{x} \left(\hat{\pi}_N^2 + (\nabla \hat{\phi}_N)^2 + m^2 \hat{\phi}_N^2 \right). \quad (6.67)$$

where $a = L/(2N + 1)$, and ω is defined as in Equation (6.63). In the limit $N \rightarrow \infty$,

$$\hat{\phi}_N(t, \mathbf{x}) \rightarrow \hat{\phi}(t, \mathbf{x}), \quad \hat{\pi}_N(t, \mathbf{x}) \rightarrow \hat{\pi}(t, \mathbf{x}), \quad (6.68)$$

and

$$\hat{H}_{\phi_N}(t, \mathbf{x}) \rightarrow \hat{H}_{\phi}(t, \mathbf{x}). \quad (6.69)$$

For illustration purposes, we note that the entanglement entropy of subregions of our lattice quickly approximates the expectations for a continuum field theory as N increases (see for example Fig. 6.3).

Regarding the interaction between the field and the detector, in the continuum, it is described by the Hamiltonian given in Equation (6.4). When we consider the discretized




Figure 6.4: Dimensionless approximated position and momentum profiles, $\lambda_p g_N(x)$ and $\lambda_q f_N(x)$, of the most mixed mode of a region comprising the central third of a box, for a 1+1-dimensional massive quantum scalar field theory of mass $m = 1/L$ in the vacuum state. The profiles were calculated using Equations (6.73) and (6.74), using the spatial profiles obtained for the corresponding 1-dimensional lattice field theories with UV parameters $N = 500, 1000$ and 2000 , and setting the parameters ω and α via Equation (6.63) with $d = 1$. The second row shows an enlarged view of the profiles close to the left boundary.

expressions for the field (6.46) and the momentum (6.49), Equation (6.4) reduces to

$$\begin{aligned}\hat{H}_I = & \lambda \chi(t) \left(\lambda_q \hat{q}_d \sum_l \hat{\pi}(t, a\mathbf{l}) \int_V d^d \mathbf{x} f(\mathbf{x}) \mathcal{S}_l^{(d)} \left(N; \frac{\mathbf{x}}{a} \right) \right. \\ & \left. - \lambda_p \hat{p}_d \sum_l \hat{\phi}(t, a\mathbf{l}) \int_V d^d \mathbf{x} g(\mathbf{x}) \mathcal{S}_l^{(d)} \left(N; \frac{\mathbf{x}}{a} \right) \right).\end{aligned}\quad (6.70)$$

Thus, to recover Equation (6.37), we simply define

$$f_l = \lambda_q \sqrt{\frac{\omega}{a^d}} \int_V d^d \mathbf{x} f(\mathbf{x}) \mathcal{S}_l^{(d)} \left(N; \frac{\mathbf{x}}{a} \right) \quad (6.71)$$

and

$$g_l = \frac{\lambda_p}{\sqrt{\omega a^d}} \int_V d^d \mathbf{x} g(\mathbf{x}) \mathcal{S}_l^{(d)} \left(N; \frac{\mathbf{x}}{a} \right) \quad (6.72)$$

Conversely, to recover the spatial smearing functions $\{f(\mathbf{x}), g(\mathbf{x})\}$ in the continuum from the spatial profiles on the lattice $\{f_l, g_l\}$, we can define

$$f_N(\mathbf{x}) = \frac{1}{\lambda_q \sqrt{\omega a^d}} \sum_l f_l \mathcal{S}_l^{(d)} \left(N; \frac{\mathbf{x}}{a} \right), \quad (6.73)$$

and

$$g_N(\mathbf{x}) = \frac{1}{\lambda_p} \sqrt{\frac{\omega}{a^d}} \sum_l g_l \mathcal{S}_l^{(d)} \left(N; \frac{\mathbf{x}}{a} \right). \quad (6.74)$$

By Equation (6.54), one can see from Equations (6.71) and (6.72) that

$$f_N(\mathbf{x}) \rightarrow f(\mathbf{x}) \quad \text{and} \quad g_N(\mathbf{x}) \rightarrow g(\mathbf{x}) \quad (6.75)$$

when $N \rightarrow \infty$, and thus the modes $\hat{Q}$ and $\hat{P}$ of Equation (6.33) converge to those of Equations (6.6) and (6.7). Note that all the convergences so far stated are achieved pointwise in general, rather than uniformly.

We can use these results to analyze what happens to the position and momentum profiles of the normal modes of a region and its complement when one approaches the continuum limit. In particular, here we will examine the most mixed normal mode.

In Fig. 6.4 we see that, in 1D, as we increase the UV cutoff—and therefore the number of sites of the underlying harmonic chain—the momentum profile $f(x)$ of the most mixed mode of a region formed by a third of the total number of sites seems to converge to a




Figure 6.5: Dimensionless approximated position and momentum profiles, $\lambda_p g_N(x)$ and $\lambda_q f_N(x)$, of the partner mode of the most mixed mode of a region comprising the central third of a box, for a 1+1-dimensional massive quantum scalar field theory of mass $m = 1/L$ in the vacuum state. The profiles were calculated using Equations (6.73) and (6.74), using the spatial profiles obtained for the corresponding 1-dimensional lattice field theories with UV parameters $N = 500, 1000$ and 2000 , and setting the parameters ω and α via Equation (6.63) with $d = 1$. Below the profiles, a zoom-in of the details of the profiles close to the left boundary.

bounded function which is peaked at the boundary and increasingly suppressed within the interior of the region. Meanwhile, the position profile $g(x)$ gets increasingly sharper at the boundary and is almost zero in the interior, hinting at a convergence towards the even sum of two delta functions supported in the two boundary points. Fig. 6.5 shows the same pattern for its partner mode in the complementary region of the chain. We also see that as we approach the continuum limit, the momentum profile localization for the most mixed mode in region A reaches a local minimum in the center of the region. For its partner mode, the momentum profile reaches a minimum in the center of the complement region. We cannot verify numerically whether there is a lower bound for this minimum, but a more detailed study of the behavior of this profile in the continuum limit should be worth exploring in the future.

The behavior displayed in 1D seems to extend to 2D. Following the parameters of the 1D case, we considered a square region with a side length of a third of that of the lattice. In Fig. 6.6, we observe that all the approximated momentum profiles are bounded, peak at the boundary and decrease towards the interior. Meanwhile, the position profiles peak at the boundary with increasingly high maxima, and are all approximately zero in the interior, hinting again at a convergence to a delta at the boundary. A similar behavior is again displayed by the partner mode in the complementary region (see Fig. 6.7).

Notice that the convergence of sequences of continuous functions to delta distributions is yet another evidence of the fact pointed out before that these convergences are pointwise and not uniform in general. Another detail worth mentioning is that the approximated profiles in position and momentum show regions in which they are negative. These only show up outside of the support that the profiles are expected to have in the continuum limit. This is an instance of the well-known Gibbs phenomenon (see, e.g., [106]), and it is a natural consequence of approximating with Fourier sums functions that have jump discontinuities, as the exact profiles are expected to have in the boundary. A simple way to overcome this issue when using the approximate profiles in practice is to set them to zero by fiat out of their known support.

6.6 Summary

In this chapter, we have tackled the problem of optimizing the coupling between probes and quantum field, in order for the final state of the probes to best reflect the entanglement structure between complementary regions of the quantum field of interest. Thanks to the fact that the ground state of the field of interest (in this case, a free scalar field) is pure and Gaussian, its entanglement structure can be very efficiently organized in terms of




Figure 6.6: Dimensionless approximated position and momentum profiles, $\lambda_p g_N(x)$ and $\lambda_q f_N(x)$, of the most mixed mode of a region comprising the central square of side length $L/3$ of a square of side length L , for a 2+1-dimensional massive quantum scalar field theory of mass $m = 1/L$ in the vacuum state. The profiles were calculated using Equations (6.73) and (6.74), using the spatial profiles obtained for the corresponding 2-dimensional lattice field theories with UV parameters $N = 15, 30$ and 60 , and setting the parameters ω and α via Equation (6.63) with $d = 2$.




Figure 6.7: Dimensionless approximated position and momentum profiles of the partner mode of the mode represented in Fig. 6.6.

entanglement between modes of the field, and the question then reduces to how one can best select which modes to couple to in each region.

We started our analysis with an approximation to the free scalar field as a lattice of coupled harmonic oscillators with finitely many degrees of freedom in any finite region of space, which can be obtained by placing a UV cutoff in the theory. However, we also showed how to carefully extrapolate the form of the optimal coupling from the lattice theory to the continuum, which allows one to reconstruct the spatial profile of any given mode of the field as the cutoff is sent to zero. This yields the position and momentum smearing functions one would like to implement in order to couple a particle detector to a field theory fully in the continuum. From a particle-detector perspective, it is worth pointing out that intuitively, particle detectors—which generally come equipped with an inherent energy gap and a minimum finite spatial extension—provide a natural mechanism for an effective UV cutoff. We only expect detectors to be physically able to probe a finite number of modes within a given spatial region, and therefore, they can be used to define an operationally motivated cutoff to the field theory. This suggests that, from the point of view of detectors, even a continuous field theory with infinitely many degrees of freedom may behave as an effectively discretized chain of oscillators, up to the resolution allowed by the detectors [85, 82, 83].

Because the joint state of a region A and its complement $\bar{A}$ is pure, we can reduce the problem of finding the set of modes that is most entangled with the complementary region $\bar{A}$ to the problem of computing the local modes of A that are *mixed*, and contain no correlations with other modes in A ; in short, the problem reduces to a computation of the most mixed *normal* modes of A . This is a very natural candidate for optimal choice of modes for extracting entanglement between a region and its complement, and we gave a proof that this was indeed the case.

As a proof of principle, we also showed how this procedure allows us to reproduce one of the fundamental properties of the entanglement in the ground state of a harmonic lattice—namely, the area law growth of entanglement entropy—with only a subset of the modes in the region. Perhaps unsurprisingly, the minimum number of required modes for a theory with nearest-neighbor interactions grows proportional to the area of the boundary that separates the region of interest from its complement; as the region grows, this number of modes corresponds to a monotonically decreasing fraction of the total number of degrees of freedom in the region. In terms of detectors, this sets the minimum required number of degrees of freedom that must be accessible to probes in order for them to recover the full entanglement that was originally present in a given region. It also sets the limits on how much of the entanglement can be extracted when the probes are limited to a fixed number of degrees of freedom, which is perhaps the most realistic scenario encountered in typical

setups in relativistic quantum information.

Part III

Complexity geometry of Gaussian states

Chapter 7

Introduction

The notion of complexity quantifies the “hardness” of performing a certain task, given some set amount of relevant resources available. The concept has its origins in computer science and information theory, where it refers to the minimum cost (as measured by resources such as memory usage or number of elementary operations) of a given computation [60]. This notion is readily exportable to quantum computing [52, 191], where a natural measure of complexity is provided by the so-called *circuit complexity*. Given some fixed set of “simple” unitary transformations (i.e., an elementary gate set), the circuit complexity of a unitary $\hat{U}$ is the minimum number of elementary gates needed to build $\hat{U}$, and the circuit complexity of a state $|\psi\rangle$ is the minimum number of elementary gates needed to prepare $|\psi\rangle$ starting from some fiducial reference state $|\psi_0\rangle$.¹

Recent years have seen a significant surge in interest in the concept of complexity, thanks to the many insights that it can provide to several areas of physics extending far beyond quantum computing and quantum information. In many-body physics, complexity plays a central role in understanding chaos and thermalization in closed quantum systems [127, 152, 107], and it also serves as a powerful figure of merit to characterize quantum phases of matter [128, 7, 115]. In high energy physics, complexity has also been suggested as an important element in establishing the regimes of validity of quantum field theory in curved spacetimes, something that has led to potential insights into the black hole information

¹If we want to be more precise in this context, we should say that the circuit complexity is the minimum number of elementary gates needed to *approximate* $\hat{U}$, subject to some error tolerance ϵ under some distance measure in the space of unitaries (most commonly the 1-norm). The reason for this is that any finite circuit built from a finite elementary gate set can only generate a finite number of unitaries, so the circuit complexity of any unitary would be infinite except for a set of measure zero in the unitary group under this definition.

loss problem and the nature of the black hole interior in quantum gravity [94, 171, 170, 172, 109, 6]. Finally, the concept of complexity is also of great importance in the context of the AdS/CFT correspondence, where a variety of different geometric quantities of the bulk spacetime have been put forward as potential candidates for the holographic dual of the circuit complexity of the state on the boundary CFT [167, 29, 28, 53, 13, 14].

However, quantitatively characterizing complexity is a very challenging task, since finding the optimal circuit that implements a given unitary or prepares a given state is highly nontrivial in general. To tackle this problem, geometric methods of describing complexity have proven to be very powerful. This has been made most concrete by Nielsen’s geometric approach [129, 132, 131], which proposes to reformulate circuit complexity fully in terms of differential geometry. The idea is that by equipping the space of unitary operators with some notion of distance, the complexity of a given unitary can be pictured as the length of a geodesic connecting the identity operator to the unitary of interest. Nielsen’s approach has recently been shown to be capable of providing lower bounds to the complexity of typical unitaries which improve upon previously known bounds [27]. It has also been very fruitful in establishing concrete notions of complexity in quantum field theory both for free/weakly interacting fields [105, 92, 17, 49] and for CFTs [46, 47], which guides much of the recent work in holographic complexity in AdS/CFT (for a recent review, see [51]). Another proposal for the geometrization of complexity is based directly on the geometry of the Hilbert space of quantum states provided by the Fubini-Study metric [50, 75]. Finally, in a similar spirit, a recent proof of an important conjecture about the growth of complexity in typical random circuits (which was originally motivated by considerations on black holes [173, 30]) was made possible thanks to tools from differential topology and algebraic geometry [93].

Despite all this progress, the inherent ambiguities in key elements of the definition of circuit complexity (namely, in the choices of reference state and elementary gate set, as well as the cost associated to each elementary gate) are sometimes seen as obstacles to a more complete understanding of complexity in QFT and the various related geometrical quantities that are conjectured to be dual to it in AdS/CFT [199]. The usual philosophy when dealing with these issues in the context of quantum information theory and quantum computing is to formulate questions about complexity that are less sensitive to these ambiguities. This is the basis, for example, of the definition of quantum complexity classes, which are only concerned with how the complexity of a family of circuits, built with a common set of elementary gates, scales with the size of the input [191]. In this thesis, however, we will adopt a complementary point of view according to which these seemingly negative aspects of the definition of circuit complexity are actually features of the formalism that can encode relevant physics. After all, if complexity quantifies the difficulty in performing a certain

task, it makes sense for it to be determined dynamically, based on what physical operations objectively cost in terms of available lab resources.

With this motivation in mind, in this thesis we will show how to formulate general Riemannian metrics on the space of pure Gaussian states (both for bosonic and fermionic systems) which are sufficiently broad to accommodate all physically relevant settings, while still having enough structure to allow us to make concrete statements characterizing state complexity. We will also propose an extension to the notion of complexity geometry which allows for a “non-reversible” cost—that is, we provide a concrete example of a modified complexity measure that can assign different costs to directions in circuit space related to each other by time reversal. The goal is to yield an operationally motivated measure of complexity that is directly connected to an experimentalist’s view of what tasks would be considered easy or hard in practice; this paper provides a step in that direction. The choice of Gaussian states is motivated by the plethora of powerful analytical techniques and results that the assumption of Gaussianity provides, as well as by the pervasiveness of Gaussian states in a variety of different scenarios of physical interest ranging from quantum field theory in curved spacetimes [10, 190] to quantum optics and quantum information with continuous variables more broadly [25, 192, 2]. We should note, however, that it is possible to extend this general strategy to other families of circuits whose infinitesimal generators form a Lie algebra.

The remaining chapters of Part III are organized as follows:

- Chapter 8 reviews the main technical tools used in this part of the thesis. We begin in Section 8.1 by presenting the basics of Nielsen’s geometric approach to circuit complexity [129, 132, 131]. As a preparation for the approach that will be taken later on, we also argue at the end of the Section why adaptations of Nielsen’s complexity geometry which relax some of the usual assumptions made about the complexity metric are reasonable and physically motivated. Section 8.2 then reviews the unified description of bosonic and fermionic Gaussian states that we will make use of, closely following the formalism of [92, 49, 90]. While not containing any substantially new material, we believe that the content of Sections 8.1 and 8.2 is useful for fixing notation and terminology, and also helps to contextualize the results to be presented later in the paper with what is already known in the literature.
- Chapter 9 contains the main original contributions of Part III. Section 9.1 describes the general conceptualization of a Riemannian metric that can yield a notion of complexity geometry for pure Gaussian states. We highlight a minimal, well-motivated assumption that the metric will be required to satisfy, and show that, provided that

the assumption holds, the state complexity defined by any physically reasonable metric on the space of Gaussian states can be characterized in a unified way. This extends known results and methods in the literature to a much more general set of metrics, with potentially arbitrary cost functions assigned to a given elementary gate set. In Section 9.2 we propose a new extension to the complexity geometry which introduces a feature of “non-reversibility” to the complexity measure by assigning different costs to a given gate and its inverse. This is another arguably very physical feature of an experimentalist’s experience of what is easy or hard to achieve in a lab, which is however usually absent in most adopted measures of complexity geometry. The particular new ingredient proposed has a very intuitive physical interpretation as a kind of vector potential, which modifies the equation of motion for the optimal paths in a way that directly reproduces the Lorentz force law experienced by a charged particle in a background electromagnetic field. Finally, Section 9.3 then illustrates the general strategy with a few simple examples where state complexity can actually be computed in closed form.

Chapter 8

Review of complexity geometry and Gaussian states

8.1 Nielsen’s geometric approach to circuit complexity

Let $\mathcal{U}$ be the set of all unitary operators that act on some Hilbert space $\mathcal{H}$ describing the physical system of interest. The first step in Nielsen’s geometric approach to circuit complexity is to replace a discrete quantum circuit with a continuous path $\hat{U}(t)$ in $\mathcal{U}$, which can in general be written as

$$\hat{U}(t) = \mathcal{T} \exp \left(-i \int_{t_0}^t \hat{H}(t') dt' \right). \quad (8.1)$$

The (in general time-dependent) operator $\hat{H}(t)$ (let us call it ‘Hamiltonian’ for short) can essentially be seen as equivalent to the tangent vector to the trajectory $\hat{U}(t)$, as Equation (8.1) is a solution to the differential equation

$$\frac{d\hat{U}}{dt} = -i\hat{H}(t)\hat{U}(t), \quad (8.2)$$

which defines the relation between the tangent vector $d\hat{U}/dt$ and the Hamiltonian generator $\hat{H}(t)$.

Nielsen’s geometric approach then replaces the notion of a “number of gates” applied at each layer of the circuit by a *cost function* $c(\hat{U}, \hat{V})$ which assigns to each infinitesimal time

step dt taken at a point $\hat{U}$ with tangent vector $d\hat{U}/dt = \hat{V}$ an infinitesimal cost

$$dC = c(\hat{U}(t), \hat{V}(t)) dt, \quad (8.3)$$

where the Hamiltonian $\hat{H}(t)$ that generates this step can be given in terms of the tangent vector $\hat{V}$ by inverting Equation (8.2),

$$\hat{H}(t) = i\hat{V}(t)\hat{U}^\dagger(t). \quad (8.4)$$

The finite version of the cost of a given circuit is thus given by

$$C = \int_{t_0}^{t_f} c(\hat{U}(t), \hat{V}(t)) dt \quad (8.5)$$

and the complexity of a unitary $\hat{U}_T$ is defined as

$$\mathcal{C}(\hat{U}_T) = \min \left(\int_{t_0}^{t_f} c(\hat{U}(t), \hat{V}(t)) dt \right) \quad (8.6)$$

where the minimization ranges over all the trajectories $\hat{U}(t)$ in $\mathcal{U}$ satisfying the boundary conditions

$$\hat{U}(t_0) = \hat{\mathbb{1}}, \quad \hat{U}(t_f) = \hat{U}_T. \quad (8.7)$$

Similarly, one defines the state complexity of a target state $|\psi_T\rangle$ relative to a reference state $|\psi_R\rangle$ as

$$\mathcal{C}(|\psi_R\rangle, |\psi_T\rangle) = \min_{\hat{U}_T} \{ \mathcal{C}(\hat{U}_T) \mid \hat{U}_T |\psi_R\rangle = |\psi_T\rangle \}. \quad (8.8)$$

The most basic requirement that we would expect from $c(\hat{U}, \hat{V})$ is that it should be a smooth and positive-definite function—that is, for any unitary $\hat{U}$ and any tangent vector $\hat{V}$ in the tangent space to $\hat{U}$, it holds that $c(\hat{U}, \hat{V}) \geq 0$. It is usual to demand that the inequality be saturated if and only if $\hat{V} = 0$. After all, it is natural to expect the cost of any given infinitesimal step to be nonnegative, and only be zero if no gates are applied at all. It also makes sense to require that $c(\hat{U}, \hat{V})$ satisfy the triangle inequality in the second argument, $c(\hat{U}, \hat{V}_1 + \hat{V}_2) \leq c(\hat{U}, \hat{V}_1) + c(\hat{U}, \hat{V}_2)$. Intuitively, this just tells us that we cannot decrease the cost of a given step by decomposing the tangent vector at that point into a given linear combination, and then evolving with each component of the linear combination separately. Finally, if we want the depth given by Equation (8.5) to be invariant under general reparametrizations $t \rightarrow f(t)$ that preserve the time orientation of the path (that is, $df/dt > 0$), we would also require that c be positive-homogeneous of degree 1 in the

second argument,¹ i.e., $c(\hat{U}, \lambda \hat{V}) = \lambda c(\hat{U}, \hat{V})$ for every $\lambda \geq 0$. Under the assumption that Equation (8.5) is reparametrization-invariant, we can always choose the parameter t to be such that the initial point of the curve $\hat{U}(t)$ is at $t_0 = 0$ and the final point is at $t_f = 1$. That is the convention that we will adopt for the remainder of the thesis.

The conditions we just provided give $c(\hat{U}, \hat{V})$ the properties of a *Finsler metric*, which then equips the unitary group with the structure of a *Finsler manifold* [155]. Finsler manifolds are a generalization of Riemannian manifolds where one can still talk about lengths between points, but the length functional is not necessarily derived from a metric tensor. The interpretation of the cost function as a Finsler metric allows us to see expression (8.5) as the length of the curve $\hat{U}(t)$ between times t_0 and t_f , and the optimal circuit that prepares a given unitary $\hat{U}_T$ is nothing more than the trajectory of minimum length—i.e., the geodesic—connecting $\hat{\mathbb{1}}$ and $\hat{U}_T$. The problem of computing circuit complexity then becomes a problem in variational calculus, where the full power of differential geometry can be put into action.

In order to put this (somewhat abstract) notion of a Finsler metric on the unitary group in closer contact with the original idea of circuit complexity, it is customary to expand the operator $\hat{H}(t)$ generating the circuit at each instant of time in a fixed basis of Hermitian generators $\hat{\mathcal{O}}_I$,

$$\hat{H}(t) = \sum_I Y^I(t) \hat{\mathcal{O}}_I, \quad (8.9)$$

and express the cost function explicitly in terms of the components Y^I in this basis. The interpretation of the generators $\hat{\mathcal{O}}_I$ is that the unitaries $\hat{U}_I = e^{i\varepsilon \hat{\mathcal{O}}_I}$ form our elementary gate set, with ε being some minimum time scale for each gate to be applied.

If our complexity measure only cares about the number of gates applied at each time step, the cost function evaluated at a given time t along the path in $\mathcal{U}$ will only depend on the Hamiltonian $\hat{H}(t)$ generating the circuit at that time. In more technical terms, this means that $c(\hat{U}, \hat{V})$ defines a *right-invariant* Finsler metric, which satisfies

$$c(\hat{U}, \hat{V}) = c(\hat{\mathbb{1}}, \hat{V} \hat{U}^\dagger). \quad (8.10)$$

If we parametrize the tangent vector $\hat{V}$ in terms of the Hamiltonian $\hat{H}(t)$ through Equation (8.4) and we expand $\hat{H}(t)$ in a given basis of generators as in (8.9), right-invariance of the Finsler

¹We only demand positive-homogeneity (instead of requiring, for instance, that $c(\hat{U}, \lambda \hat{V}) = |\lambda| c(\hat{U}, \hat{V})$ for every $\lambda \in \mathbb{R}$) because we would also like to accommodate cases where the cost is not invariant under time reversal. Although most of the examples studied in practice do assign the same cost to a given elementary gate and its inverse, in Sec. 9.2 we give an example where that is not the case.

metric as in Equation (8.10) implies that the cost function only depends on the components Y^I , and not on the point $\hat{U}$; in short, $c(\hat{U}, \hat{V}) \equiv F(Y)$ for some function F . Simple examples of cost functions $F(Y)$ commonly considered in the literature include [129]

$$F_1(Y) = \sum_I |Y^I|, \quad (8.11a)$$

$$F_{1p}(Y) = \sum_I p_I |Y^I|, \quad (8.11b)$$

$$F_2(Y) = \sqrt{\sum_I (Y^I)^2}, \quad (8.11c)$$

$$F_{2q}(Y) = \sqrt{\sum_I q_I (Y^I)^2}. \quad (8.11d)$$

The cost function (8.11a) (unoriginally named the F_1 -metric) is the one that most closely matches the original idea of counting the number of gates in the circuit. The F_{1p} -metric (8.11b) is a modification of (8.11a) that includes possibly nontrivial *penalty factors* p_I that may vary between different gates. This can be used to enforce further physical constraints in the complexity measure, for instance, by making nonlocal gates more costly than local ones. Strictly speaking, the F_1 -metric is not Finsler, since it violates the smoothness requirement; as such, it is not immediately amenable to the methods from variational calculus and differential geometry. As already explained in e.g. [129], this can be remedied by approximating F_1 by a one-parameter family of Finsler metrics that recover the F_1 metric in some limit. Another (more common) approach is to instead adopt the F_2 -metric (8.11c), which has the natural interpretation of a *Riemannian* metric that emerges from an Euclidean inner product on the tangent space at each point of $\mathcal{U}$. This is the strategy most explored in applications of Nielsen's framework to complexity in quantum field theory [105, 92, 49]. Naturally, the F_{2q} -metric (8.11d) is to F_2 the same as F_{1p} is to F_1 .

In this thesis, however, we are interested in a generalization of geometric measures of complexity where the cost function depends nontrivially on the point in circuit space where it is being evaluated. In this case, the Finsler metric is no longer right-invariant, and therefore the measure of complexity at hand is more detailed than a simple counting of the number of gates in a circuit. Physically, this is very well-motivated: after all, it is possible to consider regimes where, from the point of view of an experimentalist, the ability to reliably apply a given infinitesimal transformation will depend on the state of the system. One can imagine, for instance, that increasing the entanglement between two parties by a

given amount may become harder if the state you are acting on is already highly entangled, as compared to when the two parties start in a product state.

Additionally, previous geometric complexity measure proposals do not address the issue of non-reversibility of complexity. For example, acting on the ground state of a system to prepare some complex entanglement structure is much more experimentally involved than starting from some very entangled state and reaching the ground state, since for the latter it just becomes a matter of waiting for the system to decay to its ground state when coupled to some cold enough environment. In terms of the actual lab resources (e.g., money, graduate students, etc.) spent in the experiment, the second scenario is arguably cheaper than the first, and it would be relevant to have a measure of complexity that mimics this feature.

As we will see in the next chapter, several important features of the geometry of the optimal circuit can be stated in a way that does not actually depend on right-invariance, so it is possible to extend known results and approaches for right-invariant metrics to more general Riemannian metrics on the space of unitaries. We will also see how to modify the cost functional in order to incorporate physical features of irreversibility in the complexity measure in a very natural geometrical way. This will establish the basis on which to build a formalism for the systematic study of complexity measures that are, in some sense, closer to an experimentalist’s view of what is easy or hard to perform.

8.2 Parametrization of pure Gaussian states

Our analysis in this thesis will focus on a general description of complexity geometry for a special class of states known as *Gaussian states*. These can be generally interpreted as squeezed, coherent, thermal states of systems with quadratic Hamiltonians. Thanks to the fact that the Hamiltonian governing the dynamics of every system close to equilibrium can be approximated as quadratic, Gaussian states are ubiquitous in the study of a wide range of systems in a variety of areas of physics. Applications range from foundational aspects of quantum field theory in curved spacetimes [10, 190] to experimental implementations of quantum information protocols with quantum optics [81, 111]. The mathematical formalism for dealing with Gaussian states for both bosonic and fermionic systems is also very well-developed, and forms the basis for much of the results in quantum information and many-body physics with continuous variables [25, 192, 201].

We will closely follow a strategy employed in [92, 49] which parametrizes pure Gaussian states as equivalence classes of an appropriately chosen matrix group describing the space

of unitaries of interest. As we will see shortly, this approach is very powerful and natural for considerations about Gaussian state complexity.

Consider a quantum system with N degrees of freedom (or modes), described by $2N$ quadrature operators $\hat{\xi}^a$. These consist of a set of Hermitian operators,

$$\hat{\xi}^a = (\hat{Q}^1, \hat{P}_1, \dots, \hat{Q}^N, \hat{P}_N)^\top \quad (8.12)$$

which, in the case of bosons, satisfy the canonical commutation relations

$$\begin{aligned} [\hat{Q}^i, \hat{P}_j] &= i\delta_j^i \mathbf{1}, \\ [\hat{Q}^i, \hat{Q}^j] &= [\hat{P}_i, \hat{P}_j] = 0, \end{aligned} \quad (8.13)$$

or, in the case of fermions, the *anticommutation* relations

$$\begin{aligned} \{\hat{Q}^i, \hat{P}_j\} &= 0, \\ \{\hat{Q}^i, \hat{Q}^j\} &= \{\hat{P}_i, \hat{P}_j\} = \delta_j^i \mathbf{1}. \end{aligned} \quad (8.14)$$

Referring to the phase-space operators $\hat{\xi}^a$ as “quadratures” is standard in the literature for bosonic systems; for fermionic systems, one might find it more usual to call this basis of Hermitian generators “Majorana modes”. However, since most of what we will present will apply both to fermionic and bosonic systems in a unified fashion, we will use “quadratures” to refer to both cases at once, in line with the conventions used, e.g., in [196].

A Gaussian state $\hat{\rho}$ is fully characterized by the one-point and two-point correlators of the quadratures,

$$z^a := \langle \hat{\xi}^a \rangle_{\hat{\rho}}, \quad (8.15)$$

$$W^{ab} := \langle \hat{\xi}^a \hat{\xi}^b \rangle_{\hat{\rho}} = \frac{1}{2} (\sigma^{ab} + i\Omega^{ab}) + z^a z^b, \quad (8.16)$$

where we have defined

$$\sigma^{ab} = \langle \hat{\xi}^a \hat{\xi}^b + \hat{\xi}^b \hat{\xi}^a \rangle_{\hat{\rho}} - 2z^a z^b, \quad (8.17)$$

$$\Omega^{ab} = -i\langle \hat{\xi}^a \hat{\xi}^b - \hat{\xi}^b \hat{\xi}^a \rangle_{\hat{\rho}}. \quad (8.18)$$

Higher-order n -point correlators are then directly obtained from these one- and two-point correlators (8.16) via Wick contractions. Note that for fermions, the one-point function

always vanishes.² In this case, the formalism simplifies, as we can restrict ourselves to the two-point function only.

Considered as tensors on phase space, the real part σ^{ab} of the two-point correlator defines a metric (in the sense that it is a symmetric, positive-definite bilinear map), and the imaginary antisymmetric part Ω^{ab} defines a symplectic form (in the sense that it is an antisymmetric and non-degenerate bilinear map). Due to the commutation (anticommutation) relations of the quadratures in bosonic (fermionic) systems, only the real (imaginary) part of the two-point function actually carries nontrivial information about the correlations in the Gaussian state, with the remaining part being fully state-independent and fixed by the algebra of the quadratures. It is common to refer to the state-dependent part of the correlations in a Gaussian state (σ^{ab} for bosons, and Ω^{ab} for fermions) as the *covariance matrix*. We can then combine the state-dependent and state-independent parts of the correlations in a Gaussian state into a single phase-space operator which we call the *complex structure*,

$$J^a_b := \begin{cases} -\sigma^{ac}\Omega_{cb} & \text{for bosons} \\ \Omega^{ac}\sigma_{cb} & \text{for fermions,} \end{cases} \quad (8.19)$$

where we have defined σ_{ab} and Ω_{ab} as the matrix inverses of σ^{ab} and Ω^{ab} respectively,³

$$\sigma_{ab} := (\sigma^{-1})_{ab}, \quad (8.20)$$

$$\Omega_{ab} := (\Omega^{-1})_{ab}. \quad (8.21)$$

For pure Gaussian states—which will be our main focus here—, it turns out that the two (in principle different) definitions for J^a_b in (8.19) actually coincide: namely, for both fermions and bosons, $\sigma^{ac}\Omega_{cb} = -\Omega^{ac}\sigma_{cb}$ if the state is pure. In this case, the operator defined in (8.19) satisfies $J^2 = -\mathbb{1}$ (or $J^a_c J^c_b = -\delta^a_b$ in index notation), which justifies calling it a complex structure. For mixed states the equality no longer holds, but it is somewhat customary to keep the name regardless.

Unitary dynamics preserving Gaussianity (in short, unitary Gaussian transformations) are generated by Hamiltonians that are at most quadratic in the phase-space variables

²We should note that there is such a thing as coherent states for fermionic systems, where one can have a nonzero displacement z^a as a Grassman-valued vector. While these can be useful calculational tools, there are no actual physical fermionic Gaussian states [90]. Therefore, we will only assume nonzero one-point functions in the description of bosonic Gaussian transformations, as it is common in phase space quantum mechanics.

³For mixed fermionic Gaussian states, it might happen that Ω^{ab} is not invertible. In this case, Ω_{ab} is defined as the pseudoinverse of Ω^{ab} with respect to σ^{ab} —i.e., we invert Ω^{ab} only on the subspace that is orthogonal (in the sense of the inner product defined by σ^{ab}) to its kernel [196, 91].

$\hat{\xi}^a$, and act as affine operators on phase space. One can thus view Gaussian dynamics as implemented by pairs $(\mathbf{v}, M)$, where $v \in \mathbb{R}^{2N}$ is a phase-space displacement vector, and $M \in \text{GL}(2N, \mathbb{R})$ is some invertible matrix. A Gaussian transformation labeled by $(\mathbf{v}, M)$ will act on the quadratures (in the Heisenberg picture) as

$$\hat{\xi}^a \mapsto M^a_b \hat{\xi}^b + v^a \hat{\mathbf{1}}. \quad (8.22)$$

The vector of means therefore evolves as

$$z^a \mapsto M^a_b z^b + v^a, \quad (8.23)$$

and the matrices σ^{ab} and Ω^{ab} evolve as

$$\sigma^{ab} \mapsto M^a_c \sigma^{cd} M^b_d, \quad (8.24)$$

$$\Omega^{ab} \mapsto M^a_c \Omega^{cd} M^b_d, \quad (8.25)$$

or, in matrix notation,

$$\mathbf{z} \mapsto M\mathbf{z} + \mathbf{v}, \quad (8.26)$$

$$\sigma \mapsto M\sigma M^\top, \quad (8.27)$$

$$\Omega \mapsto M\Omega M^\top \quad (8.28)$$

where $M^\top$ is the transpose of M , with matrix elements $(M^\top)_b^a = M^a_b$. Finally, the complex structure J transforms as

$$J \mapsto MJM^{-1}. \quad (8.29)$$

Since the canonical commutation/anticommutation relations between the quadratures are state-independent, the matrices M that implement unitary Gaussian transformations must preserve either the symplectic form (in the case of bosons) or the metric given by σ (in the case of fermions). Gaussian transformations for systems of bosons are therefore represented on phase space by pairs $(\mathbf{v}, M)$ where v is a general vector in $\mathbb{R}^{2N}$ and M is an element of the *symplectic group*,

$$\text{Sp}(2N, \mathbb{R}) = \{M \in \text{GL}(2N, \mathbb{R}) \mid M\Omega M^\top = \Omega\}. \quad (8.30)$$

The set of pairs $(\mathbf{v}, M) \in \mathbb{R}^{2N} \rtimes \text{Sp}(2N, \mathbb{R})$ forms a new Lie group, with group multiplication being given by

$$(\mathbf{v}_1, M_1) \cdot (\mathbf{v}_2, M_2) = (M_1\mathbf{v}_2 + \mathbf{v}_1, M_1M_2). \quad (8.31)$$

For fermions, there are no displacements, since the one-point function of the quadratures is taken to always vanish. Therefore, fermionic Gaussian transformations are elements of the *orthogonal group* with respect to the inner product defined by the symmetric, positive-definite matrix σ ,

$$\mathrm{O}(2N) = \{M \in \mathrm{GL}(2N, \mathbb{R}) \mid M\sigma M^\top = \sigma\}. \quad (8.32)$$

The set of all unitary Gaussian transformations acting in a system of $2N$ quadratures thus forms a finite-dimensional group G , both for bosons and for fermions. In the case of bosons, the group is given by $G = \mathbb{R}^{2N} \rtimes \mathrm{Sp}(2N, \mathbb{R})$. For fermions, since we are mostly interested in the subset of orthogonal transformations that can be continuously connected to the identity, we will take $G = \mathrm{SO}(2N)$ instead of the full orthogonal group. If our elementary generators $\{\hat{\mathcal{O}}_I\}$ from (8.9) only contain operators that are at most quadratic in $\hat{\xi}^a$, then the space of available unitary transformations acting on the Hilbert space of either bosons or fermions can be identified, in a one-to-one fashion up to a global phase, with the group G , and the at-most quadratic generators $\hat{\mathcal{O}}_I$ can be mapped to elements of the Lie algebra $\mathfrak{g}$ of G . This renders the description of trajectories on the space of Gaussian transformations effectively finite-dimensional, since G is a finite-dimensional group. The simplifications that arise from this fact represent the main reason behind the power of the analytical treatments of Gaussian states—especially in the case of bosons, whose unrestricted Hilbert space is infinite-dimensional.

Now that we have mapped Gaussian unitary transformations to elements of a finite-dimensional group G , let us briefly turn our attention to states. Bosonic Gaussian states can be identified with pairs $(J, \mathbf{z})$, where J is a complex structure and $\mathbf{z}$ is the vector of means of the quadratures. In the case of fermions, the vector of means is always taken to be zero, so the complex structure J alone already encodes all the information about the fermionic Gaussian state. In either case, for any two pure Gaussian states—labeled by $(J_1, \mathbf{z}_1)$ and $(J_2, \mathbf{z}_2)$ in the case of bosons, or simply J_1 and J_2 for fermions—it is always possible to find a pair $(\mathbf{v}, M) \in \mathbb{R}^{2N} \rtimes \mathrm{Sp}(2N, \mathbb{R})$ (for bosons) or an orthogonal matrix $M \in \mathrm{O}(2N)$ (for fermions) such that

$$(J_2, \mathbf{z}_2) = (MJ_1M^{-1}, M\mathbf{z}_1 + \mathbf{v}) \quad (\text{bosons}), \quad (8.33)$$

$$J_2 = MJ_1M^{-1} \quad (\text{fermions}). \quad (8.34)$$

For the case of fermions, if we restrict G to only include elements of $\mathrm{SO}(2N)$ instead of $\mathrm{O}(2N)$, then a Gaussian transformation connecting two fermionic Gaussian states exists as long as both states belong to the same parity sector—i.e., both have the same fermion

parity number [92]. However, in all cases, fixing the initial and final states does not uniquely fix the Gaussian transformation mapping one to the other. In particular, if a given matrix M (in the symplectic group for bosons, or the orthogonal group for fermions) is such that $MJ_1M^{-1} = J_2$, then so is MS , whenever the matrix S satisfies

$$SJ_1S^{-1} = J_1. \quad (8.35)$$

Therefore, if the initial and final states have vanishing vector of means, the transformation mapping between the two is only well-defined modulo the right action of matrices M which commute with the complex structure of the initial state.

For a given complex structure J , the set of all S satisfying (8.35) defines a subgroup of the set of Gaussian transformations which we call the *stabilizer* of J , denoted by $\text{Sta}_J(N)$. More precisely, we define

$$\text{Sta}_J(N) := \begin{cases} \{S \in \text{SO}(2N), [S, J] = 0\}, \\ \{(\mathbf{0}, S) \mid S \in \text{Sp}(2N, \mathbb{R}), [S, J] = 0\}, \end{cases} \quad (8.36)$$

where the first line above applies to fermions and the second line applies to bosons.

In words, $\text{Sta}_J(N)$ simply corresponds to the subset of Gaussian transformations represented by linear operators on phase space that leave the covariance matrix of the system invariant. For a Gaussian state described by the complex structure J and vanishing vector of means, $\text{Sta}_J(N)$ simply consists of the Gaussian transformations that do not change the physical state. Since the covariance matrix defines a metric in the case of bosons and a symplectic form in the case of fermions, we conclude that, for any choice of J used as reference, the subgroup $\text{Sta}_J(N)$ is isomorphic to the intersection $\text{Sp}(2N, \mathbb{R}) \cap \text{SO}(2N)$, which in turn is isomorphic to the unitary group $U(N)$.

Note that for the case of bosonic Gaussian states with $\mathbf{z} \neq 0$, the action of $\text{Sta}_J(N)$ *will* change the physical state, as it will in general change the vector of means by $\mathbf{z} \rightarrow S\mathbf{z}$. We could have defined the stabilizer as including a nonzero displacement that undoes this shift in $\mathbf{z}$; however, we believe that the choice of stabilizer not containing any displacements is more physically motivated. The intuitive idea is that we want the action of the stabilizer to represent some kind of “free evolution” that can act on the reference state at no cost, as long as the structure of correlations present in the reference state is preserved. If we picture this “free evolution” as being generated by some Hamiltonian, we consider that having only quadratic terms in the Hamiltonian (i.e., harmonic couplings between the quadratures) is more natural than including both quadratic and linear terms (which we would normally picture as coming from some additional driving force on the quadratures). The subgroup

generated only by a family of purely quadratic Hamiltonians, *with no linear terms*, that preserves the structure of correlations in the initial state, is precisely the definition of $\text{Sta}_J(N)$ as written in (8.36). It is worth noting, however, that the formal aspects of the framework that we are about to describe would remain largely unaltered if we had chosen to define the stabilizer as the subset of bosonic Gaussian transformations given by pairs of the form $((\mathbb{1} - S)\mathbf{z}, S)$, where $[S, J] = 0$.

We thus have the following parametrization of the space of pure Gaussian states. First, without loss of generality and for simplicity, we take one zero-mean state $|J_R\rangle$ as a reference (so that the elements of the stabilizer leave it invariant). Then, we note that all pure Gaussian states can be obtained from $|J_R\rangle$ through the action of an element of G , but two elements $M_1, M_2 \in G$ characterize the same state if they differ by the right-action of some element of $\text{Sta}_{J_R}(N)$. Therefore, pure Gaussian states can be described as *equivalence classes* of Gaussian transformations, where right-action by an element of $\text{Sta}_{J_R}(N)$ defines the equivalence relation. This means that, once we pick a particular state $|J_R\rangle$ as a reference, we can endow the group G of Gaussian transformations with the structure of a fiber bundle, where the fibers are orbits of the stabilizer of the reference state, and the base manifold—which is isomorphic to the space of pure Gaussian states—is the quotient $G/\text{Sta}_{J_R}(N)$. This is the setting that we will use to parametrize the space of Gaussian states, and eventually characterize state complexity, in the following section.

Chapter 9

Generalized Gaussian complexity

9.1 General features of Gaussian state complexity geometry

So far, none of what we have said depends on a metric on the group G . The stabilizer subgroup, and therefore the characterization of Gaussian states as equivalence classes of Gaussian transformations, are functions of only the reference state. The next step in order to define a geometric notion of complexity is to equip G with notions of lengths and angles so that we can turn the question of circuit complexity into a problem in Riemannian geometry.

Let g be a Riemannian metric on G . This is a tensor field which assigns, at each point $p \in G$, a bilinear map

$$\begin{aligned} g_p : \mathcal{T}_p G \times \mathcal{T}_p G &\rightarrow \mathbb{R} \\ (V_p, W_p) &\mapsto g_p(V_p, W_p) \end{aligned} \tag{9.1}$$

that is symmetric and positive-definite,

$$\begin{aligned} g_p(V_p, W_p) &= g_p(W_p, V_p) \quad \forall V_p, W_p \in \mathcal{T}_p G, \\ g_p(V_p, V_p) &> 0 \quad \forall V_p \neq 0. \end{aligned} \tag{9.2}$$

The norm of a vector with respect to this metric is then just given by the usual expression, $\|V_p\| = \sqrt{g_p(V_p, V_p)}$.¹

¹We emphasize that this Riemannian metric g is completely different from the metric on phase space that

Having this additional structure on G , we assign the length of a trajectory $\gamma: \mathbb{R} \rightarrow G$ between two endpoints $\gamma(0)$ and $\gamma(1)$, with its associated tangent vector $d\gamma/dt := \dot{\gamma}(t)$, as

$$\ell_\gamma = \int_0^1 dt \sqrt{g_{\gamma(t)}(\dot{\gamma}, \dot{\gamma})} = \int_0^1 dt \|\dot{\gamma}(t)\|. \quad (9.3)$$

The complexity of some element $M \in G$ is then simply defined as the geodesic length between the identity and M ,

$$\mathcal{C}(M) = \min_{\gamma} \{\ell_\gamma \mid \gamma(0) = \mathbb{1}, \gamma(1) = M\}. \quad (9.4)$$

This associates a value of ‘complexity’ to a given Gaussian transformation.

However, this is not enough to talk about the complexity of a given Gaussian state. To define state complexity, a second minimization is necessary, since there is more than one Gaussian transformation that connects the reference and target states. Recall that for bosons, general Gaussian transformations take the form $M = (\mathbf{v}, S)$ where $\mathbf{v} \in \mathbb{R}^{2N}$ and S is a symplectic matrix. Now, given a bosonic Gaussian state $|J_R, \mathbf{0}\rangle \equiv |J_R\rangle$ as a reference, and taking another bosonic Gaussian state $|J_T, \mathbf{z}_T\rangle$ as a target (possibly with nonzero displacement), the state complexity of $|J_T, \mathbf{z}_T\rangle$ relative to $|J_R\rangle$ is defined as

$$\mathcal{C}(|J_R\rangle, |J_T, \mathbf{z}_T\rangle) = \min_S \{\mathcal{C}(\mathbf{z}_T, S) \mid J_T = S J_R S^{-1}\}, \quad (9.5)$$

where S ranges over the symplectic group. Note that the displacement part of the Gaussian transformation is univocally determined by the vector of means of the target state, but the symplectic transformation is only fixed up to the action of an element of the stabilizer of J_R —which is why the minimization over S above is required. For fermions, we do not have to consider displacements, and the state complexity of a given target state $|J_T\rangle$ relative to the reference state $|J_R\rangle$ is just given by

$$\mathcal{C}(|J_R\rangle, |J_T\rangle) = \min_O \{\mathcal{C}(O) \mid J_T = O J_R O^{-1}\}, \quad (9.6)$$

where O is an element of the special orthogonal group, and we assume that both reference and target in this case belong to the same parity sector [92]. Equations (9.5) and (9.6) are just the minimal geodesic length between the identity and the submanifold given by the equivalence class of Gaussian transformations that prepare the target state taking the reference state as the starting point. In all cases, the differential equations imposing the

could be defined, for instance, via the symmetric, positive-definite matrix σ associated with the reference state.

extremization of the length between reference and target (9.4)-(9.6) in any given coordinate chart γ^i for the group manifold G is simply the good old geodesic equations,

$$\frac{d^2\gamma^i}{ds^2} + \underbrace{\frac{1}{2}g^{il}(\partial_j g_{kl} + \partial_k g_{lj} - \partial_l g_{jk})}_{\equiv \Gamma^i_{jk}} \frac{d\gamma^j}{ds} \frac{d\gamma^k}{ds} = 0, \quad (9.7)$$

where s is an affine parameter along the curve $\gamma(s)$ (that is, a parameter such that the norm of the tangent vector to the trajectory is kept constant), and Γ^i_{jk} are the standard Christoffel symbols.

Given a metric tensor on the group manifold G , we can define the *Riemannian exponential map* from tangent vectors at the identity (which we identify with elements of the Lie algebra $\mathfrak{g}$ of G) to points on the manifold, which we will denote by $\exp$. More specifically, the Riemannian exponential is a function that assigns to every tangent vector $V \in T_{\mathbb{1}}G \simeq \mathfrak{g}$ a group element $\exp(V) = \gamma_V^{(R)}(1)$, where $\gamma_V^{(R)}(t)$ is the geodesic that starts from $\gamma_V^{(R)}(0) = \mathbb{1}$ with tangent vector $\dot{\gamma}_V^{(R)}(t)|_{t=0} = V$. Equivalently, $\exp(V)$ corresponds to the endpoint of a geodesic that starts at $\mathbb{1}$ with tangent vector V and has total length $\|V\|$.

The fact that G is a Lie group also equips it with another notion of exponentiation, namely the *Lie exponential map*. For a given element V of the Lie algebra, we define the Lie exponential of V as the point $e^V = \gamma_V^{(L)}(1)$, where $\gamma_V^{(L)}(t)$ is the flow line of the right-invariant vector field on G whose value at the identity is $\dot{\gamma}_V^{(L)}(t)|_{t=0} = V$.

Note that these exponentials (the Riemannian exponential here denoted by $\exp(V)$, and the Lie exponential denoted by e^V) are two distinct maps. The Lie exponential is a purely group-theoretical concept that does not depend on a metric at all, whereas the Riemannian exponential is only defined once the manifold is equipped with a metric. For a completely general Riemannian metric, the Riemannian and Lie exponentials will not coincide; in fact, they are only guaranteed to be equal if the metric is bi-invariant (that is, both left- and right-invariant) under the group action. In the case of fermions, this can be achieved by taking the metric to be the negative of the Killing form on the group manifold $\text{SO}(2N)$. For bosons, on the other hand, this is not an option, since the Killing form on $\text{Sp}(2N, \mathbb{R})$ does not have a definite sign, so it cannot be used to define a Riemannian metric.

For our purposes, however, we actually do *not* want both exponentials to match in general. Even for the case of fermions (where a bi-invariant metric can be defined), we would still like to accommodate for a complexity measure that is not right-invariant, due to the considerations made at the end of Section 8.1. Instead, all that we will demand is that the two maps coincide only for some properly chosen subgroup of the Lie algebra $\mathfrak{g}$, which

we now describe.

Consider the tangent space to the identity of G , which is identified with the Lie algebra $\mathfrak{g}$. Pick a reference state characterized by a complex structure J_R , which, as explained in Section 8.2, selects a stabilizer $\text{Sta}_{J_R}(N) \subset G$. From now on, to avoid too much clutter in the notation, we will omit the subscript J_R from the stabilizer, and just keep in mind that all statements that depend in some way on the stabilizer will implicitly depend on the choice of reference state. The stabilizer subgroup $\text{Sta}(N)$, being a Lie group of its own, also possesses a Lie algebra, which we will denote by $\mathfrak{sta}(N) \subset \mathfrak{g}$. For the purposes of defining a geometric measure of complexity, it is very natural to at least expect that the image of both exponential maps is the same when acting on $\mathfrak{sta}(N)$. After all, if this were not the case, there would be elements of the stabilizer whose lowest-cost preparation included Gaussian transformations that are not part of the stabilizer—which does not make a lot of sense, since we would like to have circuits that act nontrivially on the reference state to have strictly larger costs than those which do not.

We are now in a position to establish the one condition that we will assume the metric on G should satisfy:

- The metric must be such that, for every $V \in \mathfrak{sta}(N)$, the trajectory $M(t) = e^{tV}$ is a geodesic, and the parallel transport along that geodesic of an arbitrary vector from $\mathbb{1}$ to $M(t)$ is given by right-translation by $M(t)$.

This condition is more restrictive than just requiring that both the Riemannian and Lie exponentials span the same submanifold when acting on $\mathfrak{sta}(N)$. For one, we are prescribing how parallel transport acts also on elements of $\mathfrak{g}$, that do not belong to $\mathfrak{sta}(N)$. Moreover, we are assuming that for every V in the Lie algebra of the stabilizer, $\exp(V) = e^V$ —in other words, the images of both maps acting on $\mathfrak{sta}(N)$ are the same element-by-element. We nevertheless consider these constraints to be very natural: the stabilizer group ultimately corresponds to the “trivial” directions on the space of Gaussian operations from the point of view of state complexity. Therefore, imposing these constraints on the metric for the stabilizer is not a strong restriction, and it will still allow us to define a geometric measure of complexity in full power. Namely, we note that we still have enough degrees of freedom to meaningfully choose a complexity metric for a given general setup, since we have not specified the parallel transport along geodesics that are not contained in the stabilizer. This is precisely where we have the freedom to apply physical constraints on the available resources in order to build the complexity geometry.

Equipped with the metric at the identity, we can split $\mathfrak{g}$ as a direct sum $\mathfrak{sta}(N) \oplus \mathfrak{sta}_\perp(N)$,

where $\mathfrak{sta}_\perp(N)$ is the orthogonal complement of $\mathfrak{sta}(N)$,

$$\mathfrak{sta}_\perp(N) = \{K \in \mathfrak{g}, g_1(K, V) = 0 \ \forall V \in \mathfrak{sta}(N)\}. \quad (9.8)$$

We can then locally foliate G near the identity element in terms of the cylinder-like surfaces defined by

$$C_r = \{\exp(V)S \mid V \in \mathfrak{sta}_\perp(N), \|V\| = r, S \in \text{Sta}(N)\}. \quad (9.9)$$

Recall that due to the assumption previously made, $\exp(tV)S$ is precisely the geodesic starting from S at $t = 0$ with tangent vector equal to the parallel transport of V from the identity to S —which is orthogonal to $\text{Sta}(N)$ and has the same norm as V because parallel transport preserves inner products. Therefore, in plain English, C_r is a codimension-1 surface generated by taking the submanifold $\text{Sta}(N)$, and then shooting off geodesics of length r in the orthogonal directions to $\text{Sta}(N)$.

This setup allows us to talk about the state complexity of a given state relative to a reference state in terms of the radial coordinate r defined in (9.9) as follows. Note that the map $(V, S) \mapsto \exp(V)S$ provides a local coordinate chart from some local neighborhood of the zero vector² in $\mathfrak{sta}_\perp(N)$ (times the stabilizer group) to some local neighborhood of the identity element on the group manifold G . By the way the coordinate chart is constructed, two points in the same equivalence class by the action of the stabilizer have the same V coordinates; in particular, they belong to the same cylinder C_r . But the distance of every point of the cylinder to the identity is lower-bounded by the cylinder’s radius, and there always exist trajectories whose length is precisely equal to that radius—namely, those that emanate from the identity in the orthogonal direction to $\mathfrak{sta}(N)$ and whose tangent vector has magnitude equal to r . We have therefore shown that the radial coordinate *is* the state complexity. A geometric visualization of the strategy is displayed in Fig. 9.1.

This construction draws on and extends the strategy introduced in [92, 49] to a much wider range of Riemannian metrics defining the complexity geometry of Gaussian states. In particular, it highlights that the key feature that made the proof of their complexity bound possible—namely, the fact that every Gaussian transformation on the same equivalence class belonged to the same cylinder C_r —did not actually depend on the assumption of right-invariance or the definition of the inner product at the identity. This observation is ultimately what allowed us to leverage that construction and generalize the characterization of state complexity to more general (and potentially more physically motivated) complexity metrics.

²Whether or not the Riemannian exponential map defines a coordinate chart in a given neighborhood of the identity of G depends on whether the Riemann curvature causes caustics or not. We will make the reasonable assumption that caustics are not present within the range of target states we wish to consider.




Figure 9.1: Pictorial representation of the geometrical ingredients involved in defining state complexity from a metric in the group manifold. After fixing the reference state (and thus the stabilizer subgroup $\text{Sta}(N)$ in G), one shoots off geodesics in the orthogonal directions to the generators of the stabilizer at the identity, thus constructing the submanifold $\exp(\mathfrak{sta}_\perp(N))$. The state complexity of the target state $|\psi_T\rangle$ given by the equivalence class of some Gaussian transformation $M \in G$ is then visualized as the length of the geodesic connecting 1 to the point where the equivalence class $[M]$ intersects $\exp(\mathfrak{sta}_\perp(N))$.

9.2 A non-reversible measure of complexity

A common feature in the examples of geometric measures of complexity found in the literature as well as those mentioned so far in this chapter is that they are all “time-reversal-symmetric”—i.e., the complexity is invariant if the roles of reference and target states are reversed³. This is a simple consequence of the fact that all cost functions presented thus far assign the same cost to the direction V and to the time-reversed direction $-V$, for every tangent vector V at any given point in the group manifold G .

From an experimentalist’s point of view, however, it is natural to imagine contexts in which the “difficulty” in connecting two given states is *not* symmetric under swapping the roles of reference and target states. This may happen due to some effectively irreversible interaction with an external environment that drives the system to some preferred state: if there is some decay mechanism by which the system tends to spontaneously flow to the ground state of its free Hamiltonian, for instance, it is reasonable to conceive of a complexity measure which would naturally assign a lower cost to the circuit that starts at an excited state and ends at the aforementioned ground state than to the circuit that goes the other way around. It is therefore interesting to ask what new structures could be added to our geometrical measure of complexity, in order to encode some form of time reversal symmetry breaking.

9.2.1 Formalism for pure states

Although the loss of time-reversal symmetry is often tied to interactions with an uncontrollable environment that makes the dynamics irreversible, experimental asymmetries in the generation of different states may appear even in cases where we model state preparation with unitary dynamics. It should therefore be possible to incorporate this feature in models that retain unitarity from the perspective of the system of interest. A simple complexity measure that readily achieves this is given by the following “generalized length” functional

³Strictly speaking this is not exactly true in the example of coherent states that will be considered in Subsection 9.3.2, since we assumed by convention that the reference state would always be chosen to have vanishing vector of means. In that case, reverting the roles of reference and target states must also be accompanied by a trivial overall displacement that centers the new reference state back at the origin, which results in the modification from $(|J_R\rangle, |J_T, \mathbf{z}_T\rangle)$ to $(|J_T\rangle, |J_R, -\mathbf{z}_T\rangle)$ as reference and target. It is clear that the state complexity (9.31) is invariant under this replacement as well.

to the curve $\gamma(t)$ on the group manifold:

$$C_\gamma = \int_0^1 dt \|\dot{\gamma}(t)\| - \int_0^1 dt A_i(\gamma(t)) \dot{\gamma}^i. \quad (9.10)$$

In addition to the metric g , this cost functional also depends on a background 1-form A_i . Since the metric is assumed to be non-degenerate, this 1-form can equivalently be seen as arising from a background vector field $A^i \equiv g^{ij} A_j$, so that (9.10) becomes

$$C_\gamma = \int_0^1 dt \|\dot{\gamma}(t)\| - \int_0^1 dt g_{\gamma(t)}(A, \dot{\gamma}). \quad (9.11)$$

In order for the cost function implied by Equation (9.11) to be positive-definite, it suffices for the norm of the background vector field A^i according to the metric g to be smaller than or equal to 1. This is because the Cauchy-Schwarz inequality guarantees that

$$|g(A, \dot{\gamma})| \leq \|A\| \|\dot{\gamma}\|, \quad (9.12)$$

and therefore demanding $\|A\| \leq 1$ everywhere in the group manifold is enough to ensure that

$$\|\dot{\gamma}\| - g(A, \dot{\gamma}) \geq 0, \quad (9.13)$$

which in turn guarantees that the cost function in Equation (9.10) is always nonnegative. It is also clear that this modified cost function is smooth, satisfies the triangle inequality, and is positive-homogeneous in $\dot{\gamma}$, but it is *not* invariant under time reversal since the costs of going in the direction $\dot{\gamma}$ or its reverse $-\dot{\gamma}$ now differ by the inner product $-2A_i \dot{\gamma}^i$. Note that in the limiting case where $\|A\| = 1$ at a given point p in the group manifold, there is one nontrivial direction at that p with zero cost—namely, the trajectory with tangent vector parallel to A^i at p .

One can immediately recognize the complexity functional (9.10) as mathematically identical to the action of a charged particle moving in a background electromagnetic field, with the 1-form A_i being analogous to the vector potential. The equation of motion for the trajectory on circuit space that extremizes (9.10) thus takes the same form as the Lorentz force law in a generally curved background metric,

$$\frac{d^2 \gamma^i}{ds^2} + \Gamma^i_{jk} \frac{d\gamma^j}{ds} \frac{d\gamma^k}{ds} = F^i_j \frac{d\gamma^j}{ds}, \quad (9.14)$$

where Γ^i_{jk} is given by the same expression as in (9.7), and we have $F^i_j \equiv g^{ik} F_{kj}$, with F_{ij}

the usual standard electromagnetic field strength,

$$F_{ij} = \partial_i A_j - \partial_j A_i. \quad (9.15)$$

9.2.2 Generalizations to mixed states

The formalism from Section 9.2.1 can be readily adapted to also include mixed states. The usual strategy to define complexity for mixed states is to consider the so-called *purification complexity* [4]. Given a Hilbert space $\mathcal{H}$ describing the system of interest, for any two density matrices $\hat{\rho}_R$ and $\hat{\rho}_T$ corresponding to the reference and target states, the purification complexity of $\hat{\rho}_T$ relative to $\hat{\rho}_R$ is given by

$$\mathcal{C}(\hat{\rho}_R, \hat{\rho}_T) = \min_{|\psi_R\rangle, |\psi_T\rangle} \left\{ \mathcal{C}(|\psi_R\rangle, |\psi_T\rangle) \right\} \quad (9.16)$$

where the minimization ranges over states $|\psi_R\rangle, |\psi_T\rangle$ in an enlarged Hilbert space $\mathcal{H} \otimes \mathcal{H}'$ such that

$$\hat{\rho}_R = \text{Tr} \left(|\psi_R\rangle \langle \psi_R| \right)_{\mathcal{H}'}, \quad (9.17)$$

$$\hat{\rho}_T = \text{Tr} \left(|\psi_T\rangle \langle \psi_T| \right)_{\mathcal{H}'}. \quad (9.18)$$

If both the reference and target states are Gaussian, it is common to further restrict the minimization above to only consider purifications that are also Gaussian [38, 42]. In this case, one might call the complexity measure given by Equation (9.16) the *Gaussian complexity of purification*. It is clear that, once we have a generalized cost functional that displays some feature of non-reversibility for pure states—given, for instance, by the additional background structure described in Section 9.2.1—, the same feature will be generally transported to the case of mixed states according to the definition above. A similar logic also applies to alternative approaches to the geometrization of complexity based on directly defining a notion of distance on the space of density matrices via the Bures metric [153], which generalizes the approach based on the Fubini-Study metric [50] to the case of mixed states.⁴

⁴For a related approach based on the Fisher-Rao distance between Gaussian states, see also [62].

9.3 Examples

We will consider here four examples. For illustration purposes, the first one reviews an example already analyzed in the literature [92, 49]. The second example extends the first, by including the case of bosonic coherent states. This has also been analyzed in the past through slightly different techniques [88], but here we will be able to provide an easily computable closed expression for the geometric Gaussian complexity of general coherent states based on minimal assumptions. Finally, the last two examples consider simple extensions of the geometric complexity formalism that may capture features of the difficulty of physical state preparation in a lab that right-invariant metrics cannot.

9.3.1 Right-invariant metric on the group manifold

The first example we will present serves the purpose of illustrating how the general machinery laid out in previous sections is enough to provide closed-form expressions for circuit complexity. This example will closely follow the work in [92, 49], which was the main inspiration for the generalization we presented in Section 9.1.

In [92, 49], the authors considered pure Gaussian states with vanishing vector of means, and therefore the group G for bosons is reduced to just $\text{Sp}(2N, \mathbb{R})$. They then apply this generalized cylindrical foliation to evaluate the state complexity derived from a Riemannian metric which, at the Lie algebra $\mathfrak{g} \simeq T_1 G$ (in this case equal to $\mathfrak{so}(2N)$ for fermions, and $\mathfrak{sp}(2N, \mathbb{R})$ for bosons), is given by

$$g_1(V, W) = \frac{1}{2} \text{Tr} (V \sigma_R W^\top \sigma_R^{-1}), \quad (9.19)$$

with σ_R being the symmetric part of the two-point correlator of the quadratures in the chosen reference state. The metric is extended to all points $M \in G$ of the group manifold by demanding right-invariance,

$$g_M(V, W) = g_1(V M^{-1}, W M^{-1}). \quad (9.20)$$

In some sense, this is a natural choice of inner product in the Lie algebra, as it does not require anything beyond the group structure and some minimal information about the reference state. For fermions, σ_R is actually independent of the state, and fully fixed by the anticommutation relations of the quadratures. In this case, Equation (9.19) is nothing but the negative of the Killing form in $\text{SO}(2N)$. For bosons, the metric unavoidably depends

on the choice of reference state, as there is no canonically defined positive-definite form on $\text{Sp}(2N, \mathbb{R})$ that only depends on the group structure of the symplectic group.

It is important to emphasize once again that, in general, one should not expect this measure of complexity to accurately quantify actual physical limitations based on any realistic lab resources. We would expect, for instance, that an experimentally motivated measure of complexity would depend nontrivially on the reference state even in the case of fermions. For bosons, Equation (9.19) includes some dependence on the reference state through its covariance matrix σ_R , but the complexity measure is still not able to account for other physically motivated features such as lack of time-reversal symmetry. However, in the absence of more detailed models for the emergence of a physically motivated complexity measure, it has been common to take (9.19) as a paradigmatic definition of complexity geometry in the literature, which is why we are reviewing this example first before considering any modifications.

The choice of metric defined by Equations (9.19)-(9.20) allows for a very explicit verification of the general strategy laid out in Section 9.1. In particular, it is possible to write a closed-form expression for the shortest path connecting the identity to the equivalence class of Gaussian transformations that prepare a given target state, and thus obtain an analytical solution to the problem of computing Gaussian state complexity. Given the complex structures J_R and J_T defining the reference and target states respectively, one can always define the so-called *relative complex structure*

$$\Delta = J_T J_R^{-1}. \quad (9.21)$$

Then, the geodesic that leads to the target state can be given in the space of Gaussian transformations as

$$M(\tau) = e^{\tau \log(\Delta)/2}. \quad (9.22)$$

One can directly verify that the generator $\log(\Delta)/2 \in \mathfrak{g}$ is in the orthogonal complement of the Lie algebra of the stabilizer of the reference state J_R according to the inner product of Equation (9.19), and that $M(1)J_R M(1)^{-1} = J_T$. With this, the Gaussian state complexity of both fermionic and bosonic Gaussian systems is simply given in terms of the relative complex structure as [90]

$$\mathcal{C}(|J_R\rangle, |J_T\rangle) = \frac{1}{2\sqrt{2}} \sqrt{\text{Tr}(|\log(\Delta)|^2)}. \quad (9.23)$$

9.3.2 Complexity of coherent states

In the case of bosons, we can generalize the example in Subsection 9.3.1 to include coherent states, where the target state contains some nonzero displacement $\mathbf{z}_T$. This can be motivated, for example, by the fact that many important proposals for the implementation of quantum information in continuous-variable systems are based on encoding the logical information in the displacement vector, and not in the covariance matrix [81, 119, 174, 87]. Our next example will thus consist of a simple extension of the formulas of Subsection 9.3.1 to the case where the target state is a coherent state.

In parallel to what we did before, all we need to do now is to define an inner product at the tangent space of the identity, with the minimal modification of now including tangent vectors that generate displacements. This inner product is then transported to the rest of the group manifold by the analog of Equation (9.20), with the group element now being a pair $(\mathbf{v}, M) \in \mathbb{R}^{2N} \rtimes \text{Sp}(2N, \mathbb{R})$. A natural⁵ choice of inner product at the identity can be taken as

$$g_{(\mathbf{0}, \mathbb{1})}((\mathbf{v}, V), (\mathbf{w}, W)) = \sigma_R^{-1}(\mathbf{v}, \mathbf{w}) + \frac{1}{2} \text{Tr}(V \sigma_R W^\top \sigma_R^{-1}) \quad (9.24)$$

where $V, W \in \mathfrak{sp}(2N, \mathbb{R})$, and

$$\sigma_R^{-1}(\mathbf{v}, \mathbf{w}) = (\sigma_R^{-1})_{ab} v^a w^b. \quad (9.25)$$

Then, given a reference state $|J_R, \mathbf{0}\rangle \equiv |J_R\rangle$ and a target state $|J_T, \mathbf{z}_T\rangle$, the optimal circuit $(M(\tau), z(\tau))$ can be parametrized as

$$M(\tau) = e^{\tau \log(\Delta)/2}, \quad (9.26)$$

$$\mathbf{z}(\tau) = (M(\tau) - \mathbb{1})(M(1) - \mathbb{1})^{-1} \mathbf{z}_T, \quad (9.27)$$

where we recall $\Delta = J_T J_R^{-1}$ is the relative complex structure from the reference to the target state. Represented as a unitary operator on the Hilbert space, the trajectory above simply corresponds to the time evolution implemented by

$$\hat{U}(\tau) = e^{-i\hat{H}\tau}, \quad (9.28)$$

⁵Following the same logic as with the example in Subsection 9.3.1, this is the simplest way to define the complexity metric in the displacement sector that does not add any extra structure.

where the time-independent Hamiltonian $\hat{H}$ is defined as

$$\begin{aligned}\hat{H} &= \frac{1}{2} F_{ab} \hat{\xi}^a \hat{\xi}^b + \alpha_a \hat{\xi}^a, \\ F_{ab} &= \frac{1}{2} \Omega_{ac} (\log \Delta)^c{}_b, \\ \alpha_a &= \frac{1}{2} \Omega_{ac} N^c{}_d (\mathbf{z}_T)^d.\end{aligned}\tag{9.29}$$

In (9.29), N is given by

$$N = \log(\Delta)(\sqrt{\Delta} - \mathbb{1})^{-1}\tag{9.30}$$

and we recall that the symplectic matrix Ω_{ab} is the inverse of Ω^{ab} as defined in (8.18)—which, for bosonic systems, is state-independent and fixed by the algebra of the quadratures.

We thus obtain the state complexity

$$\mathcal{C}(|J_R\rangle, |J_T, \mathbf{z}_T\rangle) = \frac{1}{2} \sqrt{\frac{\text{Tr}(|\log(\Delta)|^2)}{2} + G(\mathbf{z}_T, \mathbf{z}_T)},\tag{9.31}$$

where $G(\mathbf{z}, \mathbf{z}) = G_{ab} z^a z^b$, and the bilinear form G is defined by

$$G = N^\top \sigma_R^{-1} N,\tag{9.32}$$

with N defined as in (9.30). The case where target and reference states share the same complex structure and only differ by the displacement $\mathbf{z}_T$ is readily accommodated as the limit when $\Delta = \mathbb{1}$. In this case, we have $N = 2 \cdot \mathbb{1}$, the optimal path on the space of displacements reduces to a straight line,

$$\mathbf{z}(\tau) = \tau \mathbf{z}_T,\tag{9.33}$$

and the state complexity reduces to

$$\mathcal{C}(|J_R\rangle, |J_R, \mathbf{z}_T\rangle) = \sqrt{\sigma_R^{-1}(\mathbf{z}_T, \mathbf{z}_T)}.\tag{9.34}$$

9.3.3 Weyl transformation of right-invariant metric

It would also be interesting to have an example of a complexity metric that does not have as much symmetry as the metrics presented in Subsections 9.3.1 or 9.3.2. After all, it is

reasonable to expect that a measure of complexity that is motivated by constraints from physical limitations will not be represented by a right-invariant metric, but will depend more nontrivially on the state on which a given elementary gate set is acting.

A very simple example of a metric that is no longer right-invariant, but for which most of the existing closed-form expressions of Subsection 9.3.1 will still be greatly useful, is obtained if we perform a Weyl transformation on the previously right-invariant metric. That is, we define a new metric in the group manifold in terms of the metric presented in Section 9.3.1 by

$$\tilde{g}_M(V, W) = e^{2\omega(r)} g_M(V, W). \quad (9.35)$$

In the Weyl transformation above, we are assuming that $\omega(r)$ is a function only of the generalized radial coordinate r —which, in this case, for any Gaussian transformation in the equivalence class labeled by a given complex structure J , is just the state complexity given by Equation (9.23) with $J_T = J$. Heuristically, the metric (9.35) describes a simple model where the cost of acting with a certain gate may change depending on how complex the circuit already is at the time the gate is applied.

Mathematically, the importance of $\omega(r)$ being only a function of r stems from the observation that, in this case, the path given in Equation (9.22) is also a geodesic of the modified metric. To see this, first note that the geodesic equation (9.7) for a curve with tangent vector V can be written in a more compact form as

$$V^i \nabla_i V^j = 0, \quad (9.36)$$

where ∇_i above denotes the Levi-Civita connection associated to g_{ij} —i.e., the covariant derivative that is torsion-free and compatible with the metric g_{ij} . Under a general Weyl transformation $g_{ij} \mapsto \tilde{g}_{ij} = e^{2\omega} g_{ij}$, the Christoffel symbols transform as

$$\Gamma_{jk}^i \mapsto \tilde{\Gamma}_{jk}^i = \Gamma_{jk}^i + \delta_j^i \nabla_k \omega + \delta_k^i \nabla_j \omega - g_{jk} g^{il} \nabla_l \omega. \quad (9.37)$$

Therefore, if the trajectory $M(\tau)$ with tangent vector $V = dM/d\tau$ was a geodesic of the metric g_{ij} , the same trajectory (with the same parametrization), for the case of the Weyl-transformed metric, will now satisfy

$$V^i \tilde{\nabla}_i V^j = 2(V^i \nabla_i \omega) V^j - \|V\|_g^2 g^{ij} \nabla_i \omega, \quad (9.38)$$

where $\tilde{\nabla}_i$ is the Levi-Civita connection associated to the new metric $\tilde{g}_{ij}$, and $\|V\|_g^2$ corresponds to the squared norm of V according to the *old* metric, $\|V\|_g^2 = g_{ij} V^i V^j$.

Equation (9.38) certainly does not look like the geodesic equation for arbitrary ω and

for a general geodesic of the old metric. However, if we take the geodesic (9.22) and we assume that $\omega = \omega(r)$, the problem simplifies greatly due to the fact that, under this assumption, we have $g^{ij}\nabla_j\omega \propto V^i$ along the trajectory (9.22). This can be seen by noting that the tangent vector V satisfies $V \propto \partial_r$, the conformal factor ω is such that $d\omega \propto dr$, and the vector field ∂_r along the trajectory (9.22) is orthogonal to the cylinders of constant r (see [92, 49]), so $g^{ij}\nabla_j\omega \propto (\partial_r)^i$. We can thus simply write

$$\nabla_i\omega = \pm \frac{\|d\omega\|_g}{\|V\|_g} V_i, \quad (9.39)$$

where $\|d\omega\|_g$ corresponds to the norm of the 1-form $d\omega$ according to the metric g , $\|d\omega\|_g^2 = g^{ij}(\nabla_i\omega)(\nabla_j\omega)$. The plus-or-minus choice in (9.39) just accommodates the cases where ω increases or decreases with increasing r , with the plus sign applying to the former and the minus sign applying to the latter. In both cases, substituting this back into (9.38), we obtain

$$\begin{aligned} V^i\tilde{\nabla}_iV^j &= (V^i\nabla_i\omega)V^j \\ &\equiv \Omega(\tau)V^j, \end{aligned} \quad (9.40)$$

which is of the form of a *non*-affinely parametrized geodesic curve, $V^i\tilde{\nabla}_iV^j \propto V^j$, with the proportionality factor $\Omega(\tau)$ being given by $\Omega(\tau) = V^i\nabla_i\omega = d\omega/d\tau$. In the last equation, $\omega(\tau)$ should be understood as the pullback of $\omega(r)$ to the trajectory (9.22). A new affine parameter s (such that the associated tangent vector $\tilde{V} = dM/ds$ satisfies the geodesic equation in the form (9.36) for the new metric $\tilde{g}_{ij}$) is then obtained by noting that, if both $\tilde{V}$ and V are tangent vectors for the same curve, we can write

$$V = \frac{dM}{d\tau} = \frac{ds}{d\tau} \frac{dM}{ds} = \frac{ds}{d\tau} \tilde{V}. \quad (9.41)$$

Plugging this back into Equation (9.40) and demanding that $\tilde{V}^i\tilde{\nabla}_i\tilde{V}^j = 0$ then gives us a differential equation for s ,

$$\frac{1}{f} \frac{df}{d\tau} = \frac{d\omega}{d\tau}, \quad f(\tau) = \frac{ds}{d\tau}. \quad (9.42)$$

We obtain a unique solution to (9.42) by demanding that $s(\tau = 0) = 0$ and $s(\tau = 1) = 1$, so that the new affine parameter also ranges from 0 to 1 between the reference and target states. This gives us

$$s(\tau) = \frac{1}{\int_0^1 d\tau' e^{\omega(\tau')}} \int_0^\tau d\tau' e^{\omega(\tau')}. \quad (9.43)$$

Putting all of this together with the results reviewed in Section 9.3.1, we conclude that the state complexity derived from the metric given by (9.35) is simply

$$\tilde{\mathcal{C}}(|J_R\rangle, |J_T\rangle) = \frac{1}{2\sqrt{2}} \int_0^1 d\tau e^{\omega(\tau)} \sqrt{\text{Tr}(|\log(\Delta)|^2)} \quad (9.44)$$

In fact, once we knew that the original geodesics were purely radial and the conformal factor was only a function of r , there was an elementary argument that would also have led to Equation (9.44). The new line element in the radial direction after the Weyl transformation is $d\ell = e^{\omega(r)} dr$, which once integrated from $r' = 0$ to $r' = r$ gives us the new proper length

$$\ell = \int_0^r dr' e^{\omega(r')} = r \int_0^1 d\tau e^{\omega(\tau)}, \quad (9.45)$$

where we have renormalized the limits of integration by setting $\tau = r'/r$, and used the slight abuse of notation $\omega(r(\tau)) = \omega(\tau)$. By then remembering that the radial coordinate associated to the target state is precisely given by (9.23), the result from Equation (9.45) is exactly what we obtained in Equation (9.44).

9.3.4 Non-reversible complexity for a single bosonic mode

Finally, as a simple example to illustrate a setting that also includes non-reversibility, let us consider a single bosonic degree of freedom, and restrict ourselves to the subset of pure Gaussian states with vanishing vector of means. In this case, the full space of pure Gaussian states under consideration can be characterized by 2×2 covariance matrices σ . Without loss of generality, we can set the covariance matrix of the reference state to just be the identity,

$$\sigma_R = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (9.46)$$

We then recall that the symplectic matrix for bosons is state-independent and given by

$$\Omega = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \quad (9.47)$$

so the reference state can equivalently be characterized by a complex structure which, in the simple choice of canonical basis where σ_R becomes the identity, is just represented as

$$J_R = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \quad (9.48)$$

Thanks to the simplicity of the setup, all of the general ingredients developed in the previous sections can be readily worked out in full detail. The stabilizer subgroup $\text{Sta}(J_R)$ associated to the complex structure (9.48) is simply the orthogonal group in 2 dimensions, whose Lie algebra $\mathfrak{sta}(J_R) \simeq \mathfrak{so}(2)$ is one-dimensional and generated by the single element $i\sigma_y$, where σ_y is the second Pauli matrix. The orthogonal complement to $\mathfrak{sta}(J_R)$ within $\mathfrak{sp}(2, \mathbb{R})$ according to the standard choice of inner product (9.19) can be parametrized as

$$\mathfrak{sta}_\perp(J_R) = \left\{ r \begin{pmatrix} \cos \phi & \sin \phi \\ \sin \phi & -\cos \phi \end{pmatrix} \mid 0 \leq \phi < 2\pi, r \geq 0 \right\}. \quad (9.49)$$

The exponentiated version of these elements of the Lie algebra then provides us with elements of the symplectic group which can be recognized as general single-mode squeezings $S(r, \phi)$, where ϕ determines the orientation of the squeezing, and r quantifies its magnitude. The parameters r and ϕ serve as coordinates on the space of pure Gaussian states of a single bosonic mode with vanishing vector of means. The fact that the squeezing magnitude is notated as r is particularly convenient, as it is simple to verify—with the complexity geometry defined by the metric of subsection 9.3.1—that r corresponds precisely to the complexity of the state obtained by acting on the complex structure of the reference state with the squeezing operator $S(r, \phi)$ as $J = S(r, \phi)J_RS(r, \phi)^{-1}$. In other words, the squeezing magnitude r is also the radial coordinate characterizing state complexity according to the right-invariant metric defined by Equations (9.19)-(9.20) in the case of a single bosonic degree of freedom. That same metric restricted to the (r, ϕ) -plane gives us the line element

$$ds^2 = dr^2 + \cosh(2r) \sinh^2(r) d\phi^2. \quad (9.50)$$

In this case, a very natural addition that reflects some level of non-reversibility would be, for instance, to assign directions of increasing r a higher cost than those of decreasing r . As suggested in Section 9.2, this can be achieved by simply adding to the cost function the contribution from a vector potential of the form

$$A = -f(r, \phi) dr, \quad (9.51)$$

for some positive function $f(r, \phi)$. In the simplest case where f is only a function of r , the vector potential is just a total derivative, and can be written as $A = -dh$ for some $h(r)$ such that $f = dh/dr$. A vector potential of this form does not change what the optimal path is, since in this case $F_{ij} = 0$ and therefore the equation of motion (9.14) is the same as if there were no vector potential at all. It does however change the actual value of the complexity (9.10) by an additional term of $\Delta h = h(r_f) - h(0)$, with r_f being the squeezing parameter of the final state relative to the reference state. If the roles of reference and target are reversed, the contribution from this term changes to $-\Delta h$, and therefore some notion of non-reversibility is already present. If f also depends on ϕ , then on top of the asymmetry between reference and target states, we also get some nontrivial modification to the equation of motion (9.14), analogous to an external magnetic field orthogonal to the (r, ϕ) -plane.

As a final note, we remark that including a conformal factor of the form discussed in Subsection 9.3.3 in this case has a very intuitive interpretation as well: namely, it means that applying the same gate to the state of the system becomes harder the more squeezed the state already is (something that any quantum optician may confirm since it becomes extremely difficult to achieve squeezing beyond 15 dB in quantum optics [120, 186]).

Part IV

Conclusions

Chapter 10

Summary and outlook

This brings us to the end of the thesis. In this concluding chapter, we will recapitulate what we have learned in parts [I](#), [II](#), and [III](#), and comment on possible directions of future work.

Discussion of Part [I](#)

In part [I](#) of this thesis, we revisited and deepened a recently proposed connection [\[137\]](#) between localized quantum field theories and particle detector models in RQI. More specifically, we showed how to relate the dynamics of fully relativistic probe systems (here modeled by quantum fields that are confined by some external potential) to the dynamics of a finite number of probe degrees of freedom mimicking the UDW model, with harmonic oscillators coupled to the target field.

For a rather general class of initial states for the probe-field system, we showed that, at leading order in the coupling constant between probe and target field, the dynamics of any finite number of modes of the probe matches the dynamics of an equal number of harmonic-oscillator UDW detectors; in the latter case, the smearing functions for the detectors are directly related to the spatial profiles of the chosen modes of the probe field. At higher orders in the coupling, the corrections to the naive UDW model that arise from the existence of the additional modes can also be written analytically, as the Schwinger-Keldysh path integral over those modes can be computed in full detail. Therefore, the fully “nonperturbative” dynamics for any finite subset of modes could also be solved exactly, and the deviations from the simpler UDW model can be computed systematically to any order in perturbation theory.

The mathematical steps performed in Subsection 3.2.2 are independent of how we choose which modes of the field should be kept as UDW detectors. For the sake of concreteness, however, a physically motivated choice would be to only include those modes whose frequencies ω_n are below a certain cutoff, $\omega_n \leq \Omega$. This gives an effective-field-theory flavor to our derivation, with the procedure performed in Subsection 3.2.2 being very close in spirit to the Wilsonian approach to the renormalization group: in this case, the frequency Ω plays the role of a UV cutoff, and the action of the finite number of harmonic-oscillator UDW detectors coupled to the field, together with the additional factor from integrating out the rest of the modes, behaves as a low-energy effective action. This is in tune with the point of view that says particle detectors should be seen as effective descriptions of systems that are fundamentally relativistic, and provides a very simple toy example where the connection between this point of view and the more technical aspects of the formalism of effective field theories can be made explicit and concrete.

This comparison to effective field theories also sheds light on some of the locality issues of particle detectors which were mentioned in Subsection 2.2. From the literature on Wilsonian renormalization, we know that integrating out high energy modes can introduce mild non-localities in the effective low-energy action, where the typical scale of the non-localities is related to the inverse of the cutoff in energy or momentum [142, 37]. As we discussed in Subsection 2.2, coupling one single degree of freedom of a detector to a smeared (non-pointlike) field operator also generally leads to violations of locality and microcausality at distance scales that are controlled by the size of the spatial smearing. Obtaining the detector model from a field theory through coarse-graining-like steps as illustrated above draws a direct connection between these two statements, and also shows how to reconcile this (only apparent) conflict between smeared particle detector models and local field theory. From this point of view, the non-locality that emerges in the detector perspective is nothing but an artifact of the coarse-graining of the probe field to a finite number of modes; by including arbitrarily many modes, of arbitrarily high frequencies, one can restore a full-fledged local field theory.

This work was largely motivated, on the one hand, by the important role that particle detectors play in several aspects of the interface between quantum information and QFT, and on the other hand, by a recently renewed interest in a fully local and relativistic measurement framework for QFT. In this context, we believe that the connection between particle detector models and localized quantum fields introduced in [137], and reinforced here, is an important first step towards a more general understanding of the interplay between the more mathematically rigorous FV framework [72, 73] and the detector-based approach to RQI.

The present work also demonstrates the usefulness of path integral methods for the

formulation of nonperturbative statements relating localized quantum field theories and particle detectors. We see this as a powerful indication of the potential insights that effective field theory tools and concepts from the renormalization group can provide to particle detector models in RQI more generally. Work in connection to this has been initiated in [56], building on a path integral formulation for UDW detectors introduced in [36]. It is clear, however, that we have only begun to scratch the surface on this topic. It would be very interesting to understand how to make more systematic use of effective field theory methods in the context of particle detectors, as a way of bridging the gap between effective models of local probes in RQI and more fundamental physics.

Another promising avenue for future work is a more careful description of the physics behind the confining potential $U(\mathbf{x})$, and what new effects can be observed in our field-theoretic probe once we allow for that potential to have some nontrivial time dynamics. In this thesis, we simply treated $U(\mathbf{x})$ as a background structure whose spatiotemporal profile was unchanged throughout the entire evolution of the probe field. In practice, however, this effective potential is also sourced by physical degrees of freedom which can evolve nontrivially with time, and it is thus important to understand the consequences of this fact to the concrete implementation of localized quantum fields as probes. An example of a physical situation where these considerations can be valuable is when we try to describe the physics of a second-quantized atomic probe that is confined in a region of space by time-dependent external electromagnetic fields, or slightly more simplified versions thereof in which the essential physics can be accommodated into a single time-dependent scalar function $U(\mathbf{x})$ which builds up the potential over time. This will also have an effect on the internal dynamics of the degrees of freedom of the probe, since time-dependent driving forces will generally induce transitions and mixing between the different modes of the field in regimes far away from adiabaticity. One possible direction worth pursuing would be to the study of how a nontrivial temporal profile to $U(\mathbf{x})$ can be related with well-established phenomena in quantum field theory such as the dynamical Casimir effect [77, 63], as the time-dependent boundary conditions implemented by the moving mirrors are nothing but highly idealized versions of perfectly confining potentials. A first step towards understanding the effects of including time-dependent confining potentials is currently underway [146], and we expect that much more can come from these kinds of questions.

Discussion of Part II

In part II we addressed two problems: first, the problem of harvesting entanglement from a quantum field using localized systems that emerged from quantum fields themselves; and second, the problem of describing the optimal coupling between a set of local probes and a

quantum field, such that the entanglement acquired by the probes best reflects the original entanglement in the field’s initial state.

On entanglement harvesting, we began by addressing the observation in [154] that the reduced state of any local mode will be mixed due to the Reeh-Schlieder theorem. We provided a quantitative analysis of how the total mixedness of local modes of a free field depends on both the size and shape of the mode, and showed that, for any fixed size of the local mode and temperature of the field, it is always possible to make the local mode as pure as we like by changing its shape. Thus, while it is true that local modes are necessarily mixed, there is no nontrivial lower bound on the mixedness in terms of either the region’s size or the ambient field temperature.

We also explicitly found examples where modes of the two localized probe fields become entangled after interacting with a free Klein-Gordon field when the regions of interaction of each field are spacelike separated. This showcases a protocol for extracting spacelike entanglement from the vacuum of a free quantum field where not only the fields, but also the probes are fully relativistic. By leveraging the results from Part I, which showed that any finite number of modes of the field will behave at leading order in perturbation theory as a finite number of UDW detectors, our analysis shows that the entanglement that the fully relativistic probes can harvest is actually bounded from below by the results obtained with regular detectors. This further legitimizes the use of particle detector models to study this kind of protocol.

In the second half of Part II, we tackled the question of what kinds (if any) of couplings between detectors and field allow for a faithful transfer of field entanglement to the detectors so that the entanglement contained in the detectors accurately reflects the entanglement previously present in the field. In particular, we laid out a criterion to select, for any given subregion of a harmonic lattice, the set of modes that best encode the entanglement properties of that subregion with its complement. We then argued that, by engineering a suitable coupling between the given subregion and a set of probes, it is possible to completely transfer that entanglement to the probe degrees of freedom. This establishes the optimal form of coupling that a probe must have in order to most efficiently capture the entanglement between the coupling region and the rest of space. With this, we have the ingredients needed to substitute the problem of computing entanglement measures in a many-body system with the much simpler problem of evaluating entanglement in probes coupled to the system.

A natural (and more challenging) continuation of this work is to understand how to optimally extract entanglement between non-complementary subregions of space when one cannot probe all the degrees of freedom in each region. This is the case of most interest for

endeavors in entanglement harvesting, where one usually aims at capturing correlations between causally disconnected regions of spacetime. Prescribing the optimal profile of the coupling for entanglement extraction between disjoint regions poses new challenges: if two subregions A and B are not complementary, the joint state of AB will be mixed, and therefore the key techniques used in Section 6.3 to select a privileged set of modes supported in the regions of interest will no longer apply in general. A powerful result due to Klco, Beck, and Savage [110] explicitly shows that it is always possible, for any bipartite mixed Gaussian state, to select a basis of local modes in either side of the bipartition, such that the entanglement in the joint state factors as a tensor product of two-mode states (with one mode supported in each region) whose negativities match the negativity spectrum of the global state of the system. The construction of [110] reduces to our result in the case where the joint state is pure, and extends it in precisely the way we want in order to be able to address this same question for arbitrary subregions of a field theory. It would be very interesting to have a formal proof that this reorganization of entanglement in a general Gaussian state is really the optimal one.

Overall, our work provides a systematic study of what kind of interaction one should aim for in order to faithfully capture the entanglement originally present in a quantum field. The conceptual framework proposed here may provide a useful organizing principle for the understanding of entanglement in quantum field theory, as well as how it can be exchanged with localized probes.

Discussion of Part III

Here we proposed a geometric notion of circuit complexity for Gaussian states of both bosonic and fermionic systems that has a built-in way of accounting for 1) the presence of possible inhomogeneities (which change the cost of a given gate depending on what state it acts on), and 2) the presence of non-reversibilities (which assign different costs to paths related by time reversal). This allows us to build a geometric complexity measure that can capture features of the physical complexity of state preparation that one can experience in practice.

Instead of focusing on one particular choice of complexity measure, we kept as much generality as possible in the form of complexity functional based on Riemannian metrics on the group of Gaussian transformations. This allowed us to highlight a set of minimal assumptions that the Riemannian metric should satisfy in order to potentially be a good physical measure of how complex a given state is. With that, we were then able to generalize a previous characterization of state complexity considered in the literature [92, 49, 90] thanks to a generalized cylindrical foliation of the space of Gaussian transformations that

applies to any Riemannian metric satisfying the assumptions of Section 9.1. This shows how to preserve/adapt a lot of the techniques and results from [92, 49, 90], while at the same time being more flexible with the specifics of the complexity metric—which, among many other things, can naturally model the fact that in quantum optics, for instance, squeezing can be more or less costly depending on how squeezed the state already is.

We also proposed an extension to the framework that includes an additional term in the cost functional which makes directions on the space of Gaussian transformations that are connected by time-reversal have different costs. We consider this to be a well-motivated feature that a physical measure of circuit complexity (understood as a quantifier of the “hardness” of preparing a given state or unitary) may display: after all, from an experimentalist’s perspective, given two states $|\psi_1\rangle$ and $|\psi_2\rangle$, it is not uncommon to encounter situations where the experienced difficulty in preparing $|\psi_1\rangle$ starting from $|\psi_2\rangle$ is not the same as starting from $|\psi_1\rangle$ and then preparing $|\psi_2\rangle$. This is realized in the model proposed in Section 9.2 through the addition of a background 1-form on the space of Gaussian transformations which behaves like a background vector potential.

Despite the fact that Gaussian dynamics is restricted to systems with quadratic Hamiltonians, the reasoning employed here may also apply equally well to more general cases (potentially including interactions) as long as the generators of the elementary gate set form a Lie algebra. It would be particularly interesting to see how this plays out in the case of CFTs, especially due to its possible repercussions on considerations about holographic complexity. In particular, this general class of complexity functionals could potentially serve as a concrete boundary counterpart to the recent “Complexity = Anything” proposal for holographic complexity in the AdS/CFT correspondence [13, 14]. In this context, the infinite family of gravitational observables in asymptotically AdS spacetimes proposed as candidate duals to circuit complexity of the CFT state would be interpreted as arising from the infinite family of complexity geometries that can be assigned to the boundary theory.

An important problem to be pursued next is to see how to provide a class of measures of complexity in quantum field theory that is explicitly derived/motivated by constraints on the physical operations (and their associated costs) that can act on the field. This was our main motivation for starting this work, and we consider that our work here establishes the first step in that direction. More specifically, we aim to understand whether considerations about the action of local probes interacting with the field (i.e., particle detectors) can provide some notion of complexity in quantum field theory. Given the power of the ideas related to local measurements with particle detectors in the conception of operationally well-motivated notions of quantum information in QFT, it is natural to consider that similar considerations about local probes may also shed light on some notion of complexity in QFT, by providing some notion of complexity associated to a given family of gates that can be

implemented in the field through local probes. It should be interesting to know whether the features of the dynamics between fields and probes can yield some equation of motion that privileges a subset of the metrics among those satisfying the constraints of Section 9.1, as well for the background vector potential introduced in Section 9.2. Finally, it would also be important to understand how such motivated measures of complexity from local probes may be potentially related to other proposals for definitions of complexity from first principles. Examples include a notion of complexity defined in connection with resource theories [200]—which is also motivated by the need to accommodate constraints on what realistic agents can accomplish—and emergent concepts of complexity in the context of AdS/CFT [45, 43, 69, 48]. We also expect this to be subject of further study in the near future.

References

- [1] Gerardo Adesso and Fabrizio Illuminati. Gaussian measures of entanglement versus negativities: Ordering of two-mode gaussian states. *Phys. Rev. A*, 72:032334, Sep 2005.
- [2] Gerardo Adesso, Sammy Ragy, and Antony R. Lee. Continuous variable quantum information: Gaussian states and beyond. *Open Syst. Inf. Dyn.*, 21(01n02):1440001, 2014.
- [3] Peter Adshead, Richard Easter, and Eugene A. Lim. “In-in” formalism and cosmological perturbations. *Phys. Rev. D*, 80:083521, Oct 2009.
- [4] Cesar A. Agón, Matthew Headrick, and Brian Swingle. Subsystem complexity and holography. *J. High Energy Phys.*, 2019(2):145, Feb 2019.
- [5] Ivan Agullo, Béatrice Bonga, Eduardo Martín-Martínez, Sergi Nadal-Gisbert, T. Rick Perche, José Polo-Gómez, Patricia Ribes-Metidieri, and Bruno de S. L. Torres. Where is the entanglement in QFT?, In preparation.
- [6] Chris Akers, Netta Engelhardt, Daniel Harlow, Geoff Penington, and Shreya Vardhan. The black hole interior from non-isometric codes and complexity, 2022.
- [7] Tibra Ali, Arpan Bhattacharyya, S. Shajidul Haque, Eugene H. Kim, and Nathan Moynihan. Post-quench evolution of complexity and entanglement in a topological system. *Phys. Lett. B*, 811:135919, 2020.
- [8] Charis Anastopoulos, Bei-Lok Hu, and Konstantina Savvidou. Quantum field theory based quantum information: Measurements and correlations. *Ann. Phys.*, 450:169239, 2023.

- [9] Charis Anastopoulos, Bei-Lok Hu, and Konstantina Savvidou. Towards a field-theory based relativistic quantum information. *J. Phys. Conf. Ser.*, 2533(1):012004, jun 2023.
- [10] A. Ashtekar and Anne Magnon. Quantum fields in curved space-times. *Proc. R. Soc. Lond. A*, 346(1646):375–394, 1975.
- [11] K. Audenaert, J. Eisert, M. B. Plenio, and R. F. Werner. Entanglement properties of the harmonic chain. *Phys. Rev. A*, 66:042327, Oct 2002.
- [12] Luis C. Barbado, Esteban Castro-Ruiz, Luca Apadula, and Caslav Brukner. Unruh effect for detectors in superposition of accelerations. *Phys. Rev. D*, 102:045002, Aug 2020.
- [13] Alexandre Belin, Robert C. Myers, Shan-Ming Ruan, Gábor Sárosi, and Antony J. Speranza. Does complexity equal anything? *Phys. Rev. Lett.*, 128:081602, Feb 2022.
- [14] Alexandre Belin, Robert C. Myers, Shan-Ming Ruan, Gábor Sárosi, and Antony J. Speranza. Complexity equals anything II. *J. High Energy Phys.*, 2023(1):154, Jan 2023.
- [15] Yoni BenTov. Schwinger-Keldysh path integral for the quantum harmonic oscillator, 2021.
- [16] Rajendra Bhatia and Tanvi Jain. On symplectic eigenvalues of positive definite matrices. *J. Math. Phys.*, 56(11):112201, 2015.
- [17] Arpan Bhattacharyya, Arvind Shekar, and Aninda Sinha. Circuit complexity in interacting QFTs and RG flows. *J. High Energy Phys.*, 2018(10):140, Oct 2018.
- [18] Eugenio Bianchi and Robert C Myers. On the architecture of spacetime geometry. *Class. Quantum Gravity*, 31(21):214002, oct 2014.
- [19] Luca Bombelli, Rabinder K. Koul, Joohan Lee, and Rafael D. Sorkin. Quantum source of entropy for black holes. *Phys. Rev. D*, 34:373–383, Jul 1986.
- [20] L. Borsten, I. Jubb, and G. Kells. Impossible measurements revisited. *Phys. Rev. D*, 104:025012, Jul 2021.
- [21] Henning Bostelmann, Christopher J. Fewster, and Maximilian H. Ruep. Impossible measurements require impossible apparatus. *Phys. Rev. D*, 103:025017, Jan 2021.

- [22] Alonso Botero and Benni Reznik. Modewise entanglement of gaussian states. *Phys. Rev. A*, 67:052311, May 2003.
- [23] Alonso Botero and Benni Reznik. Spatial structures and localization of vacuum entanglement in the linear harmonic chain. *Phys. Rev. A*, 70:052329, Nov 2004.
- [24] Nora Brambilla, Gastão Krein, Jaume Tarrús Castellà, and Antonio Vairo. Born-oppenheimer approximation in an effective field theory language. *Phys. Rev. D*, 97:016016, Jan 2018.
- [25] Samuel L. Braunstein and Peter van Loock. Quantum information with continuous variables. *Rev. Mod. Phys.*, 77:513–577, Jun 2005.
- [26] Gavin K. Brennen, Peter Rohde, Barry C. Sanders, and Sukhwinder Singh. Multiscale quantum simulation of quantum field theory using wavelets. *Phys. Rev. A*, 92:032315, Sep 2015.
- [27] Adam R. Brown. A quantum complexity lower bound from differential geometry. *Nat. Phys.*, 19(3):401–406, Mar 2023.
- [28] Adam R. Brown, Daniel A. Roberts, Leonard Susskind, Brian Swingle, and Ying Zhao. Complexity, action, and black holes. *Phys. Rev. D*, 93:086006, Apr 2016.
- [29] Adam R. Brown, Daniel A. Roberts, Leonard Susskind, Brian Swingle, and Ying Zhao. Holographic complexity equals bulk action? *Phys. Rev. Lett.*, 116:191301, May 2016.
- [30] Adam R. Brown and Leonard Susskind. Second law of quantum complexity. *Phys. Rev. D*, 97:086015, Apr 2018.
- [31] Eric G. Brown, Eduardo Martín-Martínez, Nicolas C. Menicucci, and Robert B. Mann. Detectors for probing relativistic quantum physics beyond perturbation theory. *Phys. Rev. D*, 87:084062, Apr 2013.
- [32] David Edward Bruschi, Antony R Lee, and Ivette Fuentes. Time evolution techniques for detectors in relativistic quantum information. *J. Phys. A: Math. Theor.*, 46(16):165303, apr 2013.
- [33] Dagmar Bruß and Asher Peres. Construction of quantum states with bound entanglement. *Phys. Rev. A*, 61:030301, Feb 2000.

- [34] Fatih Bulut and W. N. Polyzou. Wavelets in field theory. *Phys. Rev. D*, 87:116011, Jun 2013.
- [35] I. M. Burbano and Francisco Calderón. Self-normalizing path integrals, 2021.
- [36] Ivan M. Burbano, T. Rick Perche, and Bruno de S. L. Torres. A path integral formulation for particle detectors: the Unruh-DeWitt model as a line defect. *J. High Energy Phys.*, 2021(3):76, Mar 2021.
- [37] Cliff P. Burgess. *Introduction to Effective Field Theory: Thinking Effectively about Hierarchies of Scale*. Cambridge University Press, 2020.
- [38] Elena Caceres, Shira Chapman, Josiah D. Couch, Juan P. Hernandez, Robert C. Myers, and Shan-Ming Ruan. Complexity of mixed states in QFT and holography. *J. High Energy Phys.*, 2020(3):12, Mar 2020.
- [39] Pasquale Calabrese and John Cardy. Entanglement entropy and quantum field theory. *J. Stat. Mech. Theor. Exp.*, 2004(06):P06002, jun 2004.
- [40] E. Calzetta and B. L. Hu. Closed-time-path functional formalism in curved spacetime: Application to cosmological back-reaction problems. *Phys. Rev. D*, 35:495–509, Jan 1987.
- [41] Esteban A. Calzetta and Bei-Lok Hu. *Nonequilibrium Quantum Field Theory*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 2008.
- [42] Hugo A. Camargo, Lucas Hackl, Michal P. Heller, Alexander Jahn, Tadashi Takayanagi, and Bennet Windt. Entanglement and complexity of purification in $(1 + 1)$ -dimensional free conformal field theories. *Phys. Rev. Res.*, 3:013248, Mar 2021.
- [43] Hugo A. Camargo, Michal P. Heller, Ro Jefferson, and Johannes Knaute. Path integral optimization as circuit complexity. *Phys. Rev. Lett.*, 123:011601, Jul 2019.
- [44] P. Candelas and D. W. Sciama. Irreversible thermodynamics of black holes. *Phys. Rev. Lett.*, 38:1372–1375, Jun 1977.
- [45] Pawel Caputa, Nilay Kundu, Masamichi Miyaji, Tadashi Takayanagi, and Kento Watanabe. Anti-de Sitter space from optimization of path integrals in conformal field theories. *Phys. Rev. Lett.*, 119:071602, Aug 2017.
- [46] Paweł Caputa and Javier M. Magan. Quantum computation as gravity. *Phys. Rev. Lett.*, 122:231302, Jun 2019.

- [47] Nicolas Chagnet, Shira Chapman, Jan de Boer, and Claire Zukowski. Complexity for conformal field theories in general dimensions. *Phys. Rev. Lett.*, 128:051601, Jan 2022.
- [48] Ammanamanchi Ramesh Chandra, Jan de Boer, Mario Flory, Michal P. Heller, Sergio Hörtner, and Andrew Rolph. Cost of holographic path integrals. *SciPost Phys.*, 14:061, 2023.
- [49] Shira Chapman, Jens Eisert, Lucas Hackl, Michal P. Heller, Ro Jefferson, Hugo Marrochio, and Robert C. Myers. Complexity and entanglement for thermofield double states. *SciPost Phys.*, 6:034, 2019.
- [50] Shira Chapman, Michal P. Heller, Hugo Marrochio, and Fernando Pastawski. Toward a definition of complexity for quantum field theory states. *Phys. Rev. Lett.*, 120:121602, Mar 2018.
- [51] Shira Chapman and Giuseppe Policastro. Quantum computational complexity from quantum information to black holes and back. *Eur. Phys. J. C*, 82(2):128, Feb 2022.
- [52] Richard Cleve. *An Introduction to Quantum Complexity Theory*, pages 103–127.
- [53] Josiah Couch, Willy Fischler, and Phuc H. Nguyen. Noether charge, black hole volume, and complexity. *J. High Energy Phys.*, 2017(3):119, Mar 2017.
- [54] Gabriel Cozzella, Stephen A. Fulling, André G. S. Landulfo, and George E. A. Matsas. Uniformly accelerated classical sources as limits of Unruh-DeWitt detectors. *Phys. Rev. D*, 102:105016, Nov 2020.
- [55] Luís C. B. Crispino, Atsushi Higuchi, and George E. A. Matsas. The Unruh effect and its applications. *Rev. Mod. Phys.*, 80:787–838, Jul 2008.
- [56] Níckolas de Aguiar Alves. Nonperturbative aspects of quantum field theory in curved spacetime, 2023.
- [57] José de Ramón, Maria Papageorgiou, and Eduardo Martín-Martínez. Relativistic causality in particle detector models: Faster-than-light signaling and impossible measurements. *Phys. Rev. D*, 103:085002, Apr 2021.
- [58] José de Ramón, Maria Papageorgiou, and Eduardo Martín-Martínez. Causality and signalling in noncompact detector-field interactions. *Phys. Rev. D*, 108:045015, Aug 2023.

- [59] Bruno de S. L. Torres, Kelly Wurtz, José Polo-Gómez, and Eduardo Martín-Martínez. Entanglement structure of quantum fields through local probes. *J. High Energy Phys.*, 2023(5):58, May 2023.
- [60] Walter Dean. Computational Complexity Theory. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2021 edition, 2021.
- [61] B. DeWitt. *General Relativity; an Einstein Centenary Survey*. Cambridge University Press, Cambridge, UK, 1980.
- [62] Giuseppe Di Giulio and Erik Tonni. Complexity of mixed gaussian states from fisher information geometry. *J. High Energy Phys.*, 2020(12):101, Dec 2020.
- [63] V V Dodonov. Current status of the dynamical casimir effect. *Physica Scripta*, 82(3):038105, aug 2010.
- [64] Fay Dowker. Useless qubits in “relativistic quantum information”, 2011.
- [65] H. F. Dowker and J. J. Halliwell. Quantum mechanics of history: The decoherence functional in quantum mechanics. *Phys. Rev. D*, 46:1580–1609, Aug 1992.
- [66] J. Eisert, M. Cramer, and M. B. Plenio. Colloquium: Area laws for the entanglement entropy. *Rev. Mod. Phys.*, 82:277–306, Feb 2010.
- [67] Artur Ekert and Richard Jozsa. Quantum algorithms: entanglement-enhanced information processing. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 356(1743):1769–1782, 1998.
- [68] Artur K. Ekert. Quantum cryptography based on Bell’s theorem. *Phys. Rev. Lett.*, 67:661–663, Aug 1991.
- [69] Johanna Erdmenger, Anna-Lena Weigel, Marius Gerbershagen, and Michal P. Heller. From complexity geometry to holographic spacetime. *Phys. Rev. D*, 108:106020, Nov 2023.
- [70] Thomas Faulkner, Monica Guica, Thomas Hartman, Robert C. Myers, and Mark Van Raamsdonk. Gravitation from entanglement in holographic CFTs. *J. High Energy Phys.*, 2014(3):51, Mar 2014.

- [71] Christopher J. Fewster, Ian Jubb, and Maximilian H. Rued. Asymptotic measurement schemes for every observable of a quantum field theory. *Ann. Henri Poincaré*, 24(4):1137–1184, Apr 2023.
- [72] Christopher J. Fewster and Rainer Verch. Quantum fields and local measurements. *Commun. Math. Phys.*, 378(2):851–889, Sep 2020.
- [73] Christopher J. Fewster and Rainer Verch. Measurement in quantum field theory, 2023.
- [74] R.P Feynman and F.L Vernon. The theory of a general quantum system interacting with a linear dissipative system. *Ann. Phys.*, 24:118–173, 1963.
- [75] Mario Flory and Michal P. Heller. Geometry of complexity in conformal field theory. *Phys. Rev. Res.*, 2:043438, Dec 2020.
- [76] Joshua Foo, Sho Onoe, and Magdalena Zych. Unruh-deWitt detectors in quantum superpositions of trajectories. *Phys. Rev. D*, 102:085013, Oct 2020.
- [77] S. A. Fulling, P. C. W. Davies, and Roger Penrose. Radiation from a moving mirror in two dimensional space-time: conformal anomaly. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 348(1654):393–414, 1976.
- [78] Evan P. G. Gale and Magdalena Zych. Relativistic Unruh-DeWitt detectors with quantized center of mass. *Phys. Rev. D*, 107:056023, Mar 2023.
- [79] Daniel J. George, Yuval R. Sanders, Mohsen Bagherimehrab, Barry C. Sanders, and Gavin K. Brennen. Entanglement in quantum field theory via wavelet representations. *Phys. Rev. D*, 106:036025, Aug 2022.
- [80] Flaminia Giacomini and Achim Kempf. Second-quantized Unruh-DeWitt detectors and their quantum reference frame transformations. *Phys. Rev. D*, 105:125001, Jun 2022.
- [81] Daniel Gottesman, Alexei Kitaev, and John Preskill. Encoding a qubit in an oscillator. *Phys. Rev. A*, 64:012310, Jun 2001.
- [82] Daniel Grimmer. A discrete analog of general covariance—part 1: Could the world be fundamentally set on a lattice?, 2022.
- [83] Daniel Grimmer. A discrete analog of general covariance—part 2: Despite what you’ve heard, a perfectly lorentzian lattice theory, 2022.

- [84] Daniel Grimmer. The pragmatic QFT measurement problem and the need for a Heisenberg-like cut in QFT. *Synthese*, 202(4):104, Sep 2023.
- [85] Daniel Grimmer, Irene Melgarejo-Lermas, José Polo-Gómez, and Eduardo Martín-Martínez. Decoding quantum field theory with machine learning. *J. High Energy Phys.*, 2023(8):31, aug 2023.
- [86] Daniel Grimmer, Bruno de S. L. Torres, and Eduardo Martín-Martínez. Measurements in QFT: Weakly coupled local particle detectors and entanglement harvesting. *Phys. Rev. D*, 104:085014, Oct 2021.
- [87] Arne L. Grimsmo and Shruti Puri. Quantum error correction with the Gottesman-Kitaev-Preskill code. *PRX Quantum*, 2:020101, Jun 2021.
- [88] Minyong Guo, Juan Hernandez, Robert C. Myers, and Shan-Ming Ruan. Circuit complexity for coherent states. *J. High Energy Phys.*, 2018(10):11, Oct 2018.
- [89] Leonid Gurvits. Classical deterministic complexity of Edmonds’ problem and quantum entanglement. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, STOC ’03, page 10–19, New York, NY, USA, 2003. Association for Computing Machinery.
- [90] Lucas Hackl and Eugenio Bianchi. Bosonic and fermionic Gaussian states from Kähler structures. *SciPost Phys. Core*, 4:025, 2021.
- [91] Lucas Hackl, Tommaso Guaita, Tao Shi, Jutho Haegeman, Eugene Demler, and J. Ignacio Cirac. Geometry of variational methods: dynamics of closed quantum systems. *SciPost Phys.*, 9:048, 2020.
- [92] Lucas Hackl and Robert C. Myers. Circuit complexity for free fermions. *J. High Energy Phys.*, 2018(7):139, Jul 2018.
- [93] Jonas Haferkamp, Philippe Faist, Naga B. T. Kothakonda, Jens Eisert, and Nicole Yunger Halpern. Linear growth of quantum circuit complexity. *Nat. Phys.*, 18(5):528–532, May 2022.
- [94] Daniel Harlow and Patrick Hayden. Quantum computation vs. firewalls. *J. High Energy Phys.*, 2013(6):85, Jun 2013.
- [95] Matthew Headrick. Lectures on entanglement entropy in field theory and holography, 2019.

- [96] Laura J. Henderson, Robie A. Hennigar, Robert B. Mann, Alexander R. H. Smith, and Jialin Zhang. Entangling detectors in anti-de sitter space. *J. High Energy Phys.*, 2019(5):178, May 2019.
- [97] Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki. Mixed-state entanglement and distillation: Is there a “bound” entanglement in nature? *Phys. Rev. Lett.*, 80:5239–5242, Jun 1998.
- [98] Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki. Separability of mixed states: necessary and sufficient conditions. *Phys. Lett. A*, 223(1):1–8, 1996.
- [99] Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. Quantum entanglement. *Rev. Mod. Phys.*, 81:865–942, Jun 2009.
- [100] M. Hotta, R. Schützhold, and W. G. Unruh. Partner particles for moving mirror radiation and black hole evaporation. *Phys. Rev. D*, 91:124060, Jun 2015.
- [101] B. L. Hu and A. Matacz. Quantum brownian motion in a bath of parametric oscillators: A model for system-field interactions. *Phys. Rev. D*, 49:6612–6635, Jun 1994.
- [102] H. Häffner, C.F. Roos, and R. Blatt. Quantum computing with trapped ions. *Phys. Rep.*, 469(4):155–203, 2008.
- [103] Wolfram Research, Inc. B spline basis, 2008. Champaign, IL, 2008.
- [104] Ted Jacobson. Entanglement equilibrium and the Einstein equation. *Phys. Rev. Lett.*, 116:201101, May 2016.
- [105] Robert A. Jefferson and Robert C. Myers. Circuit complexity in quantum field theory. *J. High Energy Phys.*, 2017(10):107, Oct 2017.
- [106] Abdul J. Jerri. *The Gibbs Phenomenon in Fourier Analysis, Splines and Wavelet Approximations*. Springer New York, 1st edition, 1998.
- [107] Kazuya Kaneko, Eiki Iyoda, and Takahiro Sagawa. Characterizing complexity of many-body quantum dynamics by higher-order eigenstate thermalization. *Phys. Rev. A*, 101:042126, Apr 2020.
- [108] Leonid Veniaminovic Keldysh. Diagram Technique for Nonequilibrium Processes. *Zh. Eksp. Teor. Fiz.*, 47(4):1515, 4 1965.

- [109] Isaac Kim, Eugene Tang, and John Preskill. The ghost in the radiation: robust encodings of the black hole interior. *J. High Energy Phys.*, 2020(6):31, Jun 2020.
- [110] Natalie Klco, D. H. Beck, and Martin J. Savage. Entanglement structures in quantum field theories: Negativity cores and bound entanglement in the vacuum. *Phys. Rev. A*, 107:012415, Jan 2023.
- [111] E. Knill, R. Laflamme, and G. J. Milburn. A scheme for efficient quantum computation with linear optics. *Nature*, 409(6816):46–52, Jan 2001.
- [112] Michael Krbek, Tomáš Tyc, and Jan Vlach. Inequalities for quantum marginal problems with continuous variables. *J. Math. Phys.*, 55(6):062201, 2014.
- [113] Nima Lashkari, Michael B. McDermott, and Mark Van Raamsdonk. Gravitational dynamics from entanglement “thermodynamics”. *J. High Energy Phys.*, 2014(4):195, Apr 2014.
- [114] Shih-Yuin Lin and B. L. Hu. Backreaction and the Unruh effect: New insights from exact solutions of uniformly accelerated detectors. *Phys. Rev. D*, 76:064008, Sep 2007.
- [115] Fangli Liu, Seth Whitsitt, Jonathan B. Curtis, Rex Lundgren, Paraj Titum, Zhi-Cheng Yang, James R. Garrison, and Alexey V. Gorshkov. Circuit complexity across a topological phase transition. *Phys. Rev. Res.*, 2:013323, Mar 2020.
- [116] Richard Lopp and Eduardo Martín-Martínez. Quantum delocalization, gauge, and quantum optics: Light-matter interaction in relativistic quantum information. *Phys. Rev. A*, 103:013703, Jan 2021.
- [117] Jorma Louko and Alejandro Satz. How often does the Unruh-DeWitt detector click? regularization by a spatial profile. *Class. Quantum Gravity*, 23(22):6321–6343, oct 2006.
- [118] Jorma Louko and Alejandro Satz. Transition rate of the Unruh–DeWitt detector in curved spacetime. *Class. Quantum Gravity*, 25(5):055012, feb 2008.
- [119] A. P. Lund, T. C. Ralph, and H. L. Haselgrove. Fault-tolerant linear optical quantum computing with small-amplitude coherent states. *Phys. Rev. Lett.*, 100:030503, Jan 2008.
- [120] A. Lvovsky. *Photonics Volume 1: Fundamentals of Photonics and Physics*. Wiley, Hoboken, 2015.

- [121] Eduardo Martín-Martínez and Nicolas C Menicucci. Entanglement in curved spacetimes and cosmology. *Class. Quantum Gravity*, 31(21):214001, oct 2014.
- [122] Eduardo Martín-Martínez, Miguel Montero, and Marco del Rey. Wavepacket detection with the Unruh-DeWitt model. *Phys. Rev. D*, 87:064038, Mar 2013.
- [123] Eduardo Martín-Martínez, T. Rick Perche, and Bruno de S. L. Torres. General relativistic quantum optics: Finite-size particle detector models in curved spacetimes. *Phys. Rev. D*, 101:045017, Feb 2020.
- [124] Eduardo Martín-Martínez, T. Rick Perche, and Bruno de S. L. Torres. Broken covariance of particle detector models in relativistic quantum information. *Phys. Rev. D*, 103:025007, Jan 2021.
- [125] Eduardo Martín-Martínez and Pablo Rodriguez-Lopez. Relativistic quantum optics: The relativistic invariance of the light-matter interaction models. *Phys. Rev. D*, 97:105026, May 2018.
- [126] Eduardo Martín-Martínez, Alexander R. H. Smith, and Daniel R. Terno. Spacetime structure and vacuum entanglement. *Phys. Rev. D*, 93:044001, Feb 2016.
- [127] Lluís Masanes, Augusto J. Roncaglia, and Antonio Acín. Complexity of energy eigenstates as a mechanism for equilibration. *Phys. Rev. E*, 87:032137, Mar 2013.
- [128] Jacob Miller and Akimasa Miyake. Latent computational complexity of symmetry-protected topological order with fractional symmetry. *Phys. Rev. Lett.*, 120:170503, Apr 2018.
- [129] Michael A. Nielsen. A geometric approach to quantum circuit lower bounds, 2005.
- [130] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.
- [131] Michael A. Nielsen, Mark R. Dowling, Mile Gu, and Andrew C. Doherty. Optimal control, geometry, and quantum computing. *Phys. Rev. A*, 73:062323, Jun 2006.
- [132] Michael A. Nielsen, Mark R. Dowling, Mile Gu, and Andrew C. Doherty. Quantum computation as geometry. *Science*, 311(5764):1133–1135, 2006.
- [133] Tatsuma Nishioka. Entanglement entropy: Holography and renormalization group. *Rev. Mod. Phys.*, 90:035007, Sep 2018.

- [134] Maria Papageorgiou and Doreen Fraser. Eliminating the ‘impossible’: Recent progress on local measurement theory for quantum field theory. *Found. Phys.*, 54(3):26, May 2024.
- [135] T. Rick Perche and Eduardo Martín-Martínez. Role of quantum degrees of freedom of relativistic fields in quantum information protocols. *Phys. Rev. A*, 107:042612, Apr 2023.
- [136] T. Rick Perche, José Polo-Gómez, Bruno de S. L. Torres, and Eduardo Martín-Martínez. Fully relativistic entanglement harvesting. *Phys. Rev. D*, 109:045018, Feb 2024.
- [137] T. Rick Perche, José Polo-Gómez, Bruno de S. L. Torres, and Eduardo Martín-Martínez. Particle detectors from localized quantum field theories. *Phys. Rev. D*, 109:045013, Feb 2024.
- [138] T. Rick Perche, Boris Ragula, and Eduardo Martín-Martínez. Harvesting entanglement from the gravitational vacuum. *Phys. Rev. D*, 108:085025, Oct 2023.
- [139] Asher Peres. Separability criterion for density matrices. *Phys. Rev. Lett.*, 77:1413–1415, Aug 1996.
- [140] Eric Poisson. The motion of point particles in curved spacetime. *Living Rev. Relativ.*, 7(1), May 2004.
- [141] Joseph Polchinski. *String Theory*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 1998.
- [142] Joseph Polchinski. *Effective Field Theory and the Fermi Surface*, 1999.
- [143] José Polo-Gómez, Luis J. Garay, and Eduardo Martín-Martínez. A detector-based measurement theory for quantum field theory. *Phys. Rev. D*, 105:065003, Mar 2022.
- [144] Alejandro Pozas-Kerstjens and Eduardo Martín-Martínez. Harvesting correlations from the quantum vacuum. *Phys. Rev. D*, 92:064042, Sep 2015.
- [145] Alejandro Pozas-Kerstjens and Eduardo Martín-Martínez. Entanglement harvesting from the electromagnetic vacuum with hydrogenlike atoms. *Phys. Rev. D*, 94:064074, Sep 2016.
- [146] Boris Ragula, Bruno de S. L. Torres, and Eduardo Martín-Martínez. Localized quantum fields with time-dependent potentials, In preparation.

- [147] Jørgen Rammer. *Quantum Field Theory of Non-equilibrium States*. Cambridge University Press, 2007.
- [148] Michael Reed and Barry Simon. *Methods of Modern Mathematical Physics, IV. Analysis of Operators*. Academic Press, 1978.
- [149] H. Reeh and S. Schlieder. Bemerkungen zur unitäräquivalenz von lorentzinvarianten feldern. *Nuovo Cim.*, 22:1051–1068, 1961.
- [150] Benni Reznik. Entanglement from the vacuum. *Found. Phys.*, 33(1):167–176, 2003.
- [151] Benni Reznik, Alex Retzker, and Jonathan Silman. Violating Bell’s inequalities in vacuum. *Phys. Rev. A*, 71(4):042104, 2005.
- [152] Daniel A. Roberts and Beni Yoshida. Chaos and complexity by design. *J. High Energy Phys.*, 2017(4):121, Apr 2017.
- [153] Shan-Ming Ruan. Purification complexity without purifications. *J. High Energy Phys.*, 2021(1):92, Jan 2021.
- [154] Maximilian H Ruep. Weakly coupled local particle detectors cannot harvest entanglement. *Class. Quantum Gravity*, 38(19):195029, sep 2021.
- [155] H. Rund. *The Differential Geometry of Finsler Spaces*. Springer Berlin, 1959.
- [156] Grant Salton, Robert B. Mann, and Nicolas C. Menicucci. Acceleration-assisted entanglement harvesting and ranging. *New J. Phys.*, 17(3):035001, 2015.
- [157] S. Schlieder. Einige bemerkungen zur zustandsänderung von relativistischen quantenmechanischen systemen durch messungen und zur lokalitätsforderung. *Commun. Math. Phys.*, 7(4):305–331, 1968.
- [158] Julian Schwinger. Brownian Motion of a Quantum Oscillator. *J. Math. Phys.*, 2(3):407–432, 12 1961.
- [159] Alessio Serafini, Gerardo Adesso, and Fabrizio Illuminati. Unitarily localizable entanglement of gaussian states. *Phys. Rev. A*, 71:032349, Mar 2005.
- [160] Petar Simidzija, Aida Ahmadzadegan, Achim Kempf, and Eduardo Martín-Martínez. Transmission of quantum information through quantum fields. *Phys. Rev. D*, 101:036014, Feb 2020.

- [161] Barry Simon. Schrödinger operators with purely discrete spectrum, 2008.
- [162] R. Simon. Peres-Horodecki separability criterion for continuous variable systems. *Phys. Rev. Lett.*, 84:2726–2729, Mar 2000.
- [163] Sergey N. Solodukhin. Entanglement entropy of black holes. *Living Rev. Relativ.*, 14(1):8, Oct 2011.
- [164] Jonathan Sorce. Notes on the type classification of von neumann algebras. *Rev. Math. Phys.*, 36(02):2430002, 2024.
- [165] Rafael D. Sorkin. Impossible measurements on quantum fields, 1993.
- [166] Mark Srednicki. Entropy and area. *Phys. Rev. Lett.*, 71:666–669, Aug 1993.
- [167] Douglas Stanford and Leonard Susskind. Complexity and shock wave geometries. *Phys. Rev. D*, 90:126007, Dec 2014.
- [168] Nadine Stritzelberger and Achim Kempf. Coherent delocalization in the light-matter interaction. *Phys. Rev. D*, 101:036007, Feb 2020.
- [169] Stephen J. Summers and Reinhard Werner. The vacuum violates Bell’s inequalities. *Phys. Lett., A*, 110(5):257–259, 1985.
- [170] Leonard Susskind. Addendum to computational complexity and black hole horizons. *Fortschr. Phys.*, 64(1):44–48, 2016.
- [171] Leonard Susskind. Computational complexity and black hole horizons. *Fortschr. Phys.*, 64(1):24–43, 2016.
- [172] Leonard Susskind. The typical-state paradox: diagnosing horizons with complexity. *Fortschr. Phys.*, 64(1):84–91, 2016.
- [173] Leonard Susskind. Black holes and complexity classes, 2018.
- [174] B M Terhal, J Conrad, and C Vuillot. Towards scalable bosonic quantum error correction. *Quantum Science and Technology*, 5(4):043001, jul 2020.
- [175] Erickson Tjoa and Robert B. Mann. Harvesting correlations in Schwarzschild and collapsing shell spacetimes. *J. High Energy Phys.*, 2020(8):155, Aug 2020.
- [176] Erickson Tjoa and Eduardo Martín-Martínez. When entanglement harvesting is not really harvesting. *Phys. Rev. D*, 104:125005, Dec 2021.

- [177] Bruno de S. L. Torres. Particle detector models from path integrals of localized quantum fields. *Phys. Rev. D*, 109:065004, Mar 2024.
- [178] Bruno de S. L. Torres and Eduardo Martín-Martínez. Toward a physically motivated notion of gaussian complexity geometry. *Phys. Rev. D*, 110:025006, Jul 2024.
- [179] Bruno de S. L. Torres, T. Rick Perche, André G. S. Landulfo, and George E. A. Matsas. Neutrino flavor oscillations without flavor states. *Phys. Rev. D*, 102:093003, Nov 2020.
- [180] Jose Trevison, Koji Yamaguchi, and Masahiro Hotta. Pure state entanglement harvesting in quantum field theory. *Prog. Theor. Exp. Phys.*, 2018(10), 10 2018. 103A03.
- [181] Jose Trevison, Koji Yamaguchi, and Masahiro Hotta. Spatially overlapped partners in quantum field theory. *J. Phys. A Math. Theor.*, 52(12):125402, feb 2019.
- [182] W. G. Unruh. Notes on black-hole evaporation. *Phys. Rev. D*, 14:870–892, Aug 1976.
- [183] W. G. Unruh and W. H. Zurek. Reduction of a wave packet in quantum brownian motion. *Phys. Rev. D*, 40:1071–1094, Aug 1989.
- [184] William G. Unruh and Robert M. Wald. What happens when an accelerating observer detects a Rindler particle. *Phys. Rev. D*, 29:1047–1056, Mar 1984.
- [185] R. Ursin, F. Tiefenbacher, T. Schmitt-Manderbach, H. Weier, T. Scheidl, M. Lindenthal, B. Blauensteiner, T. Jennewein, J. Perdigues, P. Trojek, B. Ömer, M. Fürst, M. Meyenburg, J. Rarity, Z. Sodnik, C. Barbieri, H. Weinfurter, and A. Zeilinger. Entanglement-based quantum communication over 144 km. *Nat. Phys.*, 3(7):481–486, Jul 2007.
- [186] Henning Vahlbruch, Moritz Mehmet, Karsten Danzmann, and Roman Schnabel. Detection of 15 db squeezed states of light and their application for the absolute calibration of photoelectric quantum efficiency. *Phys. Rev. Lett.*, 117:110801, Sep 2016.
- [187] Mark Van Raamsdonk. Building up space–time with quantum entanglement. *Int. J. Mod. Phys. D*, 19(14):2429–2435, 2010.
- [188] G. Vidal and R. F. Werner. Computable measure of entanglement. *Phys. Rev. A*, 65:032314, Feb 2002.

- [189] Silas Vriend, Daniel Grimmer, and Eduardo Martín-Martínez. The Unruh effect in slow motion. *Symmetry*, 13(11), 2021.
- [190] Robert Manuel Wald. *Quantum Field Theory in Curved Spacetime and Black Hole Thermodynamics*. The University of Chicago Press, 1994.
- [191] John Watrous. *Quantum Computational Complexity*, pages 7174–7201. Springer New York, New York, NY, 2009.
- [192] Christian Weedbrook, Stefano Pirandola, Raúl García-Patrón, Nicolas J. Cerf, Timothy C. Ralph, Jeffrey H. Shapiro, and Seth Lloyd. Gaussian quantum information. *Rev. Mod. Phys.*, 84:621–669, May 2012.
- [193] H. Arthur Weldon. Thermal green functions in coordinate space for massless particles of any spin. *Phys. Rev. D*, 62:056010, Aug 2000.
- [194] R. F. Werner and M. M. Wolf. Bound entangled gaussian states. *Phys. Rev. Lett.*, 86:3658–3661, Apr 2001.
- [195] John Williamson. On the algebraic problem concerning the normal forms of linear dynamical systems. *Am. J. Math.*, 58(1):141–163, 1936.
- [196] Bennet Windt, Alexander Jahn, Jens Eisert, and Lucas Hackl. Local optimization on pure Gaussian state manifolds. *SciPost Phys.*, 10:066, 2021.
- [197] Edward Witten. APS Medal for Exceptional Achievement in Research: Invited article on entanglement properties of quantum field theory. *Rev. Mod. Phys.*, 90:045003, Oct 2018.
- [198] Michael M. Wolf. Not-so-normal mode decomposition. *Phys. Rev. Lett.*, 100:070505, Feb 2008.
- [199] Run-Qiu Yang, Yu-Sen An, Chao Niu, Cheng-Yong Zhang, and Keun-Young Kim. What kind of “complexity” is dual to holographic complexity? *Eur. Phys. J. C*, 82(3):262, Mar 2022.
- [200] Nicole Yunger Halpern, Naga B. T. Kothakonda, Jonas Haferkamp, Anthony Munson, Jens Eisert, and Philippe Faist. Resource theory of quantum uncomplexity. *Phys. Rev. A*, 106:062417, Dec 2022.
- [201] Quntao Zhuang, Thomas Schuster, Beni Yoshida, and Norman Y. Yao. Scrambling and complexity in phase space. *Phys. Rev. A*, 99:062334, Jun 2019.

APPENDICES

Appendix A

Appendices for Chapter 3

A.1 Computation of the Jacobian determinant $\mathcal{N}$ in Eq. (3.34)

In Section 3.2, we argued that the expansion in Eq. (3.29) giving the localized quantum field ϕ_{D} in terms of a series of independent localized modes $\phi_{\mathbf{n}}$ can be seen as a change of variables in the path integral. This change of variables then induces a transformation of the measure, which can be written as

$$\mathcal{D}\phi_{\text{D}} = \mathcal{N} \prod_{\mathbf{n}} \mathcal{D}\phi_{\mathbf{n}} \quad (\text{A.1})$$

where, as mentioned right below Eq. (3.34), $\mathcal{N}$ is formally the Jacobian determinant of the change of variables from $\phi_{\text{D}}(\mathbf{x})$ to $\phi_{\mathbf{n}}(t)$.

Since the map between $\phi_{\text{D}}(\mathbf{x})$ and $\phi_{\mathbf{n}}(t)$ is linear, $\mathcal{N}$ does not depend on any field configurations, and therefore it changes the path integral at most by some overall constant. We can thus compute $\mathcal{N}$ by comparing the results for the same path integral evaluated in two different ways, which is the strategy that we will use below.

The simplest quantity that we can compute in order to infer the factor $\mathcal{N}$ is a transition amplitude between two field eigenstates in the absence of sources,

$$\langle \Phi_f | \hat{U}(t_f, t_i) | \Phi_i \rangle = \int \mathcal{D}\phi_{\text{D}} e^{iS_{\text{D}}[\phi_{\text{D}}]}, \quad (\text{A.2})$$

where the path integral is performed with the boundary conditions

$$\phi_D(t_i, \mathbf{x}) = \Phi_i(\mathbf{x}), \quad \phi_D(t_f, \mathbf{x}) = \Phi_f(\mathbf{x}). \quad (\text{A.3})$$

If we directly perform the substitution from $\phi_D(\mathbf{x})$ to $\phi_n(t)$ as a change of variables in (A.2), the transition amplitude becomes

$$\langle \Phi_f | \hat{U}(t_f, t_i) | \Phi_i \rangle = \mathcal{N} \prod_n \left(\int \mathcal{D}\phi_n e^{iS_n[\phi_n]} \right) \quad (\text{A.4})$$

where now the boundary conditions are

$$\phi_n(t_i) = \varphi_{n,i}, \quad \phi_n(t_f) = \varphi_{n,f} \quad (\text{A.5})$$

with $\varphi_{n,i}$ and $\varphi_{n,f}$ being the coefficients of the expansion of the probe field configurations $\Phi_i(\mathbf{x})$ and $\Phi_f(\mathbf{x})$ in the basis $\{v_n(\mathbf{x})\}$ at times t_i and t_f , respectively. This is of course a consequence of the fact that every well-defined field configuration at a given constant time slice has a unique expansion in terms of the basis $\{v_n(\mathbf{x})\}$.

On the other hand, the fact that every field configuration $\Phi(\mathbf{x})$ corresponds to a unique sequence of basis coefficients $\{\varphi_n\}$ also implies that we can express the kets $|\Phi_i\rangle$ and $|\Phi_f\rangle$ as

$$|\Phi_i\rangle = \bigotimes_n |\varphi_{n,i}\rangle, \quad (\text{A.6})$$

$$|\Phi_f\rangle = \bigotimes_n |\varphi_{n,f}\rangle, \quad (\text{A.7})$$

and since all of the modes of the field are decoupled from each other in the action (3.12), the time evolution operator $\hat{U}(t_f, t_i)$ can be factored as

$$\hat{U}(t_f, t_i) = \prod_n \hat{U}_n(t_f, t_i), \quad (\text{A.8})$$

where $\hat{U}_n(t_f, t_i)$ is the time evolution operator acting on the Hilbert space of each mode ϕ_n separately. Putting the results of Eqs. (A.6)-(A.8) together, we conclude that we can write

$$\langle \Phi_f | \hat{U}(t_f, t_i) | \Phi_i \rangle = \prod_n \langle \varphi_{n,f} | \hat{U}_n(t_f, t_i) | \varphi_{n,i} \rangle. \quad (\text{A.9})$$

But now it is clear that each factor in the product (A.9) is again expressible as a path integral in each separate mode,

$$\langle \varphi_{\mathbf{n},f} | \hat{U}_{\mathbf{n}}(t_f, t_i) | \varphi_{\mathbf{n},i} \rangle = \int \mathcal{D}\phi_{\mathbf{n}} e^{iS_{\mathbf{n}}[\phi_{\mathbf{n}}]}, \quad (\text{A.10})$$

with boundary conditions set again by (A.5). By substituting (A.10) back in Eq. (A.9) and comparing it with the expression (A.4) obtained by directly changing variables in the path integral, and given the fact that these transition amplitudes are not identically zero, it then follows that we must have $\mathcal{N} = 1$, as we wanted to show.

A.2 Normalization (3.10) and canonical commutation relations

Here we will verify that the normalization condition (3.10) between the mode functions $\{v_{\mathbf{n}}(\mathbf{x})\}$ ensures that $\phi_{\mathbf{n}}(t)$ and $\pi_{\mathbf{n}}(t) := d\phi_{\mathbf{n}}/dt$ satisfy the canonical commutation relations of position and momentum, and therefore define a genuine mode of the field. This serves as a brief consistency check between the general steps presented in Sec. 3.1 and the usual story of canonical quantization, and also provides useful insight into the generalization explored in Sec. 3.3 where the probe field can be localized in multiple spatial regions.

Given the foliation of spacetime $\mathcal{M}$ by spacelike Cauchy surfaces of constant Killing time t , the conjugate momentum $\pi_{\mathbf{D}}(t, \mathbf{x})$ to the field $\phi_{\mathbf{D}}(t, \mathbf{x})$ can be identified directly from the action (3.5) as

$$\pi_{\mathbf{D}}(t, \mathbf{x}) = \frac{\delta S_{\mathbf{D}}}{\delta(\partial_t \phi_{\mathbf{D}})} = \frac{\sqrt{h}}{N} \partial_t \phi_{\mathbf{D}}. \quad (\text{A.11})$$

By expanding $\phi_{\mathbf{D}}$ as in (3.11) and assuming that the mode functions $\{v_{\mathbf{n}}(\mathbf{x})\}$ satisfy the orthogonality condition (3.10), we can write

$$\phi_{\mathbf{n}}(t) = \int d^d \mathbf{x} \frac{\sqrt{h}}{N} v_{\mathbf{n}}(\mathbf{x}) \phi_{\mathbf{D}}(t, \mathbf{x}), \quad (\text{A.12})$$

$$\pi_{\mathbf{n}}(t) = \int d^d \mathbf{x} v_{\mathbf{n}}(\mathbf{x}) \pi_{\mathbf{D}}(t, \mathbf{x}). \quad (\text{A.13})$$

Then, when $\phi_{\mathbf{D}}(t, \mathbf{x})$ and $\pi_{\mathbf{D}}(t, \mathbf{x})$ are promoted to operators $\hat{\phi}_{\mathbf{D}}(t, \mathbf{x})$ and $\hat{\pi}_{\mathbf{D}}(t, \mathbf{x})$ in the usual process of canonical quantization, the equal-time commutation relations between

$\hat{\phi}_{\mathbf{n}}(t)$ and $\hat{\pi}_{\mathbf{n}}(t)$ become

$$\begin{aligned}
[\hat{\phi}_{\mathbf{n}}(t), \hat{\pi}_{\mathbf{n}'}(t)] &= \int d^d \mathbf{x} d^d \mathbf{x}' \frac{\sqrt{h(\mathbf{x})}}{N(\mathbf{x})} \overbrace{[\hat{\phi}_{\mathbf{D}}(t, \mathbf{x}), \hat{\pi}_{\mathbf{D}}(t, \mathbf{x}')] }^{=i\delta^{(d)}(\mathbf{x}, \mathbf{x}') \hat{\mathbb{1}}} \\
&\quad \times v_{\mathbf{n}}(\mathbf{x}) v_{\mathbf{n}'}(\mathbf{x}') \\
&= i \hat{\mathbb{1}} \underbrace{\int d^d \mathbf{x} d^d \mathbf{x}' \frac{\sqrt{h}}{N} v_{\mathbf{n}}(\mathbf{x}) v_{\mathbf{n}'}(\mathbf{x})}_{=\delta_{\mathbf{n}, \mathbf{n}'} \text{ assuming (3.10)}} \Rightarrow \\
[\hat{\phi}_{\mathbf{n}}(t), \hat{\pi}_{\mathbf{n}'}(t)] &= i \delta_{\mathbf{n}, \mathbf{n}'} \hat{\mathbb{1}}, \tag{A.14}
\end{aligned}$$

which are nothing but the canonical commutation relations between position and momentum. This shows that the orthogonality condition (3.10) implies that $(\hat{\phi}_{\mathbf{n}}(t), \hat{\pi}_{\mathbf{n}}(t))$ defines a canonically conjugate pair characterizing a single mode of the field.

It is also possible to prove the converse statement: namely, that requiring $\hat{\phi}_{\mathbf{n}}(t)$ and $\hat{\pi}_{\mathbf{n}}(t)$ to satisfy the canonical commutation relations of position and momentum (A.14) forces the mode functions $\{v_{\mathbf{n}}(\mathbf{x})\}$ to satisfy the normalization condition (3.10). To see this, we expand $\hat{\phi}_{\mathbf{D}}$ and its conjugate momentum again as in (3.11),

$$\hat{\phi}_{\mathbf{D}}(t, \mathbf{x}) = \sum_{\mathbf{n}} \hat{\phi}_{\mathbf{n}}(t) v_{\mathbf{n}}(\mathbf{x}), \tag{A.15}$$

$$\hat{\pi}_{\mathbf{D}}(t, \mathbf{x}) = \frac{\sqrt{h}}{N} \sum_{\mathbf{n}} \hat{\pi}_{\mathbf{n}}(t) v_{\mathbf{n}}(\mathbf{x}). \tag{A.16}$$

Now, using the fact that $\hat{\phi}_{\mathbf{D}}(t, \mathbf{x})$ and $\hat{\pi}_{\mathbf{D}}(t, \mathbf{x})$ must always satisfy the field's canonical commutation relations and assuming that Eq. (A.14) holds, we get that

$$\begin{aligned}
[\hat{\phi}_{\mathbf{D}}(t, \mathbf{x}), \hat{\pi}_{\mathbf{D}}(t, \mathbf{x}')] &= i \frac{\sqrt{h(\mathbf{x}')}}{N(\mathbf{x}')} \sum_{\mathbf{n}} v_{\mathbf{n}}(\mathbf{x}) v_{\mathbf{n}}(\mathbf{x}') \hat{\mathbb{1}} \\
&= i \delta^{(d)}(\mathbf{x} - \mathbf{x}') \hat{\mathbb{1}}, \tag{A.17}
\end{aligned}$$

which implies that

$$\frac{\sqrt{h(\mathbf{x}')}}{N(\mathbf{x}')} \sum_{\mathbf{n}} v_{\mathbf{n}}(\mathbf{x}) v_{\mathbf{n}}(\mathbf{x}') = \delta^{(d)}(\mathbf{x} - \mathbf{x}'). \tag{A.18}$$

But note that we can always express

$$v_{\mathbf{n}}(\mathbf{x}) = \int d^d \mathbf{x}' \delta^{(d)}(\mathbf{x} - \mathbf{x}') v_{\mathbf{n}}(\mathbf{x}') \quad (\text{A.19})$$

which, using the fact that the Dirac delta can be written as in Eq. (A.18), becomes

$$v_{\mathbf{n}}(\mathbf{x}) = \sum_{\mathbf{n}'} v_{\mathbf{n}'}(\mathbf{x}) \int d^d \mathbf{x}' \frac{\sqrt{h(\mathbf{x}')}}{N(\mathbf{x}')} v_{\mathbf{n}'}(\mathbf{x}') v_{\mathbf{n}}(\mathbf{x}'). \quad (\text{A.20})$$

Now we note that Eq. (A.20) provides an expression for the mode function $v_{\mathbf{n}}(\mathbf{x})$ as a linear combination of the other mode functions $v_{\mathbf{n}'}(\mathbf{x})$. But if $\{v_{\mathbf{n}}(\mathbf{x})\}$ forms a basis, the only possible linear combination that expresses $v_{\mathbf{n}}(\mathbf{x})$ in terms of the set $\{v_{\mathbf{n}'}(\mathbf{x})\}$ is the trivial one, where the expansion coefficients are equal to 1 when $\mathbf{n} = \mathbf{n}'$, and zero otherwise. In short,

$$\int d^d \mathbf{x}' \frac{\sqrt{h(\mathbf{x}')}}{N(\mathbf{x}')} v_{\mathbf{n}'}(\mathbf{x}') v_{\mathbf{n}}(\mathbf{x}') = \delta_{\mathbf{n}', \mathbf{n}} \quad (\text{A.21})$$

which is precisely (3.10).

This explicitly shows that the set of modes constructed in Sec. 3.1 matches precisely the set of normal modes that we would have constructed in the process of canonical quantization of the field $\phi_{\text{b}}(\mathbf{x})$. Conversely, it also shows that if the no-overlap condition (3.59) is not satisfied, the amplitudes associated to each mode function $v_{\mathbf{n}}(\mathbf{x})$ cannot be interpreted as independent harmonic oscillators. One can interpret this as the version in canonical quantization version of the statement alluded to in Subsec. 3.3.1, in which the approximation (3.59) was crucial for justifying the statement that a localized quantum field theory where the confining potential had multiple minima could be seen as composed of two sets of independent modes supported around each of these minima.

Appendix B

Appendix for Chapter 6

B.1 Proof of Equation (6.30) and interlacing theorem for symplectic eigenvalues

The symplectic eigenvalues of the covariance matrix $\tilde{\sigma}_{AB}$ corresponding to the partial transpose of a bipartite system AB have a particularly simple expression when the full state ρ_{AB} is pure. Let us first consider that A and B are made of one single mode. Up to local symplectic transformations, the covariance matrix of subsystem A can be expressed as

$$\sigma_A = \begin{pmatrix} \nu & 0 \\ 0 & \nu \end{pmatrix}. \quad (\text{B.1})$$

One purification of this covariance matrix is given by

$$\sigma_{AP} = \begin{pmatrix} \nu & 0 & \sqrt{\nu^2 - 1} & 0 \\ 0 & \nu & 0 & -\sqrt{\nu^2 - 1} \\ \sqrt{\nu^2 - 1} & 0 & \nu & 0 \\ 0 & -\sqrt{\nu^2 - 1} & 0 & \nu \end{pmatrix}, \quad (\text{B.2})$$

and any other purification can be obtained from σ_{AP} by a local symplectic transformation acting solely on the purifying system P . Since local unitaries on P do not alter the entanglement between A and its purifying system, we can use the covariance matrix σ_{AP} to compute the entanglement between A and B . By taking the partial transpose on P , we

obtain the covariance matrix

$$\tilde{\sigma}_{AP} = \begin{pmatrix} \nu & 0 & \sqrt{\nu^2 - 1} & 0 \\ 0 & \nu & 0 & \sqrt{\nu^2 - 1} \\ \sqrt{\nu^2 - 1} & 0 & \nu & 0 \\ 0 & \sqrt{\nu^2 - 1} & 0 & \nu \end{pmatrix}. \quad (\text{B.3})$$

The smallest symplectic eigenvalue of $\tilde{\sigma}_{AP}$ (which is the only one that can be smaller than 1, and is therefore the only one that can contribute to the negativity) is then found to be

$$\tilde{\nu}_- = \nu - \sqrt{\nu^2 - 1}. \quad (\text{B.4})$$

We see that any $\nu > 1$ will lead to $\tilde{\nu}_- < 1$. Moreover, higher values of ν correspond to lower values of $\tilde{\nu}_-$, and therefore higher values of the logarithmic negativity. This is just a manifestation of the fact that, for an overall pure state, more mixedness in the reduced state on one subsystem signals more entanglement across the bipartition.

This logic can be immediately generalized to bipartitions with arbitrary numbers of modes. Consider now that A and B have N_A and N_B modes respectively, where we can assume $N_A \leq N_B$ without loss of generality. Due to Williamson's theorem, there is always a basis of normal modes for region A such that the covariance matrix takes the form

$$\sigma_A = \bigoplus_{i=1}^{N_A} \begin{pmatrix} \nu_i & 0 \\ 0 & \nu_i \end{pmatrix}. \quad (\text{B.5})$$

Once again, knowing that the overall covariance matrix σ_{AB} is pure and using the purification (B.2) mode by mode, we find that one can pick a basis of modes in the complement of A such that the full covariance matrix of AB can be expressed as

$$\sigma_{AB} = \sigma_{AP} \oplus \mathbb{1}_{2\Delta} \quad (\text{B.6})$$

where we now have

$$\sigma_{AP} = \bigoplus_{i=1}^{N_A} \begin{pmatrix} \nu_i & 0 & \sqrt{\nu_i^2 - 1} & 0 \\ 0 & \nu_i & 0 & -\sqrt{\nu_i^2 - 1} \\ \sqrt{\nu_i^2 - 1} & 0 & \nu_i & 0 \\ 0 & -\sqrt{\nu_i^2 - 1} & 0 & \nu_i \end{pmatrix} \quad (\text{B.7})$$

and $\Delta = N_B - N_A$.

In general, any pair of modes that purify each other, such that the two modes together are in a product state with the rest of the system, comprise what is known in the literature as a *partner mode pair* [100, 180, 181]. We can thus see the covariance matrix σ_{AP} as the covariance matrix of the system expressed in the basis of normal modes of A and their respective partners. If we now trace out the subsystem A , the covariance matrix of the remaining subsystem B becomes

$$\sigma_B = \left[\bigoplus_{i=1}^{N_A} \begin{pmatrix} \nu_i & 0 \\ 0 & \nu_i \end{pmatrix} \right] \oplus \mathbb{1}_{2\Delta}. \quad (\text{B.8})$$

This shows that the partner mode to the j -th most mixed normal mode of A is given by the j -th most mixed normal mode of B , whenever A and B are complementary subsystems in an overall pure Gaussian state.

Now, taking the partial transpose over B gives us the covariance matrix

$$\tilde{\sigma}_{AB} = \tilde{\sigma}_{AP} \oplus \mathbb{1}_{2\Delta} \quad (\text{B.9})$$

where $\tilde{\sigma}_{AP}$ can be written as a direct sum partially-transposed two-mode covariance matrices of the form (B.3). This allows us to directly infer that the symplectic eigenvalues of $\tilde{\sigma}_{AB}$ that are smaller than 1 are given by

$$\tilde{\nu}_i = \nu_i - \sqrt{\nu_i^2 - 1}. \quad (\text{B.10})$$

Because the right-hand side is a decreasing function of ν_i , if the symplectic eigenvalues ν_i of σ_A are ordered from largest to smallest, then the symplectic eigenvalues $\tilde{\nu}_i$ of $\tilde{\sigma}_{AB}$ as given above will be ordered from smallest to largest, with the smallest ones contributing more to the logarithmic negativity as given by Eq. (6.29).

Finally, it is important to mention one basic property of the symplectic spectrum of submatrices of a given covariance matrix obtained by throwing away part of its degrees of freedom. This will come in handy when going through the proof of one of our claims in Section 6.3. The result is the following:

Theorem B.1.1 (Interlacing theorem for symplectic eigenvalues [112, 16]). *Let M_1 be a $2N \times 2N$ positive definite symmetric matrix, which we can see as a covariance matrix for a system with N modes. Consider the matrix M_2 obtained from M_1 by deleting the rows and columns corresponding to any given mode of M_1 . Let $\{\nu_j(M_1)\}_{j=1}^N$ and $\{\nu_j(M_2)\}_{j=1}^{N-1}$ be the list of symplectic eigenvalues of M_1 and M_2 respectively, both ordered from smallest to*

largest. It then follows that

$$\nu_j(M_1) \leq \nu_j(M_2) \leq \nu_{j+2}(M_1), \quad 1 \leq j \leq N-2 \quad (\text{B.11})$$

as well as

$$\nu_{N-1}(M_1) \leq \nu_{N-1}(M_2). \quad (\text{B.12})$$

The proof of the theorem (which can be found explicitly e.g. in [112, 16]) essentially follows from Cauchy's interlacing theorem, plus the observation that the symplectic eigenvalues of a covariance matrix σ are just the positive half of the spectrum of the phase space operator $i\sigma\Omega$.

One very simple and intuitive consequence of the interlacing theorem for symplectic eigenvalues is that the mixedness of any single mode in a given Gaussian state cannot be made lower than that of the least mixed normal mode of the region where the mode is supported. In Section 6.3 we show that a similar application of the theorem can also bound the entanglement between any set of modes within a region and its complement, again in terms of the entanglement with the region's normal modes.